Robots in the Wild

Report III: Robots in controlled Environments

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

Introduction

This report seeks to uncover the collaboration between robots and humans in workplaces where they interact in organised everyday lives and have established practices.

Where Report I examine how autonomy is imagined and Report II explore how autonomy is disrupted and practiced in public spaces, Report III examines how autonomy is operationally replaced by human organisation, routines, workarounds, and distributed responsibility in controlled environments. Controlled environments we define as spaces where autonomy in robots is not expected, but where human–robot collaboration is actively maintained through organisational effort. In these controlled environments we find stable user groups, repeated interactions, formal responsibility chains, safety regimes and standards and ongoing learning, tinkering and adaptation.

In public spaces (Report II), robots can be freely observed and described. In closed organisational environments (Report III), companies restrict what researchers can directly observe, document, and report. Therefore, this report relies less on detailed description of observations and more on interviews and approved descriptions. We visit a company programming and selling industrial and collaborative robots (ElectroFlight), as well as a company developing and advertising such robots (MakeBot). Additionally, we present a company with plans to deploy a robot (GreenTech), along with two other companies where industrial robots are an integral part of everyday operations, though with significantly different challenges (HandsON & MedCare).

Though we learned a lot of background knowledge from observing robots in everyday lives of employees we have not been able to disclose any detailed accounts here and therefore shall rely more on the accepted interviews. The report will examine how companies relate to the previously mentioned issue that co-bots are often oversold as quasi-autonomous (see Report I) – that is, presented as more capable than they actually are – and what significance this has in practice. In addition, it will shed light on the challenge of robot autonomy in workplaces where humans constantly form a central part of the collaboration. The focus will be on how this collaboration is built up and broken down, and which mechanisms are used to facilitate a workplace with Human-Robot-Interactions (HRI).

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

The different field sites will first be introduced, and the overarching issues or areas of interest for each field will be outlined. All sites are connected to the transversal analytical categories Safety, Trust, Anomalies, Sim2real and Sabotage/Tinkering, which have been developed in connection with the RoboSAPIENS project. The report is structured such that it begins by introducing the field sites that have an overarching perspective on companies with automation ambitions. It then goes in depth with field sites where these automation ambitions unfold in practice. Finally, these fieldworks will be connected to each other and placed in the context of the future of autonomous robots in workplaces. Furthermore, the overarching issues that this report has illuminated will be presented.

ElectroFlight - Distributing Robots

The fieldwork at ElectroFlight contributes to a practical and industrial understanding of what it means to automate human work functions with the help of robots. In addition, the company's extensive experience in securing the safety of robot systems, contribute to a greater understanding of what safety actually means in robot development. The fieldwork also sheds light on a classic problem in automation: what do we do with the employees, and how do companies' expectations align with the actual possibilities? As mentioned in Sub-Report I, these prophetic ideas of what robots can do exist, generally inspired by science fiction. During the visit at ElectroFlight we also discussed these same unrealistic ideas and how the robots must be downplayed to the customers, as they expect more than what the robots can actually do. Even if these expectations can be moderated, it is critical that safety measures are in place in case customers nevertheless have other ideas about what the robots can do.

Introduction

ElectroFlight is located in a small industrial area outside a larger European city, and the surrounding companies are all IT- or industry-related. Their office building is very small, and the many desks you are met by all point towards the entrance. A single employee looks up from his screen and greets us as we step inside but then continues with what he was doing. There is no receptionist and therefore no one who particularly notices that we sit down in their reception area. We have also arrived in good time, and 5 minutes later Nick comes and receives us. When we get coffee, we jokingly ask whether this is some kind of Silicon Valley, to which he laughs and says that the IT companies located here are a bit from the old world and not so much this new software world that is booming in, among other places, Silicon Valley. After we have fetched coffee, we go out into their warehouse, and we are very surprised at how small the warehouse is. The contrast becomes especially striking because of the large industrial robot that meets us as we step inside. We also comment that we are impressed by how much they have managed to fit into an otherwise small warehouse, to which he says that this is often also a challenge for them. We walk with Nick through their warehouse while he shows us a number of different robots and we talk.

At ElectroFlight the way it works is that they are contacted by a company that has a work task or function they wish to have automated. ElectroFlight then sends some employees out to investigate what they want to have automated, and here they also align expectations with the customers about what is possible. After this they create a series of simulations of the robot in order to visualise what it is supposed to be able to do before they begin the programming. The simulations are also used to uncover the safety distance to the robot, and which additional safety measures are needed. Sometimes they also use these simulations to show customers how the robot will be working. They use two different tools to create these simulations, one of which is their own CAD drawing program, and the other is one they received together with the industrial robots they program. After this they begin the programming work of the robot, which is a lengthy process consisting of many tests and iterations of the robots. In addition to programming the robots, they also program an "HMI" (Human-Machine Interface), with the aim of making the robot as user-friendly and intuitive as possible for the customer. When they have finished working on the robot, it must be implemented at the customer's site, and here ElectroFlight is present to ensure as safe and successful an implementation as possible. Nick explains that all their solutions can be accessed remotely, so they are therefore also available to customers long after the robot has been implemented.

The robots

ElectroFlight is a reseller of industrial robots and cobots, however they mostly work with industrial robots, so all examples included in this section are accounts of industrial robots. When we arranged our visit with Nick, he told us about a robot they were working on at the time, which was supposed to place some blocks into a box. When we come out into the warehouse, this is also the robot he goes to first. Nick explains that they have equipped the robot with a vision camera which is to act as the robot's "eyes".

In the task with the blocks, the camera will be able to detect if a block is missing or if a block is turned the wrong way. However, the robot will not be able to solve the "problem"; the point is rather that it can inform an employee about it, who can then solve it.

The vision camera is supposed to make the robot more aware of its surroundings. At one point we ask whether the vision camera is also meant to feed information to the user platform, to which he says that the camera is first and foremost there for the robot's sake, so that it can see what is going on around it. He says that one can send the camera's information on to the user platform, but that this is not its primary function.

The vision cameras come in both 2D and 3D, and there is a very significant difference between the two. Whereas the 2D camera can see everything and detect shapes and figures, the 3D camera can see depth and thereby detect multiple levels. As an example, he explains that if a robot is handling test tubes and drops one, then a 2D camera will detect that it has been dropped and see where it is lying. The 3D camera would, by contrast, be able to see how large the glass pieces from the broken glass tube are in terms of depth and if there are two pieces, for example, lying on top of each other. In both cases, however, the robot would not be able to handle the situation or pick up the shards of glass, but it would instead inform an employee about the accident.

At one point we walk past a robot where there are stickers indicating that you are not allowed to touch it. We asked Nick why you are not allowed to touch the robot there, and whether it was because you might get a finger caught or similar. He replied that it was for the robot's sake that you must not touch there, because you might end up covering its vision camera and thereby make it unable to see anything.

He explains that this task of placing the blocks is normally performed by an employee and works quite well. The company that wants to have this task automated is not doing so with the wish to replace employees, but in fact because they have a shortage of employees. Nick says:

"So, the [larger companies] just snatch up all the employees, and then these slightly smaller companies are left saying, well, we can't offer twice as much in salary. But the others can. And that's what they're afraid of."

Thus, the smaller companies do not have the financial capacity to compete with the larger companies and can therefore be forced to invest in robots to increase their turnover and in the end be able to hire more employees.

Nick also explains that it is about showing that the robot can do the work just as quickly as humans, rather than humans having to live up to the robot's speed: "It's really about demonstrating that the robot can do it pretty much just as fast as a human."

However, the robot is not fully automated, because once it has finished its work, an employee has to come and approve the box and possibly check whether anything is missing. Nick also explains that in precisely this case, the robot is worth nothing unless people have given it the blocks or the box so it can begin the work of filling it:

"That means this operator who has to deliver the blocks. If he hasn't filled it up, then the machine is worth nothing, then it's just standing there waiting. So, it will park itself in that state that says, well, we're missing product. It can also be the other way round, saying, now we're finished. But the person who is supposed to fetch the box and put a new one on never came."

The robot therefore does not have value in itself, but only in connection with the employees' job functions, and it is completely dependent on there being a human present. In addition, Nick says that the robot rarely does anything on its own, but that it is always in collaboration with other machines. It doesn't have autonomy and only works in collaboration with the human employees.

At ElectroFlight they mostly work with industrial robots. The robots in their warehouse were very large, and they in a way seemed clumsier and more impractical when you saw them standing still. The robot that met us at the entrance to their warehouse surprised us quite a bit with its height, as it seemed much more imposing than some we have seen before. In the project we have seen many cobots and AGVs, and they seem "neater" in their mechanics compared to more factory-like industrial robots. When you came close to them, you could more clearly see all the nuts, bolts and mechanical parts that together make up the large robots. We felt a completely different kind of respect for the robots when we saw them up close in that way. Size can also matter when humans relate to robots, as it is clear with the *toyification* of robots, where designers try to make products more appealing by deliberately giving them toy aesthetics (Thibault & Heljakka 2018).

Cobots vs. industrial robots

Nick, however, explains that he rarely hears comments from customers about the size and appearance of the robot, but rather about their ideas and expectations of what it can do. He describes how the less experienced customers often have wild and unrealistic expectations of what the robots can do, and that he spends a lot of time downplaying the robots (trying to counter the ‘overselling’ of robots found in Report I): “*So actually it’s quite a big thing to talk the robots down. We joke a bit that some people think they make coffee in the breaks, the robots. Then they can mop the floor and stuff like that. They don’t. That’s typically people who know very little about it.*”

When customers call, they always have lots of good ideas about what the robot should do for them, and here ElectroFlight sticks to what they call the 80–20 rule as an attempt to get the customer to cut down on their expectations. In general, Nick spends a lot of time talking the robots down when they talk to new customers, in order to ensure that they do not expect too much from the implementation of them.

