Robots in the Wild

Report IV: Robots in Dynamic Environments

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

Introduction

This Report IV focuses on 2 different cases that deal with large robots in dynamic environments. These cases involve a maritime school that has worked with simulations and autonomous ships, as well as large agricultural robots. Both areas, the agricultural field and especially the sea, are very dynamic environments, and here it is not only humans that the robots have to relate to. The two different fieldworks have primarily consisted of interviews with relevant people and observation of simulations on sites. There have therefore not, as in Report III, been actual observations of the working robots that the report covers, as this has not been possible. Our fieldworks have primarily been carried out through observations and semi-structured interviews, most of them on site, some at conferences and some of which were conducted online.

If Report I shows how autonomy is imagined, and Report II shows how autonomy is challenged and collapses in public spaces, and Report III shows how it is organisationally compensated for in controlled environments, Report IV shows the limits of compensation itself. In dynamic environments such as fields and seas, neither public scripts nor organisational routines are sufficient to stabilise robotic autonomy.

A dynamic environment is in our definition a non-repeatable environment with a continuous change involving an array of non-human actors (weather, animals, terrain). It is an environment that comes with a high cost of failure because of the difficulty of making predictions and the delayed feedback (errors reveal themselves too late).

We begin with a smooth transition from cobots in Report III to the larger agricultural robots and the issues they introduce. It will then account for the central discussions in the development of autonomous ships and the sea as a dynamic environment. Finally, all the fieldworks will be linked together and positioned in the context of the future of autonomous robots in dynamic environments.

Our study covers fieldworks where there has not always been a fully developed robot to study. The fully autonomous ship or agricultural robot does not yet exist, but it is the development of them, and the surrounding discussions, we will be writing about in this

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

report. The fieldwork from the navigation school ShipGO will be used to discuss these debates and bring the experiences of experienced captains into the discussion about an autonomous ship. First, however, the report will begin by describing the fieldwork at FarmGO, which produces autonomous agricultural machines, and examine how having a robot drive around in a field already leads to some very different issues than on a factory floor or in a restaurant.

FarmGO - Robots in the field

FarmGO is a startup company whose mission is to make autonomous agricultural machines safe. In the company they offer three different packages – Integration, Detection and Safety – where each package is a step closer towards autonomous agricultural machines. They themselves describe their mission as making food production more sustainable and secure. They aim to achieve this by enabling farmers to make better use of agricultural robots and to provide a safety system for robot manufacturers, replacing the need for farmers to constantly monitor working field robots.

We first found the company FarmGO during our search for field sites, later we saw one of their robots, involved in a research project at a conference (Report I).

We contacted them via the email provided on their website, and from there set up an interview with one of the co-founders, Henry. It is through Henry's experience in the industry that he has picked up on the need for this form of robot technology. He explained to us, that the self-driving machines that exist in the industry today need to become more autonomous and better able to handle the unpredictable. Henry explains that at conferences this autonomy "stood out through its absence." He believes that it is a "worldwide phenomenon" that farmers need better machines as a relief, as they cannot get the amount of labour they need. He explains that the reason for this is rooted in the growing world population as well as an increasing demand for plant-based foods. This spiral naturally leads to a growing need for labour in plant-based food production, and with increasing regulations against pesticides, this naturally places higher demands on farmers and the agricultural sector.

In addition to our interview with Henry, we also had FarmGO's robot demonstrated at a conference, where we also met and interviewed Andrew (See Report I). He was working on a research project related to agricultural machines (including machines that involved FarmGO). When we met him, he was standing outside at a stand next to their very large agricultural robot. We did a short semi-structured interview with him, and this will be included when relevant.

Job functions and the people at the site

The company, FarmGO, consists mainly of engineers, where one of the founders is also an engineer, and the other, Henry, is an economist and farmer. They have not spoken with any farmers in the development of their software, because farmers are not their customers. It is primarily manufacturers of agricultural machines and system integrators who make up the company's customers, and it is therefore also them they communicate with, even though in the end it will be farmers who use the products.

According to Henry, it is crucial to talk precisely with the manufacturers, as they need to understand the "brain" of the machines in order to implement the software correctly and, for example, make a manned tractor autonomous. At the same time, it is central to live up to the applicable ISO standards and the Machinery Directive, since the product is being developed for commercialisation. According to Henry, farmers have limited knowledge of these standards and regulations and therefore cannot contribute significantly to the development process.

In order to meet the requirements, it is also necessary that the manufacturers (e.g. of robotic tractors) that FarmGO collaborates with also follow the relevant standards. FarmGO develops software for robots, and not the robots themselves. For that reason, it is particularly important for them to secure that they have the right customers among machine manufacturers. In addition to standards and regulations, they are also dependent on knowing the various machines' system architecture and interfaces, so that their technology at the software level can be compatible.

Since their software technology is still in the startup phase, they carry out tests in different places around the world. At the moment they have eight systems in operation, both with manufacturers and with end users such as farmers in different parts of the world. So far they have received positive feedback from customers, but the whole project is still very new, and it is therefore too early to say anything definitive about the technology. This is, however, a deliberate choice on FarmGO's part, as they want to be involved early in the integration phase, where the challenges that arise can actively contribute to the further development of their technology.

In contrast to Henry, Andrew has focused primarily on collaborating with farmers rather than manufacturers. This is because they are mainly interested in how to make agricultural robots safe, and therefore they judged that it was most important to ask farmers what the biggest concerns and challenges they encounter are. Here they found that one of the most important things is which kind of sensors you use, since fields are a much more dynamic environment than normal industrial settings, and therefore there is a need for more advanced and detailed sensors than, for example, in a warehouse.

Technical description of the robot

FarmGO describes their goal as enabling safe autonomy, and for this they have the following three "products": Integrate, Detect and Safe. The three products are to be understood as packages that companies can buy and that correspond to how far the individual companies have come in their robot development.

The Integrate package is the first step towards making agricultural machines autonomous, which among other things involves placing cameras and sensors on the machines. The next step in development is Detect, where one carries out "prototyping" in order to teach the robot its environment and identify potential challenges (this will be elaborated later in the report). The last package, Safe, involves ensuring that the robot lives up to various safety standards, as well as FarmGO's own defined safety performance levels.

They want to create a technology that will make self-driving agricultural machines more self-driving and autonomous. Henry himself describes that their software aim to make the machines capable of handling the unpredictable, which currently involves reacting to unforeseen events. Not only must the robot learn to react to this unpredictable event, it must also send the information about the event and the reaction on to a farmer, who can supervise the machine's movements. The idea is therefore not that their software should go in and replace the farmer, but that the farmer should instead be able to control the machine remotely.

This is to free up more time for the farmer, for today, as Henry explains, the self-driving agricultural machines must be followed "like a dog on a leash." Instead, they want to free the machine from the farmer and the farmer from the machine, so that it can simply be remotely controlled with a tablet.