In terms of safety, Nick emphasises that there is a big difference between how you secure an industrial robot and a cobot. Cobots do not need to have a safety fence around them; instead, scanners are used that emit infrared beams which register if there is a person walking in front. The safety beam works such that the robots stop immediately if the beam registers something entering in front of it. When you then leave the robot’s area again, and thus move away from the safety zone, you have to acknowledge that you have left the place, so that the robot is told that you are gone, by the press of a button. If a new person enter just as another person leaves the area, then the robot cannot register that, and the new person is put at risk. Some cobots are physically fenced in, but Nick explains that it is often not practical that they are fenced too much, since they are precisely collaborative robots and are supposed to work together with humans. Industrial robots are much more static than cobots because they require much more safety and indeed require this complete fencing.

Nick at one point refers to another robot manufacturer and their marketing of cobots, where he to a certain extent believes that they have been oversold, as customers have an expectation that they can do much more than they actually can. Here he also talks about various robot conferences where you see the robots walking around among people and handing out chocolate, pens or other things. This leads people to develop unrealistic expectations of what cobots are suited for (Nick’s statements are a confirmation of what we have written about in Report I). For even though a cobot could technically walk around and hand out food all day in an office, in terms of its price this would be an inappropriate and economically inefficient use. Conversely, Nick also sees cobots being worn out because they are given tasks they in fact cannot or should not perform. In addition, customers become dissatisfied with how slowly it works, but if they instead want an industrial robot that works much faster, they lose the interaction and closeness with the robot that was what drew them in. The many fairs and conferences contribute, in Nick’s view, to oversimplify and confuse, because people find it harder and harder to distinguish between a cobot and an industrial robot. But the difference is significant, because whereas a cobot is designed not to harm humans since it is to collaborate with them, industrial robots are not at all designed to collaborate with humans, which is why they are fenced in. If you are not explicit about this difference, you risk putting your employees and yourself into a potentially life-threatening situation in which you do not have a realistic idea of what machines you are actually working with.

In general, they do not sell that many cobots, and according to Nick this is probably because people are quickly impressed by them, but when they then discover how little they can and that they must now work side by side with them every day, they back out. Despite this, there is a tendency for cobots to gain a foothold in modern work environments, and this places demands on both the interaction and our expectations of the technology

Nick also explains that you have to be certified in order to be allowed to call a robot collaborative. Cobots are characterised by moving much more slowly than industrial robots and by lifting far fewer kilos. In addition, cobots are programmed such that they can feel if resistance occurs, and you can set how much resistance they have to feel before they stop. Industrial robots do not have this feature, since they are not intended to work side by side with humans. The assessment of how much resistance the robot has to feel before it stops is made by the robot manufacturers and is thus something the robot is "born with" when ElectroFlight buys it. Taken together with the previously described HMI solution, where screens are bought from an external company, but the user interface is developed internally, it becomes clear that these robots are based on distributed technologies. They are not built from scratch in one place but as a sample kit with part from many places. This is seen again in the fact that the respective robot manufacturers create a series of risk assessments which ElectroFlight works further with and then sells on to their customers. So, it is not only the robot part that are distributed. The ethical responsibilities for the workings of the robot's parts are also distributed in each part as well as their assemblage.

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 - ElectroFlight

We ask Nick whether they experience risks that are only discovered "after the fact", to which he surprisingly says no and *"there shouldn't be, really."* He also says that they do not experience having overlooked risks, and this is probably connected to their extensive experience in the field. Nick explains that one of their biggest problems in terms of ensuring safety in their robots is actually human unpredictability. He describes how this unpredictability, among other things, arises when employees feel that they can work more productively if they compromise on safety. This situation is well-known within the world of safety analysis in organisations, and the philosopher Erik Hollnagel (2009) calls it the ETTO principle: Efficiency–Thoroughness-Trade-Off. In relation to that we ask Nick whether they have ever tried to reduce safety in order to make something more user-friendly, to which he promptly replies that you are not allowed to reduce safety. Here he is referring to the *Machinery Directive* and the rules dictated therein. The Machinery Directive is a transnational standard that only applies to member states within the EU. The directive is developed by the European Commission and is intended to ensure the safety of machinery. If all the requirements of the directive are met, a product can obtain a CE certification, which means it can be sold in Europe. Thus, this standard is not only intended to ensure the safety of machinery, but also to ensure fair competition on the European market. In continuation of the above, Nick also says that customers in general are quite annoyed about the many rules and standards. The rules tend to make development more difficult, especially since new initiatives are constantly being added to the directive:

"It puts a stop block on a lot of things. Because we have so many different ways of doing various safety measures. But it's always connected with some kind of expensive cost or something that hinders production. That's how it is."

Besides being cumbersome, Nick also describes here that in general it is quite expensive to acquire these safety measures. It is therefore not only productivity that you may have to compromise on, but also the budget.

Nick explains that risk assessment is necessary for their robots to be CE-marked and thus sold in Europe. There are public standards for risk assessment, such as the ISO standards and the *Machinery Directive,* and these are the ones they follow at ElectroFlight.

Connected to this are ways of designing your way out of the different risks, such as the aforementioned fences. In connection with this, Nick says: *"Fences are not put up so that the robot cannot get out. They are put up so that employees cannot get in."*

Safety is not just about the humans but also robots that must be kept safe from human tinkering. The risk assessment of the robot is in fact also to protect the robot from human interference. Nick elaborates that this is because humans do not take the safety measures seriously and think they can just go in to the robot or move the fence. When we talk about the different standards and regulations that exist in relation to safety, Nick explains that as a manufacturer he has the responsibility that no one gets hurt with his products, and therefore they make sure to follow the rules. However, he immediately adds that he cannot control what the customers then choose to do afterwards. He compares this to a car manufacturer and seat belts: *"What they then come up with afterwards, just like if you drive your car and say, well, we can't be bothered to use the seat belt. You can't blame the manufacturer for that. It was there, it worked, but you're not using it."* Here Nick points out the challenge of working with these distributed technologies, since as a manufacturer you can end up being held responsible for something that happens further out in the distributed chain. This can also lead to doubts about responsibility, and, in some cases, responsibility being pushed away, since the many elements of the chain make it more difficult to maintain an overview of those involved and their responsibilities.

Besides using the standardised procedures for risk assessment, Nick also explains that they have a great deal of experience in risk assessing robots, and much of it is therefore based on this experience. There are a number of preliminary safety measures they carry out, and these are carried out based on the physical reach of the robot.

They take the robot arm and make a circle around it, which is then its reach all the way around. This area is secured with fences or beams. Besides the robot itself, it is also relevant what tool it is to handle. Nick makes the point that there is a big difference between whether it is holding a stick or a knife, and that if it is holding a knife then suddenly it is the item it is holding that is dangerous rather than the robot itself.

These risk assessments are therefore very unique and depend on which sector you are in. They also depend mostly on which robot you are dealing with, for Nick says*: "So that's why* *cobots cannot always just be defined as being safe. It's the tool that helps determine whether they are safe."*

The robots come with what is called a "teach pendant", which can best be translated as a handheld programming unit and is used to program and control the robot. At first glance it is not very intuitive for non-technical staff, as many of the buttons have symbols they don't recognise. The teach pendant is also developed such that you cannot use the buttons on the front without holding certain safety buttons down at the same time. Even though this is there to ensure safety in the robots, it is not very user-friendly for the ordinary operator. Nick explains that they use it when they program and test, but that when it is to be sent out to customers, they "hide" it away. We asked Nick whether any of their employees have previously been injured in working with the robots. Here he refers back to this teach pendant, as it is programmed such that they hold a safety handle in with a very specific pressure to prevent harm. If they press too hard, the robot stops, and if they do not press hard enough, the robot also stops, so they are therefore kept safe in their interaction with the robot if they trip or are inattentive.

As mentioned earlier, ElectroFlight programs its own HMI, which they do on screens they buy from another company, where they then create all the software behind the screen. Nick explains that their customers experience a kind of "joy of recognition", because if they buy a new robot from them with a new function, then the HMI will be the same.

In addition, the large screen is also something the customers want, as it makes it more operable for non-technical staff. Because ElectroFlight can themselves control all the software behind the screen, they can also adapt it to the individual customers' needs. Some customers need a lot of information, while others just want the absolutely most important messages displayed. In principle the robot could inform about everything it does*: "So it could give information every time it has done something; that's rarely something you're interested in. Because that's just part of the process, you sort of expect that."* Here Nick is hinting at information overload (that too much information can be a nuisance and not a help), He also explain that the costumers have different needs for information. Some are interested in one kind of information, which others do not care about. Their customers cannot themselves change this HMI, and Nick explains that this is because it is part of the product they as a company deliver to ensure safety.

Nick explains that they have not previously experienced the robots doing something they were not supposed to, and he accounts for this by saying that they have programmed the robot to do specific tasks, and therefore it cannot do anything else *"because fundamentally it's completely stupid."* He describes them as *"slaves to the software they have made",* and that if the robots have on rare occasions made a deviation, they can always trace it back to a "bug" in the software. If something unexpected were to happen, then according to Nick it would be due to a special scenario where something has happened to a sensor, and therefore the robot would be unsure what to do. That is why it is extremely important for them that they have thoroughly tested the robot and its functions before it is to be implemented with customers.