At the robot conference, the following agricultural machine was on display as an example of robots in the field:

The machine is about 3 x 3 metres. On top it has solar panels. The wheels sit on a metal frame and are designed to drive on the uneven soil in fields. The red cord functions as a safety measure that gives the robot a signal if there is a large object in front of it that it cannot drive over. It is equipped with several different kinds of sensors that have different functions. Underneath it, sprayers with weed spray and crop spreaders are mounted. It drives around in fields and sows crops both on organic fields and with pesticide, where, according to Andrew, it uses 95% less pesticide because it sprays individual weeds instead of the entire field. It can do this because it is equipped with cameras that can spot weeds. When it drives on organic fields, it does not spray.

Unlike many cobots, the machine is not equipped with LiDAR alone. While LiDAR is effective in controlled environments, it cannot distinguish between different types of objects and therefore registers a tall grain stalk, a stone, or a human in the same way. In agricultural settings, where the environment is highly dynamic, this limitation necessitates more advanced and context-sensitive sensor systems.

A dynamic environment

Henry explains that the core product at FarmGO is their object recognition software, but that they still run into anomalies that arise from defective sensors or reduced visibility in the fields.

They have a major challenge with the environment their software will work in, as nature is a dynamic environment that is in constant change and where there can be many unforeseen challenges. Henry describes the robots as "dumb as a doornail," because they are completely dependent on the software and the right algorithm functioning optimally. He emphasises that it is crucial to minimise the number of false positives, as these make the robot stop repeatedly and thereby reduce efficiency. This far from the media robots discussed in Report I and the discussion around R2D2. It shows how real robots are often very dumb compared to the fictional robots like R2D2 from Star Wars.

In real life robots must prove themselves. If the robot is not efficient, it is a bad investment, and the idea of freeing the farmer from the machine disappears, since he must constantly go out to check why the robot has stopped. It is therefore crucial for the software that these anomalies are handled correctly, and that the robot/machine is able to filter between the different anomalies' degree of importance. For that reason, they also cannot just use LiDAR, which registers that there is "something" in front of the machine – it needs to know what it is so it can be handled correctly.

To achieve this, Henry explains that the better you know your environment, the better the conditions you have for distinguishing between false and real positives. Despite his experience in the agricultural sector, however, the environment is still extremely complex, precisely because it is so dynamic, where it can, for example, be difficult to distinguish clearly between individual blades of grass.

Henry himself compares the agricultural environment with an industrial environment and points out that the latter potentially contains fewer dangerous obstacles because it is a less dynamic environment. In the examples mentioned in Report III, it is much easier to control the environment. Often it is also humans who are the most dynamic element in, for example, a warehouse, whereas in a field you also have to relate to nature, which is incredibly dynamic.

The dynamic nature

As mentioned at the beginning, one of the central issues in this report is robots in very dynamic environments. Agricultural machines are an example of robots that must work in such a dynamic environment.

Even on a conventional farm, nature still cannot be controlled. There are variables in a field that cannot be known and accounted for. In addition, agricultural machines are often larger than the robots we have presented in the other Reports (II, III), and in the next part of this report we will be working at an even larger scale (a whole ship as a robot) – both in terms of what is to be made autonomous, and in terms of the environment the robots are in.

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 - FarmGO

At FarmGO they have defined two kinds of attention zones, where one is a warning zone and the other is a hazard zone. The difference between these zones is the robot's speed when an object is registered in the zones. Henry did not specify what distance there is from the warning zone to the hazard zone, but based on his explanation we assume that this depends on which machine you are dealing with and which environment you are in.

The hazard zone is the first safety zone the machine activates when it approaches an object, and here the speed is reduced significantly. Henry explained that if the machine, for example, registers something at a distance of 15 metres, it will reduce its speed from 10 km/h to 3 km/h. In addition to reducing its speed, the machine sends a message to a tablet, which is monitored by a farmer, where it simply informs about the situation. If it continues to register something in the hazard zone, it maintains the speed, and if the object moves away from the hazard zone, the machine speeds up again.

If the object comes closer, it moves into the warning zone, which means that the machine stops completely. When something is registered in the warning zone, the farmer must check whether what is in front of the robot is moving, or whether it is an object that is standing still, and if it is, they must remove the stationary object or in some other way help the machine around it. This of course is complicated if the object is an animal pretending to be an unmovable object.

This means that the farmer has to stay attentive to the tablet even as they attend to other tasks. None of the two safety zones are carried out without human assessment, but they can, however, be carried out without the human needing to be physically present. The machine is therefore self-adaptive, but not fully autonomous.

In the interview with Andrew, we gained an insight into which objects may be involved in the above situation. He explained that farmers are most concerned about fawns (baby deers). In the above situation of a machine encountering an unmovable object, it could be a stone or a tyre. This is something farmers have to be aware of, as these hard objects can damage the agricultural machines, but the object could also be a small animal. This is also where the importance of not only using LiDAR comes in. Fawns do not move. Their survival technique consists in lying completely still and trying to resemble a stone, so that predators do not discover them. Therefore, they also will not move even if a huge agricultural machine is driving towards them. This means that several thousand fawns are killed by agricultural machines every year.

The farmer needs to know what it is he has to deal with when his agricultural machine encounters an immovable object in its path. Andrew told us that farmers therefore had equipped some machines with an infrared camera that could detect temperature differences and thus register whether the object was a fawn or just a stone.

In the risk assessment of their technology, FarmGO has taken as its starting point the ISO standards and the Machinery Directive, and on that basis, they have interpreted what this requires of their technology in order to live up to it. In the standards there is already an assessment of which risks are placed in which categories, and here they insert their technology and assess what significance these risks have for it, or vice versa. Henry also explains that since they are still in the process of developing their product, they do not yet have a finished risk assessment.

Living up to the risk assessments is not only important for securing certifications or for ensuring the best performance of the robots. Henry explains that it is important to have the right documentation, because if something goes wrong, you must be able to show that you have considered possible accidents and carried out the necessary risk mitigation. If you can document that you have done everything that was possible, an occupational safety authority cannot regard an accident as negligent.

FarmGO's technology is not yet able to distinguish between a human and a tree. But according to Andrew this is not something that would be difficult to introduce to them. Henry explained in connection with this that it is not part of the standards that it must be able to do that, since the standards do not include AI and machine learning, and therefore they have not begun working with it yet.

The research project Andrew is working on has a different focus than the standards and has therefore made other safety assessments, which is why they are looking into developing agricultural machines that can distinguish between trees and humans. The standards do not function only as a checklist for safety and an approval stamp but are also used actively to assess where developers should prioritise their efforts.

FarmGO is a commercial company that has to sell its product, and Andrew is working on a research project that examines safety in relation to agricultural machines. Henry said directly that they are waiting for it to become part of the standards so they can begin developing it.

This reflects an innovation logic in which development is based not only on technological possibilities or ambitions, but on external standards and regulations. The standards become a kind of roadmap for developers and help define which areas are considered necessary to spend time on. At the same time, henrys statement points to a tension between safety and regulation, because even though human recognition could obviously contribute to increased safety; development only begins when the standards make explicit demands for it. Henry's remark thus supports a picture of standards as a strategic guideline that reduces risk and ensures market relevance, but which can also slow innovation and leave a temporary gap between technological potential and actually implemented safety.

Trust - FarmGO

Henry explained that their solution arises from a need to free farmers from constantly having to walk next to the robot and supervise its work. He explains that today they have what is called "within line of sight," which he describes as the farmer having to walk beside the robot as with a dog on a leash. He believes this makes no sense, because no farmers want to do that or have time for it.

For these self-driving machines to have any value, they must therefore be released from this physical supervision. This does, however, require a good deal of trust from the farmer, since instead of being able to see everything, the machine does at all times, they are removed from the equation and can only follow it remotely on a tablet.

This trust in the machine is one of the things Andrew examined and sought to improve, which is precisely why they asked farmers what is most important to them. Here they found that one of the most important things was that the machines are able to spot fawns to ensure they don't kill them.

These self-driving and autonomous machines also reflect a radical shift in the agricultural sector. Previously the farmers themselves sat in the machines and carried out the work, whereas humans they became an integrated part and extension of the agricultural machine working in the field. It shifts the attention of the farmer from being in the field and paying attention to the machine in the environments to a different bodied position attending to both the tablet and some other tasks at hand (multi-tasking).

Henry told us that the feedback they have received from farmers so far is predominantly positive towards being removed from the field, as they, according to him, have long since come to terms with the fact that these robots must be able to drive themselves and be autonomous. He actually believes that farmers expect the machines to be able to drive independently, and that being present in physical supervision is therefore experienced as an unnecessary burden they would like to get rid of.

At the same time, Henry emphasises that in the end it is operational reliability that is decisive for farmers, because without it the machines give them no real value. So even though the autonomous and self-driving aspects are well received, it is never at the expense of operational reliability for them. The machines are also not cheap, and Henry explains that the farmers' reluctance is also connected to the fact that when they have to invest several million, "then it bloody well has to work." There is thus an important element of "return on investment" for the farmers.

Overall, Henry believes that what is decisive for successful implementation is the farmers' mindset in relation to technology. Farmers cannot simply treat the machine as a new employee to whom they can explain what to do and what to pay attention to. The machine must be carefully set up anew each time it is introduced to a new task. This can be considerably more time-consuming, as it requires them to plan and think differently about their work.

The interface and the system itself must therefore be user-friendly and intuitive, so that they can easily figure out how to use it in their otherwise busy everyday lives. Henry also explains that there are differences between farmers in terms of how receptive they are to new technology. Even though many of them are very conservative, there are also quite a few who think it is exciting that something new is happening technologically.

In general, Henry also argues that farmers are actually quite technically skilled, as they work daily with technically complicated machines. In order to make the interface user-friendly and intuitive, it would therefore be obvious to make it comparable to the systems and user interfaces they are already using now.

Henry also believes that those farmers who find new technologies too complicated or who oppose development are in reality speaking out of ignorance about what the technologies actually involve. Henry believes that they have built up a futuristic and unrealistic idea in their heads about what it means that new technologies are emerging in the agricultural sector (See Report I).

According to Henry, however, the vast majority still expect this technological upheaval, as they are aware of the major developments taking place in the field and therefore know that it is only a matter of time. This is in line with a text by Linda Reissig and Michael Siegrist from 2025 and one by Théo Martin et al. from 2022, which examine farmers' relationship to digitalisation/robots and how it can affect their trust.

The first text, by Linda Reissig and Michael Siegrist (2025), examines Swiss farmers' attitudes to digitalisation and future agricultural technologies with special focus on virtual fencing (infrared beams) in livestock production and fully autonomous hoeing robots in arable farming. The study, which is based on a survey of 939 farmers, shows that acceptance of new technologies depends on a number of factors such as age, level of education, digital skills, technology affinity, economic resources and farm size. For virtual fences, the farm's labour and resources play a central role, just as farmers' perception of risks and benefits is decisive.

For the autonomous hoeing robots, education, general attitudes towards digitalisation and the assessment of risks and benefits are of greater importance. In general, the study points out that farmers' acceptance of digital technologies is specific to each individual technology and largely depends on a weighing of benefit and risk.

Farmers' trust in these new forms of robots is closely linked to their perception of the risks and benefits of the new technologies. The study emphasises that virtual fences, meaning fences that exists for the robots via their internal maps and not fences that are visible in the real world, are only accepted if farmers experience them as reliable and ethically acceptable, especially in relation to animal welfare. For autonomous hoeing robots, it is particularly farmers' education, digital competence and experience with technology that strengthen trust.

Overall, the study concludes that it is the weighing of risks and benefits that has greatest importance for farmers' acceptance of the technology. If the technology is experienced as more risky than beneficial, this affects farmers' trust in the technology and thus, ultimately, their acceptance.

If, on the other hand, the technology is experienced as very beneficial with only minor risks, these risks are seen as a necessary evil. However, there is also the element of farmers' own digital skills and experience, where more knowledge and technological familiarity create greater trust, while lack of knowledge and low digital competence create scepticism.

The second text by Théo Martin et al. takes a broader perspective and shows how robots change the meaning of work and relations, as farmers' reduction of physical work and supervision requires an increase in trust in the machines and alarms. This form of dependence can potentially create uncertainty and "technostress" if farmers do not trust enough in the reliability of the system.

The article shows that the relationship between humans and animals changes when technology takes over parts of the work, which also affects the farmer's sense of autonomy and control over the farm. Martin et al. call for this loss of autonomy to be examined further, as it, according to them, has major consequences for the farmer and their acceptance of technology.

Trust in the robots therefore is not only about the technology as such, but also about whether farmers can trust that the robots support their professional identity and values in their work. They conclude that robots not only change the content of work, but also the social relations and identities in agriculture, and that future research should include these social and cultural dimensions in addition to the technical and economic aspects.

In summary, agriculture is an area where there is great potential for using autonomous robots, and where a lot of development and research is already taking place on how the implementation of robots can work, and how farmers relate to this development.

Anomaly - FarmGO

Henry explains that the core product of FarmGO is their object recognition software, and that the most significant types of anomalies therefore arise from defective sensors or reduced visibility. However, they face a major challenge in the environment in which their software is expected to operate, as nature is a dynamic environment that is constantly changing and presents many unforeseen challenges. Henry describes the robots as "as dumb as a doornail," because they are entirely dependent on the software and the correct algorithm functioning optimally. He emphasizes that it is crucial to minimize the number of false positives, as these cause the robot to stop repeatedly and thereby reduce its efficiency.

If the robot is not efficient, it becomes a poor investment, and the idea of separating the farmer from the machine disappears, since he constantly has to go out and check why it has stopped.

It is therefore crucial for the software that these anomalies are handled, and that the robot is able to distinguish between different anomalies based on their level of importance.