Trust - ElectroFlight

Nick explains that some of their customers do not want their production employees to find out that they are collaborating with a robot automation company. They are worried that the employees will fear being replaced, even though Nick points out that this is rarely the purpose of automation. Nick also reiterates that automation doesn't necessarily mean that the robot will completely take over everything, as was clear in the case of the robot described in the intro. The employees therefore do not trust the robots from the very beginning, as they see them as competitors to their work. Nick, on the other hand, sees robots and automation as an opportunity for employees to gain more competencies, either because they can get new tasks or because they acquire more IT expertise. He believes that it is the older generation and less IT-familiar employees who are most frightened. He also believes that if companies are not open to automation, they will not be able to survive in the industrial sector today. In continuation of this, Nick explains that these robots are programmed to perform very precise tasks, and it is perhaps when this programming fails that we become worried and lose trust. People have an expectation that robots just work, especially the non-technical staff, for whom the robot's interior is just a black box. Their mistrust perhaps stems from not knowing what is going on inside this black box when something goes wrong, and therefore they do not trust that nothing worse can happen. In contrast, you can simply ask a new colleague *"hey, what happened?",* to which they can reply that they were probably just nervous or unsure what to do. In relation to safety, Nick explains that they have ensured the robot to such an extent that even if it drops something, nothing will happen to anyone. The employees' mistrust is therefore not based on fear of physical unsafety, but perhaps more on a feeling of uncertainty around the robot and its black box.

We asked Nick directly whether he trusts the robots, and here he explains that he does, because the robots are fundamentally just stupid, and it is them (the programmers) who decide what they are to do. He also says that he trusts that it will not do anything they had not expected, because they have decided everything it is supposed to do. He says that if it were to do something unexpected, then it would be because they as programmers have made a mistake. Here one can deduce that his eventual mistrust of the robot in fact lies with the programmers rather than with the robot itself. In connection with this he also says that just because he trusts their programming does not mean that it should not be completely thoroughly tested. His trust in the robot's programming is not the same as it being safe, and therefore it is important that it also lives up to all the safety standards, despite them trusting it.

They also experience that the older generation is very sceptical from the start, and if the robot makes the slightest mistake, they say *"That's crap, that will never work, and now it dropped that thing."* That is to say, these employees' trust is practically non-existent from the start, and it only takes the slightest misstep for it to disappear completely. Nick compares it to getting a new colleague who also needs time to settle in and will inevitably make mistakes or drop things. What he is pointing to is that you wouldn't usually judge a new human colleague as harshly, or get as annoyed with them, as he feels some of their customers' employees do with the robot. There is, in other words, a difference in trust: these employees seem more inclined to believe that a human who makes a mistake will learn and not repeat it, whereas a robot that makes a mistake is seen as less trustworthy and more fundamentally flawed. Perhaps precisely because robots themselves don't "learn" from their mistakes, yet.

Over time, however, employees' trust grows, Nick explains, and he says that after about three months following implementation, the employees begin to see all the positive things about the robot. The employees even start to become happy about the robot, especially when they realise that all the tedious work they previously had to do can now be done by the robot. Slowly they also begin to take more interest in the robot and explore more closely how it works. Nick explains: *"So when the boss comes out, it's cool to be able to say 'right, now you just have to see here. Yeah. We can show you the statistics, how it went yesterday.' Then suddenly they have something else to show."* The employees' trust in the robot thus grows with their own IT interest and competencies.

Experimenting with the technology in this way can be called "tinkering", and it can be seen as a way of expanding practical knowledge about the technology in question. Through the employees' tinkering with the robot, they actually strengthen their trust and collegial collaboration with it. However, not all employees engage in this form of tinkering, and here Nick especially refers to the older generation. According to him they are less IT-familiar and therefore also more resistant to having to experiment and try something new with the technology. According to Nick you will never be on friendly terms with the technology if you do not try things out – that is, if you do not tinker. Later, in the section on sabotage, we will, however, address how tinkering can also have negative consequences. In the implementation and training of employees, Nick explains that they spend a maximum of a couple of days on the "school bench", because what is best for employees is precisely to get started using the machines. Their experience and competencies must therefore be built up in practice, and you could say that Nick has enough experience to know that this is also the best way to build trust.

When Nick at one point shows a video of an industrial robot, we notice that the video is filmed very close to the robot, and we ask whether you are allowed to stand that close. He says that it is one of their own employees who filmed it, and that of course you are not allowed to, and later he refers to their close dealings with the robots as an *"habitual damage"*. By this he means that they are strongly influenced by their surroundings and established practices, making them accustomed to a specific way of working. This means that the employees a He does, however, say that new employees have a bit more respect for the robots and are very aware of how much weight they can lift and thus how dangerous they are. The employees at ElectroFlight therefore generally have a lot of trust in the robots, perhaps because they trust what they themselves have programmed and created. But the trust is not there from the start; it is something that comes with time and experience, after which they almost trust the robots too much and therefore go too close to them.

ElectroFlight makes it clear that there is a world in which robots can be helpful in a workplace without it meaning that you have to replace and fire employees. In fact, Nick makes a point of explaining how little the robot is worth in itself and how much it requires a functioning network around it in order to be useful. He says: "So it's the whole combination that does it. The robot is just a single partial combination in it." Despite this, he talks about an example where the surroundings around a work function were directly built for humans, but the robot still managed to perform the task just as well as a human. Here he explains that the robot does in fact go in and replace the human in the task, which means that this human either has to be fired or given a completely new task. It is precisely this context that situates ElectroFlight in the world today: will these employees who are replaced even want or be able to manage the new tasks?

Anomalies - ElectroFlight

When we spoke to Nick about possible anomalies for the robot, he repeated each time that if something were to happen to the robot that was unforeseen or not planned, it would simply stop and inform an employee. We asked whether it might be advantageous to make the robot able to react to these unforeseen events, to which he replied that you could indeed program it to do that, but it would cost more than it was worth, and therefore he did not think it would pay off. In addition, he saw a challenge in the fact that the robots are not smarter than the information you feed them with; that is, they do not learn anything by themselves. Their robots could not come up with doing something unpredictable, and even if it has to choose between things, the choices and decisions are still something they have fed it with. This can be seen as the unpredictability paradox, since you have to teach the robot about the unpredictable, but the point of unpredictability is that it cannot be predicted.

He does, however, believe that customers would like the robots to be able to react to these unforeseen events, but that they would not like to pay for it. According to Nick, in order for it to make sense, you have to find a specific area where there is potential in being able to react to unforeseen events. He therefore does not believe that this is something needed in all automation tasks. In the context of the robot with the blocks in the box, Nick's point is that an employee has to come and approve the box and replace it anyway, so if something they had not foreseen, or an anomaly, were to occur, then this employee could quickly fix it. We ask about a specific example where you could imagine that the robot might drop something, and whether you could then program it to pick it up. Nick believes that it is simply too expensive to train the robot in this, and that it makes far more sense to have an employee quickly come and pick it up for the robot. In addition, they also did not experience the robots engaging in abnormal behaviour, since Nick said that it only does what it is programmed to do. The following two sections will examine these observations in depth and will also include Nick's perspective on fully automated robots.

Sim2Real - ElectroFlight

Nick already told us in our initial phone call that they make extensive use of simulations, and we therefore asked whether they use them in connection with their risk assessment. Here he explains that they use the simulations to see whether the customers' wishes can be realised in practice. This is because it is quicker, easier and cheaper than having an electrician and a mechanical engineer spend several weeks building and programming it. Often customers also have specific machines or tools that they simply do not have the opportunity to get hold of, and here it is really important that they can simulate it beforehand so that they do not show up at the customer and then it does not work. Nick says directly: "*We simulate as much as possible.*" He does, however, also say that the simulation is not 100% identical with reality, but that then you just make small "touch-ups", so it is not very difficult to correct it when you have first made the simulations in advance. The customers in fact sometimes ask for the simulations in order to get a taste of the work. In addition, ElectroFlight also uses them now and then either to persuade a customer about a function or to talk a customer out of a function which is unrealistic or troublesome. Nick explains that customers often come with new ideas or wishes for the robot on an ongoing basis, and here the simulation tool is once again invaluable for them. They can easily add extra tools or functions to the program once they have the strong foundation with the robot simulated.

Because of the accuracy of the simulation tool, it also means that they can begin on the software and programming work before they have received the hardware they need. In this way they save time and resources in the development of the robots. We were very impressed by the accuracy of the simulation tool, and here Nick said that the program is basically like Excel or Word. He says that the programs are of course good, but they are only as good as the person using them and what they feed the program with. It therefore requires quite a bit of expertise to be able to use these tools in such a way that they are as relevant and accurate as they are at ElectroFlight.

We asked Nick whether they simulate humans, to which he replies that they do not, because it is not relevant for the robot. At first this sounds a bit naive, especially when he has previously explained that the robot and the human are supposed to collaborate to perform the work tasks. His point, however, is that this is because they specifically only simulate the robot's work functions and tasks and not the environment around it. In addition, this safety circle around the robot creates a field where the robot will stop completely if anything enters it. Therefore, he does not think it is relevant to simulate humans, because they would never move into where it is working anyway. Nick also says that the robot should preferably be self-operating, and therefore there is no functional point in simulating humans, since they are not to contribute to its work anyway. This doesn't mean that he thinks humans are completely excluded from the robot's area, but it means that the technical act of moving objects from one place to another, is supposed to be done without human intervention.

Sabotage and Tinkering - ElectroFlight

In this section we elaborate on aforementioned examples of tinkering, but here we set them in context of sabotage. The distinction between these two can be seen as a difference of intentions. While tinkering is often done to expands one's knowledge, the sabotage will, in our cases, mostly be done in order to stop or hinder the robot's workflow. This is by this distinction we structure the report's sections on sabotage/tinkering.

As mentioned earlier, they have experienced having to tighten their risk assessments after implementation because employees circumvent the safety measures. In these instances, Nick explains that sabotage is a well-known phenomenon for them. Here they specifically experience that employees put tape around the safety beams, trick the sensors or make special tools so they can get close to the robot. The employees thus become very creative in their sabotage and even make special tools they can use in this sabotage. Misuse, unintended use. Experimenting with the technology in this way can be called "tinkering", and it can be seen as a way of expanding practical knowledge about the technology in question. The employees sabotage the robot's safety systems and, in that way, learn about the robot's blind spots and general functionality. They also explore their own abilities by thinking beyond the existing technologies and in a way innovating on them. "Tinkering" thus helps employees expand their knowledge about the robot, the technologies and their own abilities. However, this is seen as sabotage in relation to the robot and from developers like ElectroFlight and is therefore something that must be corrected. The question is, though, whether it is even possible to avoid this "tinkering" or whether humans will always find a way to experiment with the technology. We asked Nick about the motivations behind the sabotage examples and whether he thinks it is because employees do not trust the technology. He replies: *"They probably can't fully grasp what it means. So, they think we can just… It usually just runs there, so we can just manage to walk over here."* He thus has the perception that employees do not understand the seriousness of the robot technology and therefore carry out this sabotage. In addition, he also believes that it is because they are overconfident, and that when they think they are tricking the robots, they are in reality tricking themselves. In the end, a large industrial robot will always win over a human in terms of strength, and therefore he believes it is overconfident and even dangerous of them to try to trick them.

When it comes to which age group is most inclined towards this sabotage, Nick does not immediately have the impression that one age group is more inclined than another. He does, however, explain that the older generation, which he also describes as less IT-familiar, finds it harder to be on friendly terms with the technology and that they are generally less receptive to new initiatives.

For a deeper understanding of the development of cobots and the assessment of risks and safety, as well as the collaboration with customers and their expectations, we will now present our fieldwork at the company MakeBot.

MakeBot – Making Robots

The fieldwork in this report consists of a visit to the company's premises, where two unstructured interviews were conducted with the employees Annie and Martin. After that, we took a free online safety course that they offer on their website. This course led us to contact a safety expert from the company who helped develop the course in question. Here we conducted a semi-structured interview over Zoom with the employee Rosa.

The fieldwork at MakeBot contributes to a broader industrial understanding of the field of cobots and the challenges manufacturers face in developing them. The company has extensive knowledge about what it takes to develop cobots, what considerations must be made, and what demands come from the industry and from customers. In addition, the company's risk assessment course and the interview with the safety expert have provided invaluable insight into the field of safety assessment for robots.

The findings from the fieldwork also deepen our understanding of the "language barrier" that exists between engineers and designers or customers. The fact that they have had to develop specific methods to translate customer needs into something engineers can work with shows how difficult this tension is, and how much attention and work it requires.

Introduction

MakeBot is a company that specialises in the development and production of collaborative robots (cobots), which are designed to work side by side with humans without the need for extensive safety measures. These robots are used primarily in industry to automate repetitive and precision-demanding tasks such as assembly, packing, palletising and quality control. MakeBot emphasises flexibility, user-friendliness and accessibility in their products, which has made them popular for both large and small companies wanting to implement robot technology in their production. During our visits, they asked us whether we were dealing with both software and hardware in the project, which we nodded yes to*. "The more hardware, the longer it takes*," [to develop]" said one of the male employees. They also told us that they have previously employed several anthropologists to investigate *"what matters to people?"* and how people position themselves in relation to a robot, because there are different approaches here.

Back in the meeting room Annie explains that they have an overarching vision at MakeBot, that they want to *"work with robots and not as robots"*, and here she is referring to employees doing assembly-line work that wears them out. Annie said that they considered it important to *"work with robots, not as robots*", and that robots should take over the monotonous work so that people can have a meaningful working life. They say that MakeBot aims to create robots that fit many different environments and then leave it up to the user to specialist. They explain that in the beginning MakeBot robots were hired and sold on their "playfulness", but that they have since made them more robust. There were three employees who received us on the day, one of whom, Annie, was head of UX design. In addition, there were two men, one who focuses on user experience, who had only been in the company for half a year and had no knowledge of robots before that. The other, Martin, was a visual designer who focuses on industrial design and the user's perspective rather than the engineers. Martin generally had many reflections on working with engineers, and as Annie said many times*, "they just think in a completely different way"*. They believe instead that people's competencies should be used elsewhere and that you should therefore make use of *"what they are actually good at"* rather than the simple tasks a robot can solve. It is therefore their vision to make industrial work better for the coming generation and create a *"good and meaningful working life"* that is not about wearing out your body by, for example, lifting boxes all day. Annie explains that they are very inspired by Apple and Tesla, and that they want to create a unique user experience in order to break through on the market as much as possible.

We ask who mounts the robots, and here they answer either a partner, themselves or integrators, or the company itself, which in the end is also the one that stands with the responsibility. Integrators are the group of people who help with the integration of the robot; they stand between the company that sells the robot and the company that buys the robot. An integrator can be in the company itself, in the company's "partner network" where MakeBot itself recommends them, or external.

Understanding the people who are to interact with the robots

In her presentation, Annie explains that they have developed archetypes that are meant to make communication with the engineers simpler and create a shared frame of reference. These archetypes are developed and used internally and are confidential. We therefore only describe them briefly and at an overall level, in order to preserve the essence without revealing too many details. The archetypes are based on semi-structured interviews as well as visits by employees to the companies' factories. It is through MakeBot's salespeople that they have gained contact with these companies, and they have tried to contact both large and small ones to ensure that they obtained a broad target group, from companies all over the world. In addition, it is a mix of customers who have bought or wanted to buy the robot, but also companies that merely have an idea about how one might use a robot for automation in an industrial context.

The archetypes range from those who learn by doing and are not afraid of making mistakes, to those who, in teams, use their different competences to convince management of the value of the technology. Then there are the "experts" with deep technical insight, who can work independently even when things are tricky. In addition, there are those who need help all the way through and prefer that others take care of the technical side. And finally, there are the enthusiastic advocates, who love the technology and function as its ambassadors. Martin said that he did not immediately have an overview of whether there was a predominance of certain types rather than others. The number of the different archetypes was decided so that they would have enough to represent everyone, but not so many that there were any overlaps in behaviour. The archetypes show that there can be great differences between the people who are to become part of the robots' everyday life, across the robots' work tasks. MakeBot is clearly very aware of these differences and even recognises that it is relevant for them to relate to how they interact with the robots.

Martin explains that they have not experienced cultural or organisational differences in how companies approached the robots; instead, they experienced that the companies with more money were actually less careful with the robots. As an example, they once experienced a company that did not have as much financial leeway packing one of the robot's parts in a suitcase because they were so afraid it would break. He therefore says that richer companies treat them more carelessly than companies with few resources: *"Smaller companies see them as more valuable."*

In addition to talking to customers and companies, they also visit them at their sites, in order to get an actual sense of the environment in which the robots are to be used and thus the needs that exist. They are very focused on customer involvement and are constantly sending out questionnaires. This also means that they get a lot of knowledge and feedback from their customers, and therefore Annie explains that they have chosen to say that an insight is when several customers share the experience: *"(...) it counts as an insight we can document, that there are several who have this repeated experience with our robot."*

One insight they have gained is, among other things, that the lack of haptic feedback makes it difficult for customers to "feel their way", that is, because it is a touch screen they constantly have to take their eyes off the robot in order to use the screen. There were also a number of safety-related pitfalls in this, as it meant that you did not look at the robot as much as you would if you could feel your way by means of buttons.

At MakeBot they have developed customer journeys as a way of showing how a customer moves from A to Z in a company and thus which choices and activities they undertake along the way. However, in the company they want to create a loop in which the customers actually come back and buy more robots and become an active part of their business. They have therefore designed this customer journey more like an infinity symbol, where they show how the customer ideally moves from A to A rather than stopping at Z.

As with the archetypes, the details of the customer journey are confidential, and once again we will therefore make a brief and overall presentation of it. A customer journey can, for example, start with a company manager lacking labour and placing a job ad in the newspaper, but receiving no responses, because no one wants to do monotonous work for 40 years, as Annie says. Then he suddenly discovers that there is a gap, a need for labour. This may set thoughts of automation in motion, and at this point they contact MakeBot. After implementation, the user then experiences lots of problems for which they need help from MakeBot. We ask which kinds of problems this might be. Annie says that it could be anything and asks Martin for an example; he replies, *"for example this thing about how sensitive the robot should be."*

We ask how often errors occur. They say that it varies a great deal, but they try to prevent further errors through support, and that they do not always know whether errors occur, because the companies do not have the competencies to handle them. They emphasise that *"training is a big part of MakeBot's customer journey."* So, the user has to be ready to take in new knowledge.

After a successful implementation, the company begins to automate several similar tasks. Over time they develop a more creative approach and explore what else can be automated. Finally, and hopefully, the customer becomes an advocate for the technology and recommends MakeBot to other companies.

For MakeBot, a successful customer journey is precisely that the customer reaches this stage where they recommend their robots to other companies. With the different phases, they can see in which exact phase a customer becomes dissatisfied. They experience that if customers become too skilled at handling the robots and do not need help from them at MakeBot, then they "lose" the customers because they do not hear from them. This does not mean, however, that they do not want to make their customers skilled at using the robots, because they still place great emphasis on training the customers as much as possible. Annie explains that the need for robots in fact arises when an employee "breaks down" and shows that their task is too grinding to be performed by a human, and then this employee is instead given the opportunity to do something else. In contrast to humans, a robot can work 24 hours a day, but this is most relevant for larger companies that need this form of labour.