To achieve this, Henry explains that the better one understands the environment, the better the conditions for distinguishing between false and true positives. Despite his experience in the agricultural sector, the environment remains extremely complex precisely because it is so dynamic. It can be difficult to clearly distinguish between individual blades of grass. Here, it is again noteworthy that they have not spoken with any farmers, whose experience and extensive domain knowledge could have contributed to identifying and distinguishing between different types of anomalies. Henry also compared the agricultural environment to an industrial setting and pointed out that the latter potentially contains fewer dangerous obstacles because it is a less dynamic environment.

Sim2Real - FarmGO

As mentioned earlier, FarmGO divide their software into different packages, where the Detect package includes the process of "prototyping." Henry explained that at the start of a collaboration with a new customer, they receive drawings from them, which they use to develop a digital solution.

The digital solution is a simulation of the machine they are working with and the environment the machine is to drive in, which they use to assess where the sensors should be placed on the machine to work optimally in relation to their technology. After this they order a "sensor package" which they test on the machine in question and then help to integrate with the customers, so that the machine can go out into the field and drive in order to collect data.

"Prototyping" is about the machine gaining a sense of the environment it drives in, and about FarmGO gaining a sense of potential blind spots and challenges.

The optimal thing for their technology is therefore to start with simulations in order to assess sensor placement and then to send the technology out into its actual environment and test it. This is because the environment they work in is very dynamic and therefore difficult to generalise.

On the basis of this, Henry explains that their simulations are not 100% comparable with reality. That is also not the point, as a complete reproduction would be far too complicated and in practice impossible. He mentions, for example, that in a vineyard the vines can be anything from 1 metre to half a metre high, which makes it difficult to draw a realistic environment for a machine in the simulation program. The simulation process is therefore iterative and based on trial and error, where errors that appear in reality are used to adjust the simulations.

Even though the simulations are not completely accurate, they save considerable time in the start-up phase and contribute to developing the technology on the best possible foundation. Henry believes that simulation is a valuable tool for everyone, and as he says, even if it is not 100% it is better to have 80% than 0%.

This also speaks to one of the central issues in Sub-Report IV: autonomous robots in a dynamic environment. The more dynamic an environment is, the harder it is to implement robots in it. In Sub-Report I it is described how some robot developers have problems with implementing robots in hospitals, as they can have different beds, which makes it difficult because robots ideally have to be able to be standardised.

Fields, however, are practically impossible to standardise, which only makes it harder to develop autonomous robots for them.

Sabotage and Tinkering - FarmGO

Since we didn't conduct any observations of actual agricultural robots, we haven't observed any sabotage, and we didn't discuss it with Henry or Andrew in detail. Therefore the following section will be mostly theoretical and based on other empirical data, namely Linda Reissig and Michael Siegrist's article (2025), discovering farmers in Switzerlands opinions on digitalisation and future agricultural technologies, and Theo Martin's et al. (2022) literature review of 90 different studies on how robots change and affect work in agriculture.

In Reissig and Siegrist's text, they emphasize that the acceptance of new technologies largely depends on farmers' digital competencies, education, and perceptions of risks and benefits (Reissig & Siegrist 2025). In this context, sabotage may arise if the farmer or their employees perceive the technology as more risky than beneficial. Regarding the farmer's own digital competencies, our empirical data has repeatedly shown that employees who do not understand various technologies completely refrain from using them. Although this is not active sabotage, the rejection of the technology is also detrimental to its development and integration into work practices.

Martin et al. demonstrate that robots change both work organization and the meaning of work. With automated milking systems, the physical workload is reduced, but new tasks emerge, such as monitoring alarms and managing data, which can lead to "technostress" and increased mental strain. When farmers and employees experience that work becomes more surveillance-oriented and less craft-based, it can create a sense of loss of identity and autonomy. In such a context, sabotage can be seen as a protest against the changed conditions, and here too we have empirical evidence of employees sabotaging robots on this basis.

In a broader perspective, robots in agriculture also raise questions about ownership, control, and power distribution. Martin et al. point out that robots can reinforce division of labor and power relations on the farm, for example by shifting responsibility from employees to the owner, who must handle the robots' alarms (Martin et al. 2022). This form of redistribution may be perceived as unfair and thereby create grounds for resistance. Sabotage can therefore be understood as a social reaction to technological changes that are experienced as imposed from outside or as threats to established practices. Overall, this shows that sabotage functions as an indicator of the social tensions that may accompany the implementation of robots in agriculture.

As shown, agricultural robots work in a dynamic environment that is constantly changing, and therefore the robots need to be able to adapt as well. Henry also compared them to robots working in factories, and how agricultural robots will encounter more dangerous and complex situations than industrial robots. The next part of this report will be describing robots working in an even more dynamic environment and in a larger scale, namely autonomous robots on the sea. Therefore, the next part of the report will also include a sensory description of the sea as an environment.

ShipGO – the Sea as Dynamic Environment

ShipGO is a European navigation school where students and course participants receive instruction in seafaring. Part of the school buildings used to be a machine factory, where, among other things, engines for ships were produced. Teaching takes place both via simulations and more practical experience. The following part of the report covers a field visit, during which we visited the school and its training sites. The fieldwork consisted of guided tours of their campus and semi-structured interviews with two of the instructors. During the tour, we observed a simulation session and the training sites.

ShipGO was visited because they both have simulators that they use to teach their students shipping/navigation, and in addition they are relevant in relation to one of our project's use-cases, which investigates the possibility of developing an autonomous ship. We interviewed two of the instructors at the school about autonomous ships, and whether, in their view, this is a possibility in the near future, as well as what it would require to develop an autonomous ship.

Over the course of the Sub Report, we will include relevant discussions from an internal fieldwork in the project that concerns the development of autonomous ships.

The dynamic sea

We will first give a description of the sea as an environment, as this will help create a better understanding of the issues we address later in this Sub-Report.

In the course of the fieldworks, we have been out to many different locations to study robots. Everything from restaurants, conference halls, warehouses and fields has been examined. Something that recurs is how little variation it takes to make the implementation of a robot difficult. The difference lies in how easy it is to reduce the variations, an example of this is the example of a hospital looking to implement cleaning robots, and how difficult this was, because the hospital has 12 different kinds of beds, and that is simply too much variation for a robot to deal with. Everything needs to be standardised.

When you are in a warehouse, it is easier to see the places where a cobot will encounter problems, and then do something about it. The employees who work in the warehouse can be trained on how to interact with the robots. In a restaurant the unpredictability of the humans, becomes the biggest challenge, while the environment remains (somewhat) static. When you let robots loose in nature, such as a field, the dynamic environment introduces more challenges, since it is inherently unpredictable and vast. Out at sea, it only becomes more complicated.

In addition to a robot having to be able to navigate a changing environment that cannot always be predicted, it also has to be able to manoeuvre around other ships, such as sailing vessels and tankers, and here it must be able to follow the rules of the sea. In addition, it must also be aware of squalls, whales and heavy waves. The sea is not forgiving. It is a physically hard environment that takes its toll on the technology and, as we will also show later in the report, one of the crew's biggest tasks is maintenance.

Oversold technology

Curt also says that, in his view, many of these new technologies are oversold, just as we have experienced at robot conferences. When the technologies then come out to sea on board a ship, it turns out that they are simply not good enough yet.

Curt explains this both by saying that the technologies have to be sold to shipyards, who will not buy them if they cannot promise a lot of improvements, and because the engineers and "technical people," in his view, do not understand the sea:

"They don't understand well enough the environment they're trying to sell technology into when we're talking about shipping." (Interview with Curt)

The difference between testing technologies in a controlled environment on land and when they then come out at sea is so great that it is difficult to develop technology that has been tested well enough and that can handle the unpredictable sea. Therefore, it often ends up that a given technology has to go through two or three generations/versions before it starts to work reasonably reliably.

Field visit at the school

The field visit consisted, as mentioned, of a tour where we were shown around the whole of ShipGO and, among other things, went in to see a simulation in which some students had to navigate around in a specific body of water, where one of the ships controlled by the instructor started to drift off course and the students then had to react. In this concrete example it ended with one of the students colliding with the instructor's ship, because they did not manage to change course or reduce speed in time.

It is primarily men who work at the school, but of the two instructors we had longer interviews with, one was a man and one a woman, named Curt and Hanna. Both come from a maritime background. Both were trained at ShipGO and both have sailed for many years on several different types of ships.

They are both instructors at the school and therefore often interact with the simulators and the remote-controlled ships that the students themselves have to build.

The simulators that the students use when they are learning to sail are to function as a simulation of real situations that the students could encounter at sea in reality. There can be several students sailing in the same "simulation" at once, which means they have to communicate with each other. The instructors are often part of the simulations, controlling other ships, which makes it possible for them to expose the students to various situations where they suddenly have to react quickly, or where their attention skills are tested.

The simulator is to resemble a real ship's bridge as much as possible. It has six screens that function as the windows on the bridge. In addition, there are six smaller screens under the window screens, which show sea charts and other data about the ship itself and the waters it is sailing in. Beneath them, on desks, the simulators have various buttons, controls and of course a wheel to steer the ship. All of this helps create as realistic a simulation of reality as possible, so that the students learn as much as they can. The sessions they have in the simulators can also last several hours, so that they can learn what it is like to sail. Furthermore, the students also sail small model boats in one of the buildings on campus, where they get to experience sailing in real seawater, just scaled down. Here they get a sense of reality that gives a different bodily experience than the simulations do.

The students learn to combine simulations with real world experiences, when they sail the small boats. They need to improve their small boats, and they have to pass certain tests. In this they also learn about financing and gain valuable experience for when they have to sail with real sized boats.

The robot/simulator

ShipGO does not have a robot in the same way as we have been out to observe elsewhere, but instead they have a lot to say about both the use of simulations in general and whether it would in fact be possible to have autonomous ships. As a case, it also lies very close to a use-case on the development of autonomous ships. In addition, it is also interesting because many of the robots in our project are used in controlled environments, whereas a case like these deals with how robots would function in very dynamic environments.

As Curt describes the problem with robots at sea versus in a warehouse building:

"If it's out at sea, the level of noise and vibrations and moisture and all that kind of stuff is just significantly higher, and there's just nothing you can do about it." (Interview with Curt)

The sea is outside our control, and this makes it much harder to actually develop robots that can navigate in it, as there are so many factors that humans simply cannot control. Even in a small body of water such as, for example, the North Sea, it is impossible to control, and it can also be difficult to predict.

"The North Sea is terrible, because, well, the geological conditions mean that even though it isn't big, it isn't an ocean where the waves can really build up enormously, the geological conditions mean that you get these irregular, hard and high seas." (Interview with Hanna)

Hanna then described a concrete experience she had in the North Sea, where the ship she was on, had to lie still for 12 hours because the waves were so strong that they could not sail one way or the other. This is also a situation that a potential autonomous ship could encounter.

It is not just about being able to sail the right way at the right speed or being able to navigate around other ships. It is also about being able to judge what is best to do in an unforeseen situation on the very unpredictable sea.

It is not enough to have a ship that can hold its course. Curt believes we already have the technology for that. He specifically describes that there is, as such, no need for a crew when the ship is sailing on the open sea. He even suggested that it could be a possibility to have a crew sail the ship out of the harbour and onto the open sea, where they then would disembark and sail back to shore, while the ship continued towards its destination, for example from Japan to the USA across the Atlantic. When the ship then approached the USA, a new crew could board and sail the ship into port.

This of course leads to other issues in the form of responsibility and monitoring of the ship's cargo, as very few companies would be comfortable with there being no one watching over their valuable cargo while the ship is sailing on the open sea. In addition, there is a problem concerning who bears responsibility if the ship is involved in an accident.

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 - ShipGO

The following section will deal with safety on board an autonomous ship – both in relation to the autonomous ship itself and the various risks involving other ships.

First and foremost, it is important to mention here that a fully autonomous ship would not be allowed to sail legally today. This is both because, according to current maritime law, there must be two independent methods by which a ship can determine its position (Interview with Hanna), and because there must be a captain who is responsible. Maritime law simply does not include autonomous ships in its legislation.

The requirement for the independent position-finding methods is so that if some kind of failure occurs in the ship's navigation, the crew can use another method to determine their location and course. The two methods must not be connected to the same network and must be independent of each other.

One example Hanna had was the system called ECDIS (Electronic Chart Display and Information System), which modern ships use to navigate and plan routes with, where the instructors sometimes program errors into the students' ECDIS in the simulators so that they learn that the systems do not always work.

"But it shows something else, that they're somewhere else. So all of a sudden they don't know where they are. And then they have to get the old methods going and find something, and look out the window." – Interview with Hanna

This is based on the fact that when you are sailing around on the open sea, there is not always anything that indicates where you are. If ECDIS then fails, which can happen, and the ship therefore thinks it is somewhere other than where it actually is, the crew must use other methods that are not based on ECDIS and its data to determine their position so that they can correct their course, etc.

If you therefore have an autonomous ship, its location must be determinable by an independent method, and according to Hanna it must be possible to control it from someone on land – or on the ship, for that matter. There must be someone who can always say exactly where the ship is, and it is not enough to just trust a GPS.

If the ship had only a GPS and nothing else, and that failed for one reason or another, you would only discover that something was wrong when the ship ran aground or similar. It is therefore extremely important that ships have more than one method that can determine their location.

In addition, there is an issue of hacking an autonomous ship. If it can be hacked and remotely controlled to sail somewhere other than its planned course, it could potentially cause great damage to other ships, ports, drilling platforms, etc. This is not something we will go much further into.

Another problem is how to ensure that the ship avoids colliding with other ships. The two instructors specifically brought up a thought experiment they called "the 3-ship situation." It is a situation Curt usually asks developers of anti-collision technologies how their software would solve. It cannot be solved with a clear rule for what to do but must be solved on the basis of Rule 2 of the COLREGs (International Regulations for Preventing Collisions at Sea), which says that one must always show good seamanship and make a correct decision. (Interview with Curt).