Cobots from MakeBot

The robots from MakeBot, also called collaborative robots or "cobots", are developed to be able to work directly together with humans. In contrast to traditional industrial robots, which often stand behind fences and require special programming equipment, cobots from MakeBot are designed to stop automatically if they encounter an obstacle, which makes it possible for humans and robots to work side by side in the same work area without large safety barriers. All the robots are built as flexible robot arms, which gives them a mobility akin to a human arm. This means that they can perform many different kinds of tasks, ranging from assembling small components to lifting and moving heavy items. The robots can be mounted on the floor, wall or ceiling and can be moved between different workstations. They are programmed via a user-friendly touch screen, where the operator can guide the robot arm through the movements it is to perform and save them as a task, so in this way advanced technical knowledge is not required to get them up and running.

MakeBot offers several models with different strengths and ranges, so they can be adapted to different types of production. The smallest models can handle tasks with small items and light lifting, while the largest models can move items of up to 30 kilos and reach almost two metres.

They sell a kind of “package” that consists of a control box, a robot arm and an “iPad”. In the showroom we saw a demo of a robot arm that was to palletise. Here boxes stood neatly and precisely on a pallet, which the robot then had to lift from one pallet to a new one, and here they programme a palletising pattern, “and then it can just palletise”. After that we saw a demonstration of various AGVs in different sizes, which also work closely together with other robots, for example robot arms. They have chosen not to specialise their robots too much and instead leave it up to the individual customers if they need a very specialised tool on their robot. In that way they can also reach as many different “special” customers at once.

Annie explains that they want to make the robot of the future, and that the first robots they developed were actually perceived as too “playful”. She explains:

> *“It’s also a bit compared to Fisher-Price, and that means it’s right down at toy level. (…) But it has been really good for us in the beginning to be there, because that’s what people have been able to buy into. They’ve said, that looks easy, that thing with robots, it’s just a toy.”*

This helped them gain a lot of new customers who did not have much experience with robots, which of course was very good for their turnover. At the same time, however, this was not how they wanted to be perceived; they would rather be taken more seriously as a brand and not have their robots perceived as toys. In addition, she explains that because they were up against some large competitors in industrial automation, they had to make their robots more robust and move away from this perception as toys. Despite moving towards a more robust robot, they still wanted to hold on to this playfulness, because they had experienced that it made people more inclined to trust the robot and made it easier to use. Annie explains that it is important for them that their robots are easy to use, and that they must be minimalist and reliable.

The control system for the robot is designed as a tablet, and Martin explains that this is done with a view to making the user interface more intuitive and user-friendly. He says that most people by now are used to using tablets, and that there is a high level of recognisability in tablets, which makes it more likely that people will be able to work out how to use it. Martin describes their new user interface as follows: *“What you meet today is closer to what you have in your everyday life on your phone than what it used to be.”* As mentioned earlier, they say that they are very inspired by Apple, and this is probably because they want to create something that is recognisable and resembles something you already use in your everyday life.

A shared language

The different archetypes and customer journeys they have developed have been a way for the designers to communicate better with the engineers, Annie explains. She says that engineers often need something more concrete and context-based in order to understand what the designers mean:

"And one of the things engineers need, because they typically think with a slightly different half of the brain than designers do, is to have some concrete material to work from."

These various framework tools they have developed help them establish a shared language with the engineers, instead of misunderstanding each other and becoming frustrated. Annie also says that it has sometimes been challenging for the engineers to work with them, because they found it difficult to understand what the designers actually meant. More precisely, Annie says that *"they need to have things they can put two lines under"* (i.e. clear, fixed conclusions).

A colleague explains that you can clearly sense this in the way MakeBot's engineers now talk and think:

"Now that we work as engineers and have worked with some of your engineers, it's obvious they've taken it on board in the way they work – they talk about the archetypes and about who we're designing for and so on."

He explains that this actually makes it easier to work with the engineers, and that you can sense a shift happening within engineering as a profession.

Martin says that the company has moved away from their old interface, because you could clearly tell it had been built by engineers for engineers, and it was not very intuitive for non-technical staff. The way you could feel this "engineer-to-engineer" design was that the user was exposed to an enormous amount of information all at once – the interface was marked by a kind of information overload. Martin describes the problem like this:

"And a robot engineer can cope with that better than a non-robot engineer. (...) The information is too overwhelming when you see it all at once. (...) So, what we've done is say that some information is more important than other information at a given point in time."

They have therefore chosen which information is necessary for operators to see, and which is not. This can, however, be problematic, because in this way they are filtering out information that workers might actually benefit from having. We ask Martin whether it is only MakeBot that decides which information workers get, or whether the workers themselves have had any say in it. He answers that it is MakeBot who decides which information is most important to highlight.

This can be seen as a form of *techno regulation*, since the developers are effectively deciding (and thereby regulating) which information workers should and should not have access to. It is tied to the concept of information overload.

Martin points out that on their previous user interface there were about 20 different options and arrows that operators could press, and that this was far too overwhelming. He adds that an engineer is in a much better position to understand all these options, because they are used to working with such interfaces and are “trained in it”. At the same time, by narrowing down the number of options in this way, they also take away some agency from the workers.

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

Safety - MakeBot

While we are waiting for the green light to see the showroom, we talk with Martin and Annie about, among other things, risk assessment. They say that *“a versatile environment makes risk assessment difficult”,* because every time the cobot is moved around, it has to be safety-checked again. Small and medium-sized companies often do not have the same tasks every day, so here the cobot is moved around on a trolley to other tasks. It is used for tasks that are difficult (in terms of precision) and physically demanding, and it is a waste of resources if the cobot cannot take care of several tasks in small to medium-sized companies.

They also explain that “end tools”, e.g. grippers, are bought externally, and these of course also affect the risk assessment. We ask whether the companies prioritise safety, and they say that there is great interest from companies in getting CE marking, even though it rarely gets checked by authorities.

They explain that the robot itself is not dangerous; Like Nick they underline it is the tool it is given – the end tool – that can make it dangerous. The robot’s own surface has no sharp edges or other hazardous features. They have deliberately made the design as minimalist as possible to increase safety: the very clean and simple shape makes it easier to see what is moving when, and you can place your hands almost anywhere on the robot without getting hurt. As Martin says, *“The fewer edges, the safer it is.”* In visits to companies, they also found that this kind of design was in demand. On a very practical level, it also makes the robot easier to keep clean, especially because it often has to stand in warehouses where there is a lot of dust.

We also saw other robot arms that were equipped with safety scanners. If a certain threshold is exceeded – for example, if a person comes too close – the robot slows down. The rotating robots use power-and-force limiting to ensure that they do not injure people or damage objects. When we were “testing” one of the robot arms, one of us put our hand in front of or alongside the robot, and it did not stop, so we asked why. They said they had deliberately chosen not to make it “too sensitive”, so it wouldn’t stop at the slightest touch, which is an advantage in close human-robot collaboration (HRC). Here there is a potential trade-off between safety and sensitivity: you don’t want the robot to be so sensitive that it constantly brakes and halts production. One of the people presenting in the showroom remarked that the fact the robot does not stop its movement to the left, even though my hand is in the way, doesn’t actually harm my hand. However, it still feels very transgressive that it continues regardless of my hand being there, and that precisely illustrates the difference between *being* physically safe and *feeling* safe.

They themselves say that there is a trade-off between sensitivity and reliability. When we ask what they mean by reliability, they explain that reliability is relative, and that there are acceptable and unacceptable errors. This is also relative, because a robot “failing” and stopping production is, from a safety perspective, not an unacceptable error – but it is unacceptable if it fails in a way that injures a person. For a production manager, delays are not an acceptable error, and this illustrates that different user groups accept different trade-offs.

Safety course

The people present were not safety experts in the sense of being hired by companies to install robots in full compliance with safety regulations. But MakeBot offers a free online risk-assessment course on their website, which we took in order to better understand the kinds of safety considerations that are made. The course is based on ISO standards, including ISO 12100, which is described on the ISO website as: *"[...] specifies basic terminology, principles and a methodology for achieving safety in the design of machinery. It specifies principles of risk assessment and risk reduction to help designers in achieving this objective."* As part of this, a table is introduced that starts from identifying a set of risks and then assessing the severity level of each individual risk.

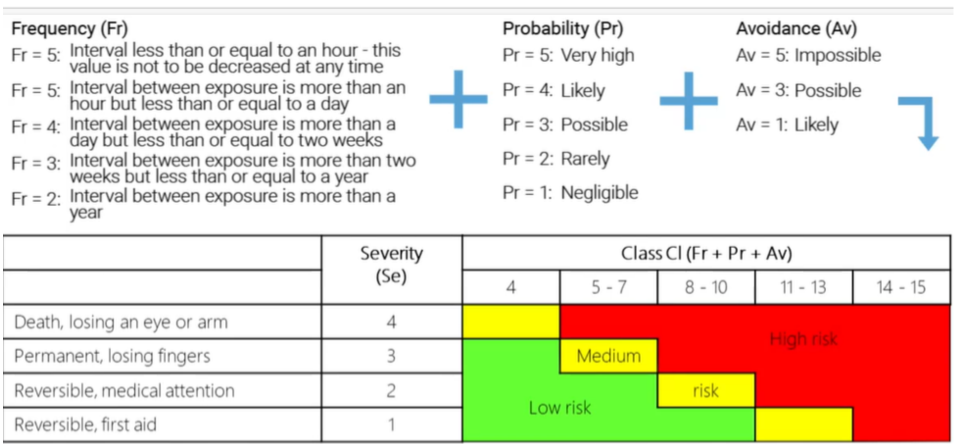

The lowest level of severity corresponds to something that can be fixed with first aid – for example a scratch or a minor wound – while the highest level of severity corresponds to death or "losing an eye or an arm". After that, you assess how often it might happen, the likelihood that it will happen, and how avoidable it is. These three numbers are then added together, and you can place the risk in the table to see whether it is rated as low, medium or high risk. The risks that end up in the red area are the ones where it is most important to put safety measures in place.

It was on the basis of this course that we contacted Rosa, who helped develop it. Rosa explains that they use ISO standard 10218, an international standard that sets safety requirements for industrial robots and their integration into the work environment. It consists of two parts: one for robot design (part 1) and one for safe robot systems and installations (part 2), and its aim is to protect humans and ensure correct risk assessment. The standards focus on how to ensure the safe use of robots in practice, with an emphasis on implementation.

Rosa explains that robot development has changed enormously over the years, and that robots have become far more electrical where they used to be hydraulic. In addition, she notes that the hydraulic robots of the past were much more dangerous, because their components were much larger and heavier than the parts used today. Previously, the only way to protect humans from robots was to separate them physically, which made human–robot interaction and cobots very difficult. Rosa explains that one of the biggest developments in robotics has therefore been the ability to protect human safety without separating people from robots:

"So back then, what really the only thing we could do was try to separate people and robots because that was the only way to protect them. And what's happened over time is that, well, you know, we have gotten that capability to provide protection for people without necessarily them being separated."

She also says this is exactly the area where MakeBot really thrives with its robots within the market, but she adds that customers' wishes cannot always be fulfilled for safety reasons.

In relation to this development, we ask whether this means that standards are actually developed "after the fact". She replies that they are not, but that they follow what is considered "state of the art". This means you cannot predict what will be needed in the future; you can only respond to the needs that exist here and now, with the technology currently available. She also notes that the things she approved years ago, when she sat on the standards committee, are things she would most likely not approve today, because they no longer make sense. That does not mean she was ignorant or careless at the time – only that there has been so much development in the meantime that it is necessary to rethink and revisit earlier decisions.

We explain that after taking their risk-assessment course, we realised that risk assessment is never completely objective but actually depends on who is doing the assessment. She agrees but says that everyone is working from a shared logic around risk assessment precisely because you cannot avoid people being subjective. This shared logic is based on the principle that the most important thing is to make the robot safe for humans – to avoid harming people. She also says that this logic helps create a kind of hierarchy of what has to be prioritised first in risk assessment. For example, you would always prioritise a situation where someone's head could be injured over a situation where someone might get a paper cut:

"And it could be more difficult to tackle that and easier to tackle getting papercuts. You have to put your attention to where you could have the worst harm."

There is therefore an inherent prioritisation decision in risk mitigation: the risk that can cause the greatest harm must be addressed before the smaller ones.

We ask Rosa whether she has ever been injured by a robot. She replies that she does not develop robots, she only does maintenance work, so she has never been in a situation where she could be hurt. She has been a safety consultant, where her job was to ensure that robots were safe, but again she has only been an observer who tells operators what they should do. Rosa summarises this by saying:

"The people who are at risk are the people that are using them."

Despite the confidence in the robots' safety measures described above, it is still the users of the robots who are most at risk. Her point is that the more you are exposed to hazards, the higher the probability that you will eventually be injured.

The standards are constantly evolving because the industry keeps getting smarter; the standards written in the past are therefore not as robust as those developed today. There is a kind of chicken-and-egg problem in robotics when it comes to safety. Before you can develop a standard, you have to be sure that is relevant technology on the market that warrants development of such standards. For example, it would not make sense to sit down and write requirements and standards for flying cars, because they are not yet on the market, and it is therefore impossible to imagine all the hazards involved.

The positive side in the robotics industry, she says, is that most actors are willing to make big changes in order to ensure better performance: *"(…) performance improves, safety improves."* We ask for her opinion about the standards – are they too rigid and should they be loosened? She replies: *"I'll tell you, anybody who's involved with using robots doesn't want the safety standards. They say, I just want to do what I want to do."*

She adds that there *are* people who care a lot about safety, but overall, as she puts it, *"we can't be enabling stupid*": safety exists to prevent the people who interact with robots from doing something foolish. She explains that if an integrator or manager does not have a sufficient understanding of hazards and consequences, they can make decisions about robots that, in the worst case, harm other people. For that reason, safety has to be ensured so that people do not harm themselves or others by making careless or reckless decisions. In the end, they do everything they can to build in safety – but users also have to *use* those safety features and not try to bypass them, which blurs the lines between user responsibility and the manufacture's. This question of responsibility is prominent in multiple fieldworks, and will therefore be further explored in Sub-report V.

Rosa explains that the biggest safety issue for mobile robots is: what happens if they lose power? At MakeBot, which mainly develops robot arms, they handle this differently than with walking robots. Their robots have pins that interlock to brake the robot's movement. The safety challenge is not as big for MakeBot, because if one of their robot arms loses power, it simply drifts slowly downwards. Sometimes they activate the pins deliberately to stop it immediately if needed – a kind of emergency stop. This kind of technology is harder to implement in "legged" robots, that is, robots walking on legs, because there are many more joints and movements that must be secured if they stop. And a two-legged robot has much less inherent stability than a robot arm.

We explain RoboSAPIENS' goal of developing safe self-adaptive robots, and she immediately says that this will not change the existing safety measures on the robots. She elaborates that even if the robot can choose new paths or avoid things it has not avoided before, it will still be operating within the same fundamental safety framework. In other words, the robot may become autonomous and self-adaptive *within* that predefined safety envelope. Rosa actually states quite directly that they do not care whether the robots are autonomous, because that has nothing to do with safety.

Trust - MakeBot

As mentioned earlier, their first robots were perceived as toys, and although this was not how the company wanted to be seen, it did have some positive effects. Annie explains that it actually meant some people were less afraid of the robots and dared to approach them more. She compares it to toy brands like LEGO and Fisher-Price, and says people would even say things like, "Robots look easy – it's just like a toy." In these cases, people's trust in the robot grows because it resembles toys they already know and trust. The fact that the robots carry a degree of familiarity makes customers more inclined to trust them, even if they don't understand the technology behind them.

Martin tells us that the first time he had to interact with the robot interface, he didn't dare touch anything on the screen because he was afraid, he might break something. He explains that his lack of technical knowledge about the robot meant he didn't know what would happen when he pressed the buttons, and so he was afraid of doing something wrong:

"So, my first experience with it was that we didn't dare touch it. We were afraid something would break if we started."

He received no onboarding or training but was simply told to go and "play" with the robot and try pushing buttons. Martin did not feel safe with this kind of "play" because of his limited technical knowledge. Experimenting with technology in this way can be called *tinkering*, and it can be seen as a way of expanding one's practical understanding of a given technology. Martin is encouraged to tinker, and the assumption is that this will expand his practical knowledge. However, he is not interested in doing so, because he does not trust what it means to press the various buttons. His trust in the robot is therefore constrained by his knowledge of it, and he is unwilling to engage in exploratory work with the robot without sufficient onboarding and training.

In contrast to himself, Martin believes that engineers benefit more from tinkering, because they have a basic understanding of the robot's technical components. He describes the robot's movement as being based on *"an invisible coordinate system"* and says that, given his position and educational background, he does not have knowledge of this – but the engineers do. Here he is most likely referring to the earlier notion of the "black box". It is this lack of insight into the "black box" that produces mistrust: he knows something is going on inside it, but he simply doesn't know what that something means. He also stresses that it is all about knowledge; there is nothing magical or out of his hands about it, it is just knowledge he has not yet acquired because he has not been properly trained.

In connection with the archetypes, he explains that if one of the customer types is confronted with too much information, they pull back because it becomes overwhelming and unmanageable:

"If they encounter too much information or something, they back off."

This align with what we have heard before about information overload – too much information is not a help but a nuisance. Even though he doesn't explicitly identify himself with this customer type, he still compares it to his own approach to tinkering with the user interface. He was confronted with too much information that he did not know how to relate to, and so he backed away.

Anomalies - MakeBot

When we further ask Rosa about robots becoming more autonomous, she explains:

"Well, they can be as autonomous as they want. It's just that the safety that's associated with their operation cannot be affected by evolving. It cannot itself evolve in its behaviour. How it

functions must remain completely deterministic and understood. So, and that's a requirement of the safety standards. At some point in time, more might be allowed, but it is not at this time."

What Rosa is saying here is that the robots are fully deterministic and cannot do anything they have not been programmed to do. In doing so, she indirectly rejects the whole idea of anomalies: in her view, anomalous behaviour is not a relevant concept *yet*, because under current safety standards the robot's behaviour is not supposed to evolve in unpredictable ways.

Sim2Real - MakeBot

During a demonstration of one of the robots, an employee explains that it uses its sensors to navigate. We ask where these sensors are placed, and they tell us they are positioned "all over", all the way around, and describe the robot as "almost trying to dance". Up to this point, they have not said anything about simulations, and the interface does not display a simulation of the arm, but rather a coordinate system. The robot thus uses its sensors to sense and respond to obstacles in the real environment

Sabotage and Tinkering - MakeBot

The formulation *"reasonably foreseeable misuse"* is, Rosa explains, taken directly from the *Machinery Directive*, and the idea is that you first identify all the ways in which the robot *could* be used "incorrectly". These possible misuses must then be tested to see whether, if the robot is used in that way, it still operates within its safety limits. If this misuse remains within the safety envelope, nothing further is done, and it is categorised as "unintended use". This means that something that originally counts as a "wrong" way of using of the robot, can be reclassified as merely 'unintended use', if it does not pose a danger to the robot or its surroundings.

If, however, the misuse falls outside the safety limits, the manufacturer must record that type of misuse in the product documentation and explicitly state that the robot must not be used in that way. So where does sabotage fit into this categorisation? We can assume that sabotage would be when someone *knowingly* uses the product in one of the "wrong" ways that have been clearly marked as prohibited. But what about all the other forms of "wrong" or "unintended" uses of the robot?