Normally the rules of the sea are fairly deterministic: if A, then do B, etc., but because not all situations you can encounter at sea are deterministic, the rules also have to reflect that.

The situation is that you have three ships, where Ship A is on a collision course with Ship B. According to the rules, Ship A must turn away and avoid collision with Ship B, but this will mean that the ship is now on a collision course with Ship C. Ship C must therefore now, according to the rules, turn away from Ship A and into a collision course with Ship B, and so on and so forth.

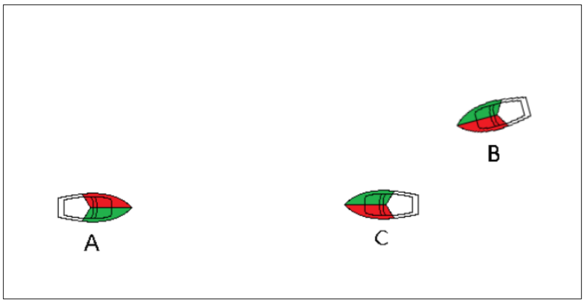

It is a situation that shows the complexities of sailing, as there are rules for which ships must turn and in which direction. If an autonomous ship is programmed to follow the rules, it would not be able to solve the 3-ship situation, because no matter what it did, it would break the rules.

If it were three autonomous ships, it is not hard to imagine that they would "burn out" in the attempt to follow the rules of the sea while at the same time solving the problem. It is a situation where the only right thing to do is to break the rules:

"So either you break the rules with regard to him and sail into him. Or you break the rules with regard to him when you turn towards him. Which, of course, is the right thing to do. Then the rules of the sea say that now he has a ship – now he has that white one there – a ship on his starboard side that he has to give way to. So he now has to turn over here." – Interview with Curt

This is an example where the otherwise very deterministic rules of the sea, where there are concrete solutions for most situations, cannot be followed. This is also why there is a rule in the COLREGs that one must always show good seamanship.

However, this rule is immediately far more difficult to have programmed into an autonomous system. It should also be mentioned that the rules apparently exist to be able to assign blame. In principle, most situations at sea should be manageable with good seamanship, but if something goes wrong, someone has to be "punished" for it.

It is also why it will always be the captain who, in the end, is the one who bears responsibility. The rule on good seamanship means that in all situations someone can be found to carry the responsibility. This problem of responsibility is something we encountered in an interview with an experienced captain called Frank. He was quite sceptical of new technologies, since they in his view often don't make life easier for the crews on ships. Furthermore, he outrighted rejected the idea of autonomous ships, and exclaimed they would never exist, since no software developer would want to take responsibility if an autonomous ship crashes. In his view, it is too convenient to have a captain who legally is the responsible one, if anything goes wrong, and if you remove the captain you have to find someone new to take the responsibility if something goes wrong, and he doesn't believe there exists a single software company that would want to do that.

He argued for this by describing how, in reality the captains often have to answer to people on land, and that he has often experienced how higher ups in the company he works at, will tell him what to do, and what routes to take and so on, but when it comes to taking responsibility for any delays or accidents, they are nowhere to be found, and it always falls on the captain.

Another interesting discussion in relation to autonomous ships is maintenance. Many of the tasks that the crew perform on a modern ship are maintenance.

This maintenance is, however, primarily of the technologies and equipment that are on the ships because there is a crew. Therefore, there would not be a need for a large part of the maintenance that exists on ships if the crew were not on board – and in that way an autonomous ship makes good sense.

But it is not quite that simple. Curt told about a thesis that dealt precisely with this discussion, and which also concluded that much of what the crew does now would not be necessary on autonomous ships. The student argued that if, instead, you had three of all the technologies, such as GPS, ECDIS, etc., then you would not need to maintain them before the ship reached port, since the ship would be able to sail even if one of the technologies broke down.

There are, however, several problems with this. The first, according to Curt, is that there probably is no shipyard or shipping company that would let a multi-million-dollar ship sail without knowing what is actually happening on board, even if it "just" sails across the Atlantic and turns up in port 10 days later. That is simply too much money. This also relates to trust, which will be touched on later.

In addition, very few shipowners would be comfortable spending that much money on a ship without having anyone to make sure it arrives in port in good condition. Finally, there is a problem with the technology that would be required to have an autonomous ship, as technology does not hold up very well at sea:

"All experience with sending film mechanics out to sea is that it works terribly, if we use a good expression, right? So, I think that there will be a new maintenance item at that point, which will also include some skills that people don't have today. So that's really where the film breaks for me in relation to this fully autonomous ship. It's that this whole question of maintenance, this whole business case question, I haven't seen anyone prove that it can make sense." (Interview with Curt)

The new technology that would be needed on the ship for it to be autonomous would also mean that there was new technology on board that would need to be maintained. And since all experience with taking, for example, film equipment out to sea shows that it breaks because of the hard environment, it would require that it can be maintained by a crew, who would also have to be trained in it.

The two major problems with autonomous ships, according to Curt, are maintenance and money. It takes a lot of maintenance to sail a ship on the open sea, and in addition to the ship itself costing a lot of money, the cargo ships carry is often very valuable.

One area where Curt, on the other hand, believes autonomous ships make a lot of sense is for small measurements. Specifically, he mentions how a major energy company has some drones that they use to take measurements for offshore wind farm projects:

"...a drone like that at 100,000 dollars is doing a job that used to cost 50 million dollars. And that's why people very often chose to say, well, we just won't do it. When we're designing a wind farm, we make an educated guess based on some statistics, but we don't investigate the actual conditions. Because it's so damn expensive." (Interview with Curt)

Here is an example where drones or autonomous vessels allow various companies to save a lot of money. These measurements also do not take nearly as long to do, whether it is with a ship with crew or drones. It is, in a sense, also a more controlled area, since it is not far out on the open sea, and it is an area where the risks are smaller in relation to losing a drone.

This also shows one of the differences between robots in industry on land and robots on ships. The robots mentioned in Sub-Report III and some of those in Sub-Report I have all contributed to reducing the financial burden for some companies.

A good example is the cleaning robot from Sub-Report I that supposedly led to the cleaning staff being cut in hours. It is cheaper to run in the long term than staff and possibly just as efficient. Even agricultural robots can make economic sense, as they can save farmers the working hours they would otherwise spend driving on the fields.

But at sea, with ships, it is apparently not that simple. Just the price of a potential autonomous ship would be extremely high, not to mention the investment required to develop it. The crew on ships nowadays is already so small that it hardly pays to replace them. Not to mention the amount of trust the shipowners, customers and consumers would have to have that the autonomous ship can deliver their goods properly, without damage, and that someone will be held responsible if the ship fails.

This leads the report on to the next topic.

Trust - ShipGO

Trust was a theme that turned out to be very relevant in relation to the automation of shipping. It is something that is already being worked on, not just in our project. Many new technologies have been developed to improve navigation on ships.

One of the problems with implementing new technology is that ship officers must both learn to trust the technology in order to use it, and at the same time not trust it too much, as it can also be wrong.