Rosa explains that they have a test department that does everything it can to try things out and to try to break things, in order to see where failures may occur. In this way they capture many forms of misuse in advance.

However, it is one thing to write down what you may and may not do with the robot, and another thing altogether to stop people from doing it anyway. Rosa tells us that at one point they discovered that a beta version they had sent out was not safe, so they wrote to the buyers and told them they either had to return it or destroy it because it was dangerous. None of the buyers sent the robot back, and Rosa says that the only thing they could do was to continue informing them that it was dangerous and that they should not use it:

"I don't think anybody did return them. They said, well, what do you mean? This is handy to have. I'm going to keep using it. (...) So, but we've told them everything. There's not anything else we can do. We can't send the police in and confiscate. We can't do it."

She also explains that sometimes a safety issue can be solved with a software update, and then they write to customers to tell them they must update their robots – but it is rare that everyone actually does it. It is therefore clear that customers do not necessarily understand the seriousness of using robots incorrectly or using them when they are unsafe, and that they sometimes do it anyway, for one reason or another.

Whether this should be seen as sabotage depends to a large extent on the customers' motives and intentions: are their actions driven by overconfidence, by ignorance, or by something more deliberate? It is not obvious that customers themselves recognise that what they are doing with the robot might be perceived as sabotage. Rosa says that, from their perspective, the important thing is to keep end-users continuously informed and to ensure that they themselves meet the standards and requirements – beyond that, customers' actions are out of their hands.

The next field site we shall present is a company called GreenTech, which works with refurbishment of electronic waste. Here, they are interested in implementing a robot to help automate their initial waste-sorting process.

GreenTech – Refurbishing Robots

The fieldwork at GreenTech provides an example of a company that sees possibilities for implementing a robot and automating tasks. It thus contributes to an understanding of which considerations are relevant *before* purchasing such a robot, where economic and social aspects play a major role. Management sees many opportunities for automation but does not expect to have to carry out major rounds of layoffs. Since they have not yet acquired the robot, we cannot say whether these hopes or expectations will turn out to be accurate.

Introduction

GreenTech is located in a large European city and, although their premises are actually quite spacious, the place looks fairly modest from the outside. The contact to GreenTech was established at a robotics conference, where we ended up talking to one of the company's founders, Christopher. He told us that his company was preparing to get a robot.

We visited the company twice, right after the contact was initiated, where we were given a tour and heard about their initial thoughts on getting a robot. A year later we contacted them again to hear how far they had come in the process. Once more we were shown around the facilities, while Christopher explained what had changed since our last visit.

We began the tour at the very beginning of the process: the goods reception, where all incoming e-waste arrives. From there we followed, step by step, the journey the material takes if it is sorted for reuse. First, the waste is assessed for how reusable it is. Anything classified as "non-reusable" is separated out and dismantled. As we walked through the different stages, Christopher pointed out all the places where he could imagine a robot taking over tasks, even though most of these areas are not yet being automated in practice.

Christopher also explained GreenTech's approach to refurbishment. Rather than only accepting voluntarily delivered equipment, they acquire electronic waste in two main ways:

1. by purchasing it from large companies and recycling centres, and
2. through agreements with companies and municipalities, who give them their electronic waste free of charge.

The electronics they receive then move into several possible pathways. If a device has sufficient resale value, it is refurbished and sold on—either on the Danish market or abroad, depending on demand and price. If it is deemed not worth the time and cost of refurbishment, it is dismantled; functional components are extracted and sold, while the remaining material is prepared for smelting. Certain products, such as worn-out PC monitors, are particularly problematic: they neither generate profit nor can easily be taken apart into saleable components and thus represent an economic burden.

Reusable equipment first passes through a "cleaning" stage, where stickers and other surface markings are removed. It then proceeds to technical processing, where units are tested for faults and deficiencies. Faulty devices are identified using error codes and component measurements on electronic diagnostic stations. Defective modules—such as screens, batteries or motherboards—are replaced with reused or factory-refurbished spare parts. After repair, the units undergo functional testing (functionality, performance and quality assurance) and are finally packaged for resale.

GreenTech's largest supplier is a recycling company that collects waste from recycling centres and resells it to them. GreenTech then handles cleaning, serial number registration and initial sorting. Around 2% of this stock is sold directly to consumers; the remainder is exported, primarily to Eastern Europe. Anything that GreenTech cannot reasonably reuse or refurbish is sent back to the recycling company for breakdown into materials. A favourable agreement allows the electronic waste that GreenTech itself collects to be returned for a relatively low fee. When GreenTech picks up equipment directly from companies, older and heavily worn hardware often comes along as well; this is typically not suitable for reuse and is therefore handed over to the recycling company.

The environmental motivation behind this model is strong. The CO_2 savings from reuse, compared to full material recycling, are so substantial that it can be economically worthwhile to transport IT equipment over long distances—for example from Norway—rather than sending it straight to the shredder. GreenTech's practices are thus both economically and ecologically framed.

Smaller, everyday practices also reflect this orientation toward reuse. GreenTech produces its own packaging filler from old cardboard boxes using a simple "shredder". They collect surplus cardboard from customers, neighbours and suppliers, process it into cushioning material, and use it to pack outgoing goods for customers "with green glasses on" who appreciate visible reuse. This approach is not only environmentally appealing; it is also significantly cheaper than buying new packing materials.

However, basic logistics remain a major challenge. The company handles an extremely large and varied assortment of products. Some items, such as stage lighting, are particularly difficult to pack safely and thus become disproportionately expensive to handle if one wants to avoid damage. For most common product types, GreenTech has developed standardised packing solutions tailored to typical forms and dimensions, which helps keep both costs and handling efforts down. But because they receive almost every imaginable type of electronics, new packing and transport problems continually emerge.

Electronics that cannot be reused are dismantled on disassembly lines. Here materials are separated into metals (such as copper, gold and aluminium), plastics, glass and hazardous substances like battery acid or mercury. Metal fractions are melted down or refined in collaboration with certified smelters, while plastic and glass parts are sent to specialised recycling plants to be processed into new raw materials. All these technical processes are continuously monitored through an ERP system (for logistics and inventory management) and an MES system (for production monitoring, which stands for Manufacturing Execution System), with the aim of maximising the reuse rate and ensuring traceability.

In summary, GreenTech combines advanced e-waste handling and diagnostics with tightly organised logistics and process management systems. The goal is to maximise reuse, safeguard data security and support as closed material loops as possible within the constraints of a highly heterogeneous waste stream.

The dream of a robot

The company does not yet have a robot, but they want to implement one. However, even after a year they have still not fulfilled the dream.

In the sorting hall, preparations are being made for the implementation of a 12-metre conveyor with five robot arms equipped with cameras and AI. The vision is to automate the fractioning of tech waste into up to ten categories and allow the robots to recognise model numbers, look up specs via online searches, and even take pictures of defective screens during a preliminary inspection prior to the auction. The point of the selection is to distinguish between technology waste that is potentially valuable enough to be reused, or if it needs to be sent to separation and then destruction.

The project promises to be able to increase throughput by a factor of five without hiring more staff – but this is also where doubt begins: the system is delivered *"with an empty AI brain",* we are told, and all knowledge must be encoded manually by employees. They have also considered using robots in other parts of the company and are generally very open to expanding through robotics.

The robot they want in the sorting department would initially just be a camera equipped with AI, which should be able to sort according to their internal sorting guide. Only later would it perhaps be the plan to add robot arms. The system is to be trained by their staff, who will have to step in whenever it scans something it does not know where to send. Christopher described it like this:

"But it will need to be fed with information. All the time. So that camera has to constantly see what is going across the belt. And then the employees have to enter whether it should be green or red, right?"

The information it is to be fed with is their internal sorting guide: a series of pictures of computers, phones, screens and other electronics, each with brief descriptions of what should be sorted out, and so on. For example, computers with an old operating system, or screens without a stand.

They already have a hoist system that brings containers of electronic waste in and dumps them onto a belt, where staff then start sorting. The current hoist was designed by a company that also produced the plans for the sorting robot they would like to have. However, they assessed that it was too expensive for their current volumes of e-waste, and that it is only something they will implement once they handle larger quantities.

The hoist uses safety sensors, but when we observed the workers starting the sorting work, it did not look as if they were taking any special safety precautions, even though the electronics in the container weighed over 200 kg in total, and large boxes were just tumbling out onto the belt.

The warehouse with prepared goods is a large room with rows of shelves where all the electronics is stored. They have also discussed the possibility of using robots here, but for now employees walk around and pick every single item that needs to be sent to a customer. In addition, they have a long conveyor installed that runs along the rows of shelves, on which employees can place items to be shipped, and then they are transported down to packing. It is not automated; it is only there as an aid. We did not see it being used while we were there.

Internal knowledge

This empty AI brain that the robot is supposed to have is because it must first be trained on the company's own sorting guides. Employees will then be able to help it whenever it scans something it does not know how to sort, and only once it is well trained would robot arms be added. Christopher also said that they have no plans to fire any of their current employees after the robot is implemented.

However, he also said that if they were in a situation where they needed to hire 5–10 new employees to keep up with the volume, and then the robot came in and helped, they might be forced to lay some people off – but that the intention is not to fire any of their current staff.

During the visit we focused on two areas: their workshop, where technicians inspect, test and document faults on equipment (an experienced technician handles around 55 desktop computers a day, while high-end devices under five years old are repaired), and their manual sorting process, which will soon be supplemented by a robot-based solution.