One of the reasons for this is that when, for example, ECDIS was introduced, it was based on charts from 1949, whose data was based on measurements from before the 1920s. As Hanna described it:

"The data was the same from 1970, and there was no better depth than what was measured back then. So there were actually, you could say, ECDIS-related groundings."

This points to situations where people have had too much trust in a system that did not have better or more precise data than the charts they themselves had used for many years. Because they then possibly did not look out of the windows at their surroundings either but just assumed that ECDIS was a 1:1 version of reality, they then ran aground.

In the interview with Hanna, she herself told how she has learned during her education that technologies are never 100% precise:

"This distance measurement can have some percentage error." (Interview with Hanna)

Here she is talking about radars that send out waves so that you know how far away another ship is. It is, however, not completely precise, and therefore it can be dangerous as a captain to trust 100% in the technologies.

Autonomous ships will therefore also require enormous trust from the shipowners' side that they will not run aground, that their software is advanced enough, and that their data is up to date. In addition, there will also be a need for the ship to have safety measures that can be activated should the ship's primary navigation systems fail. As mentioned in regards to safety, one of the issues with autonomous ships, is how expensive they would be, without even accounting for their cargo, and as Curt mentioned, no shipyard in the world, would trust an autonomous ship to sail across the ocean without any supervision and a crew to make sure nothing goes wrong, and if there is a crew, what's the point in investing millions or potentially billions in new technology, when the current technology works just fine?

On a ship with a crew there still exists a classic authoritarian leadership style:

"Well, it's a very classic authoritarian leadership style with the captain more or less on a par with God (...) And that still exists legally." (Interview with Curt)

This is not to say that everything the captain says is correct, but that in the end he is the one who bears responsibility. That is also what Curt means when he says "legally." If something goes wrong – if the ship runs aground or has a collision with another ship – it is ultimately the captain who is held responsible for this.

It is also the captain who must approve the voyage planning and be able to stand by it, since according to maritime law it is he who has responsibility for the cargo as well. In practice, many shipping companies actually approve the voyage planning, and the captain just has to give it a rubber stamp, but it will legally still be his responsibility. The same applies to how heavily the ship is loaded, which to a greater extent is determined by a decision-support system that must then convince the captain that it is a good decision.

So there is also a trust issue in relation to new technologies and systems. Curt described it as follows:

"The computer system is very often in the same situation as a newly graduated officer. They come out and then there's this old guy sitting there saying: yeah, yeah, I've been doing this for a lifetime. Now you just show me what you can do." (Interview with Curt)

The aforementioned captain Frank is one such captain, who himself described how he only trusts new technologies when they can show him that they can improve safety or make his life easier on sea. He has experienced a lot of new technologies being introduced on ships, and the biggest change for him seems to be that he now spends most of his time in front of a computer, answering various, in his view, irrelevant emails and providing data, available through other channels, to engineers and software developers. That being said, he is not against new technologies, but he will always be sceptical towards them at first, and even hinted at a potential walkout or strike, from captain and sailors if too many autonomous systems are implemented, because they wouldn't trust them either.

This also leads to the implementation of new control systems being extremely slow, and Curt himself also told us about a system called SPOS, which was to help with planning routes across the sea based on, for example, weather data. Here Curt says that it took several years before they really started using it, because at first they just drew their own routes across the sea, let SPOS make its own, looked at them and thought they were interesting routes it suggested – but took their own first.

This is also relevant in relation to one of our other cases, where we learned about, among other things, "alarm fatigue" and how captains on modern ships get so many alarms that they stop registering them.

Specifically, an instructor we spoke with told how he had experienced being on board a ship where an alarm started sounding. As the sound used for alarms on the ship was the same for all alarms, he did not know what it could be. This led him to check through all the critical alarms, then all the less critical ones, and there were no faults anywhere. In the end it turned out to be an alarm for a freezer that was not completely cold.

This is also a form of mistrust of the different systems and a possible over-trust in one's own abilities as captain to judge which alarms are serious and maybe even an ability to spot possible problems oneself. It can lead to one beginning to ignore alarms because one doesn't think they are important.

Therefore, one of the first steps in the automation of ships could also be some kind of filter-to-filter alarms according to their importance. This would however also require a lot of trust in the technology that a captain like Frank doesn't have yet. To him a lot of technologies are simply developed and implemented for technology's sake. Meaning that instead of figuring out if there is a need for the technology, he instead feels that the technology is developed because it can be.

Anomalies - ShipGO

Anomalies were not, as such, something that figured much in relation to the simulators themselves or the remote-controlled ships. Rather, it was something that was mentioned in relation to the navigation technologies used by the different ships, when they sometimes displayed incorrect data or in other ways did not work as intended.

To the two instructors this sounded like an almost constant feature at sea: you cannot always rely on your instruments and data, because the sea is very unpredictable. Both instructors had experienced situations on board a ship where their technology did not work as it was supposed to.

This speaks to how difficult it is to develop an autonomous ship, as they require it to be able to handle unforeseen events. And it requires that the ship has independent methods to determine its location and calculate its course. And if the two independent methods have to exist on an autonomous ship, would that also mean that it needs to have two "autonomous captains" that can control the ships independently of each other?

Sim2Real - ShipGO

This is possibly the most relevant topic in relation to RoboSAPIENS. The primary reason for our field visit was precisely that they have these simulators, which they use for teaching. Particularly relevant for other of our cases is the discussion we had with them about digital twins.

This is something that is used in several of the project's use-cases and something ShipGO also has experience with. The school worked together in a project with the municipality in which the school is located, which was about building the world's longest-sailing electric ferry. Before they had finished building the ferry and launched it, they trained the crew and staff via a simulation of the ferry.

So they had programmed a digital twin of it, which they tested in different conditions. In addition, they also discovered an error in the simulation where some parts of the ship did not fit. When they mentioned this error, the initial reaction was that it must be a fault in the simulation, but it turned out to be an error in the construction of the ship.

They discovered this when they had the drawings checked that they had fed into the simulator. Unfortunately, they did not discover it quickly enough to stop the ship from being built, but they did manage to inform them before the ship was launched.

This is quite simply an example of how good simulations can actually be, as they can tell you something about reality, and they can be used to discover errors and run tests before it is too late.

However, not everything can be captured by the simulations. Hanna also told us how there are some things that are difficult to simulate, including icebergs. Icebergs are generally difficult to navigate and handle for ships. She said this is why ships in Greenlandic waters still have the wheelhouse at the front. They have to be able to see the icebergs, and even radars today find it hard to see icebergs.

Therefore, there are also two crew members whose job is to spot icebergs and say which way to go around them. Hanna also told how the captains who sail a lot in Greenlandic waters know the ice and the icebergs. She described this knowledge almost as bodily and as something that would be almost impossible to train a robot to have.

An experienced captain could get within what looked like a few centimetres of an iceberg without problems, because he knew the area and knew exactly where it was safe and dangerous to sail, she explains.

An interesting point Curt mentioned was that they had tried to explore the possibility of using the data they collect during various teaching sessions with students. He thought that this kind of data should be almost worth gold for developers of autonomous ships or navigation technology.

Their students have and have had many thousands of hours of simulations, and all the data they collect about how they act, or how the ships act, could be used for the development of better ships and give insight into captains' decision-making.

Sabotage and Tinkering - ShipGO

During our fieldwork the issue of sabotage, was not something that was discussed in great detail. However, we did discuss tinkering and how sailors and captains have to be comfortable doing a lot of testing to get acquainted with new technologies.

During the interview with Curt, he told us that because there is no standardised implementation of technology on ships, the first two weeks on board a new ship are spent, as Curt describes it, "playing 'guess a new button'."

In other words, it is almost a necessity for the crew to tinker with the technology on a ship. Because they do not know exactly what all the technology is, they have to learn it on the job. This is also something we have seen in other situations where people are about to start working with new technology. One of the ways we have observed people develop trust in robots is through tinkering.

This could also be a potential problem for the implementation of autonomous ships. If there is too much variation between ships, it will require each shipyard to develop its own autonomous ships, as it will not be possible to create standardised robot technology. So, to create truly autonomous ships, it might be necessary to make entirely new ships, where everything can be standardized.

Something we did discuss in relation to sabotage, was the potential for sabotage of autonomous ships. Specifically in regard to hacking of the ships, by foreign powers, either to completely guide them out of course, towards a different harbour, or just simply stopping it, and boarding it.

Furthermore, Frank hinted at a potential sabotage, if they start making autonomous ships, and how he believes most captains would refuse to work on such ships, since they would not want any responsibility if the ship went aground if its controlled autonomously. This refusal of technology can be seen as a form of sabotage. Because it will also halt the development of the technology.

Conclusion

Developing autonomous robots for dynamic environments like a field or the ocean, is incredibly complex. The environments themselves offer plenty of challenges, both furthermore the current laws existing for these environments in the case of autonomous ships, make it legally impossible. There simply needs to be a change in the law of the sea, before autonomous ships can exist, and furthermore, there needs to be a cultural change regarding responsibility.

This report has shown that the more dynamic an environment gets, the more difficult it will be to develop fully autonomous robots, and perhaps it will even be impossible as shown with the 3 ships problem. However, it has also shown that while fully autonomous might not exist, AI and more autonomous decision making, working in collaboration with captains, might have a lot to offer. One of the big issues we encountered in this fieldwork, was alarm fatigue, and the captains becoming almost indifferent to the alarm. Here it has been proposed, on numerous occasions, that AI could be a helpful tool, since it can filter large datasets and classify them. However, this raises a myriad of new issues, such as the degree of transparency and the distribution of responsibility.

This report has examined autonomous robots in two markedly different, yet analytically aligned, dynamic environments: agricultural fields and the open sea. Across both cases, autonomy emerges not as a stable technical achievement, but as a fragile and continuously negotiated arrangement shaped by environmental volatility, organisational constraints, regulatory regimes, and human judgement. While agricultural robots and autonomous ships differ greatly in scale, purpose, and context, they confront a shared problem: dynamic environments resist standardisation, prediction, and full control.

The empirical cases presented here show that the more dynamic an environment becomes, the less feasible it is to stabilise robotic autonomy through technical design alone. In agricultural settings, robots must contend not only with human actors, but with weather, terrain, vegetation, and animals that behave in ways that are difficult to anticipate or formalise. At sea, this challenge is amplified further: the environment is physically unforgiving, continuously changing, and governed by rules that explicitly rely on human judgement rather than deterministic decision-making. In both contexts, errors often reveal themselves too late, when the cost of failure is already high.

Across the report, we have shown that autonomy in dynamic environments does not remove humans from the system. On the contrary, it redistributes and often intensifies human responsibility. Farmers remain responsible for supervising machines remotely, intervening when anomalies occur, and ensuring that robots operate safely and efficiently. Captains remain legally and morally responsible for ships, even as navigation systems become more advanced and automated. In both cases, responsibility cannot be fully delegated to machines, developers, or algorithms. Instead, autonomy concentrates accountability at specific human positions—most clearly at the farmer and the captain—while developers and organisations often remain at a distance from legal and moral liability.

The report also demonstrates that safety and trust are not properties that can be engineered once and for all. Rather, they are ongoing achievements that depend on situated practices, experience, and confidence built over time. In agriculture, trust hinges on operational reliability, ethical concerns such as animal welfare, and the robot's ability to support farmers' professional identities and values. At sea, trust is shaped by long histories of technological failure, alarm fatigue, and an acute awareness that digital systems can be wrong. In both cases, trust is fragile and reversible: it is earned slowly and lost quickly.

Simulations and digital twins play a crucial role in making autonomy thinkable and testable in such environments, yet this report has also shown their limits. Simulations are indispensable tools for training, prototyping, and risk reduction, but they cannot fully capture the complexity of dynamic environments. Rather than bridging the gap between simulation and reality, Sim2Real practices function as filters that select which aspects of reality can be acted upon, while excluding others that remain dependent on human judgement, embodied knowledge, and experience. The cases of agricultural fields and ice-covered seas illustrate that some forms of environmental knowledge remain deeply situational and resistant to formalisation.

Finally, the report has highlighted that resistance to autonomy—whether in the form of scepticism, refusal, tinkering, or even sabotage—should not be understood simply as opposition to technology. Instead, such responses often signal underlying tensions related to responsibility, identity, power, and control. Where autonomy threatens to redistribute risk without adequately redistributing authority or accountability, resistance becomes a meaningful social reaction rather than an obstacle to be overcome.

Taken together, Report IV shows that dynamic environments mark a critical boundary for robotic autonomy. They expose not only technical limitations, but conceptual ones: situations in which autonomy itself becomes difficult to define, implement, or justify. This does not mean that automation and AI have no role to play in such contexts. On the contrary, the report suggests that partial autonomy, decision-support systems, and human–AI collaboration may offer significant value—particularly in filtering information, managing alarms, and supporting human decision-making under conditions of uncertainty.

In this sense, Report IV completes the empirical arc initiated in Reports I–III. If autonomy is imagined in Report I, challenged and disrupted in Report II, and organisationally compensated for in Report III, then Report IV shows where compensation itself reaches its limits. The findings presented here thus prepare the ground for Report V, which will turn explicitly to autonomy as a concept, asking not only whether autonomous robots can exist in dynamic environments, but what kinds of autonomy are socially, legally, and ethically desirable—and at what cost.

References

Reissig, L., & Siegrist, M. (2025). From the attitude towards digitalisation in agriculture to the acceptance of future agricultural technologies. Smart Agricultural Technology, 12, 101095. https://doi.org/10.1016/j.atech.2025.101095

Martin, G., Rives, J., Cavet, C., Moraine, M., Bohanec, M., Debaine, F., … & Hostiou, N. (2022). Robots and transformations of work in farm: A systematic review of the literature and a research agenda. Agronomy for Sustainable Development, 42, 96. https://doi.org/10.1007/s13593-022-00796-2

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.