They have built an internal error-code system over the course of half a year, but with the very varied mix of products (from mobile phones and printers to power tools and household robots) and the relatively small batch sizes, automation with robots is currently not practically feasible. With around 40 technicians on the floor, however, a robot solution could significantly reduce wage costs. Christopher's expectation is that the future automation cell will only really be "run in" after about six months.

Products suitable for reuse go to a data-erasure process: hard drives and flash media undergo certified overwriting or physical destruction of data. Components are then physically cleaned (for example with ultrasonic cleaning, dust removal and degreasing) and electrically tested on specialised test benches. Hard drives are thoroughly wiped with certified equipment before being reinstalled in computers; ordinary "delete" is not enough to guarantee that data cannot be restored.

Solid-state drives (SSD) are preferred, as they are faster, cheaper and much quieter during read/write and erase processes; the technology has developed considerably over the last 15 years. Data deletion on mobile devices (phones, tablets) is time-consuming: some models can be partially automated, but most require manual work via factory resets. The technician connects a PC, runs special software, and often has to follow a unique sequence of restart presses and menu navigation – instructions that may only exist on YouTube for that specific model. Each operator typically resets 150–200 phones per day, while other types of devices are handled by specialists.

Over time, they have built up a bank of procedures for different models, but it is surprisingly difficult to make the system fully searchable and to produce clear, user-friendly instructions. Electronics that cannot be reused or resold is dismantled into metal components (circuit boards, aluminium, etc.) for traditional scrap recycling.

If robots are implemented without opening new role opportunities for these groups (for example by training them in robot maintenance or managing data-erasure processes), there is a risk of pushing precisely the most vulnerable workers out of the labour market. Ethically, a parallel social strategy should be planned that actively retrains and reallocates current employees so that technology strengthens inclusion rather than undermining it.

Over time, a strong dependence on the robots' teach-in could create a situation where all sorting rules and technical know-how are encoded in the AI model rather than in employees' heads. From an ethical standpoint, the company should ensure transparency in AI logics, carry out ongoing audits of the sorting algorithms, and maintain a manual fallback option, so that critical knowledge is not lost if the technology fails or changes.

Christopher told us that they have no plans to fire their current employees. At the same time, he expressed many wishes and kept pointing out different places in the company where, in his view, tasks could be optimised with robots or other automation. This gave the impression that his dreams of implementing robots in multiple parts of the company did not entirely align with his ideals of not laying off any existing staff.

This impression was reinforced in the conversation with the sorting department manager, who said he did not look forward to possibly having to fire some of his employees when the robot is introduced.

"Depending on how efficient it is, it may well be that some employees will have to leave the shop. And of course we would be sorry about that," the sorting manager said, expressing some of his concerns about the robot. It sounded to us as if he also thought it might affect current employees, and not only hypothetical future hires. However, he also said that he did

not think this would be necessary in the near future – perhaps both because he does not expect to see the robot any time soon, and because he imagines it will take a while before the robot reaches a point where there is no longer a need for so many employees.

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 - GreenTech

The introduction of robots could change the safety landscape on several levels in the company, but as the robot is not yet fully implemented this is not yet known. Here are some speculative reflections:

Physical safety (safety) concerns how the robot arm could cause harm. For example, if an operator comes too close to a gripper arm without proper light curtains or guarding, or if a rapid change of gripper tool triggers an unexpected movement. In practice, you have to relate to three actors: the robot itself, the people around it, and the physical surroundings (belt, pallets, floor markings).

Even after they got the hoist system, they have no floor markings for safety – the markings that exist only indicate where pallets should be placed – but there are cages around the hoist itself.

Perceived safety would be the operators' feeling of safety, which may be low at first, especially if they do not understand emergency stop procedures or perceive the robot as unpredictable, even if the objective risks (physical safety) are controlled through speed limitations and sensor systems. In addition, it is relevant to note that some of their employees may find it harder to adapt to various changes.

Companies will most likely follow current standards and regulations such as ISO 10218 for industrial robots, ISO/TS 15066 for collaborative robots and relevant CE-marking requirements, so that the design of safety zones, installation of emergency stops and ongoing risk analyses all meet regulatory requirements. But whether this will create real, as well as perceived safety, across all employees is unknown.

Overall, there was not much safety around the belt and the hoist, even though the hoist was fenced in. This might indicate a work culture where employees feel they know the work so well that safety is not something they actively think about. There was also a sign warning against placing hands at the start of the belt to avoid getting fingers caught, but this too seemed to be ignored, as we saw an employee put their hand in to help the electronics down onto the belt.

It could be relevant to investigate this further using Stian Antonsen's theory on the relationship between safety and culture in the workplace safety (in his case on ships). In his view the workplace culture is defining what safety amounts to in a workplace. And GreenTech is a special type of workplace that, among other things, employs refugees and people with mental and motor challenges, which may contribute to a different kind of culture – one that is constantly up for negotiation and in flux – which is probably not what ISO standards are written with in mind.

If we look at the culture we observed during our visit, it seems to indicate that employees place a great deal of trust in themselves and ignore potential risks. In their use of the hoist, for example, it was clear that employees rely on their own judgment and therefore feel able to ignore warnings, such as the risk of getting fingers caught at the start of the belt. They also stood very close to the point where the electronics were tipped out onto the belt, even though the contents of the container consisted of large, heavy components.

Trust - GreenTech

We assume that trust in the robot is a dynamic process, where operators' expectations and experiences are constantly being adjusted. At the beginning, there may be a tendency either to trust the system too little – for example by stopping the belt at every minor uncertainty and thereby interrupting the flow – or to trust the AI recognition too much and thus overlook errors when the robot blindly relies on its previously learned rules.

Trust is built through relationships and repeated successful experiences: every time the robot correctly identifies and handles a unit without human intervention, the operator's trust presumably grows, while errors or sudden "freeze" situations create mistrust and lead to more close monitoring of the system.

The ongoing interaction between robot and human, and the way successes and failures are shared within the team, could be crucial in a trust trajectory that only really takes root once both technology and users have adjusted their behaviour in step with one another.

In addition, the specific employees they have may also matter for trust in the robot. For example, we were told that when they first got the conveyor with the hoist, one employee was unhappy with it from the start, both because it was a change to what he was used to and because it did not work perfectly in the beginning. It therefore took him longer to get used to it and to start trusting it. If they later introduce a robot that will take even longer to fine-tune for its tasks than a simple hoist and belt, this could lead to an even longer adjustment period.

Anomalies - GreenTech

Due to observation in fieldwork, we began to question if what the engineers referred to as 'anomalies' always had negative consequences. Of course, seen from the robotic systems, there is nothing negative just as there is nothing positive. However, robots react to anomalies by seeing them as something negative (in the sense it does not know what to do about it) to be remedied to receive balance. A negative anomaly could be the robot reading a barcode incorrectly or classifying a product as a category that does not exist in the database and therefore stopping the belt for operator intervention. If the AI model is trained with incorrect teach-in examples (for instance, a wrong category), it may systematically repeat the same mistake many times – a hidden error that is only discovered later. However, what is seen by the robot as negative anomalies may be seen from the human perspective as something to be welcomed.

They already experience situations where something that should have been sent for material recycling is instead sent to refurbishment, and they only discover this later in production or in the warehouse with goods ready for sale. Some of the employees we spoke with expect that similar things will happen with the robot. However, employees may also find ways to 'push' the robot in the right direction and even if this 'push' is also seen as a negative anomaly from the robot's perspective, these human interventions is what makes the workflow/improves the workflow.

Sim2Real - GreenTech

There is nothing to suggest that Sim2Real methods are part of the current setup. Instead of training the AI in a virtual simulation (Sim) that is then transferred to the physical conveyor (Real), the entire "teach-in" phase takes place directly in the real robot cell.

The system is delivered with an "empty AI brain" and learns solely through interaction with real items on the conveyor. Every time an unknown unit appears, the robot stops and asks the operator for the category, and this feedback is then stored directly in the model. The constantly changing products (old laptops, specialised tools, smartphones from many manufacturers) make it difficult to build a representative simulated world that covers all variations. The primary data source is manual teach-in.

This slow, human-driven learning (6–12 months) functions as both training and validation, so there is no separate virtual test phase. Setting up a faithful digital twin with accurate physics simulation, camera models and robot collision checks would require extensive development and maintenance and would not be of much help here. The closest they come to simulation is their sorting guide, which they plan to use to train the robot – but they are also aware that this will not be enough, and they expect that the robot will need to ask for help.

In this context, it would make sense to train this "AI brain" on synthetic data, as this could speed up the robot's learning and reduce errors. Synthetic data is data created artificially using various computer-based models that generate data from statistics or real observations.

As one of the informants mentioned, he himself finds it difficult to recognise all iPhone models, and he often has to look at their connectors to know exactly which model it is. This is something you could train the AI on using synthetically generated images of all phone models, computer models and other types of technology that appear. With such data you could also train it to recognise stickers, if these help identify the model or condition of a device.

Sabotage and Tinkering - GreenTech

Although Christopher did not give specific examples of situations that could be seen as sabotage, or generally express dissatisfaction from employees, we still wonder whether there is potential for sabotage. As mentioned above, we observed employees putting their hands under the hoist, even though this was against the safety rules. This suggests that staff are already bypassing certain safety measures for various reasons, which raises the question of whether they might also break the rules around other robots.

One of the department managers also told us that an employee was unhappy when the hoist was introduced, which likewise suggests that other forms of automation might not be welcome in his view. There are therefore indications that some employees do not want automation/robots to affect their everyday work, which could potentially result in mistrust, dissatisfaction and thus sabotage. Of course, this is not certain, but based on experience and insights from other, similar field sites, we make these assumptions.

The following fieldwork at the company HandsON concerns a company that has successfully integrated robots as a fixed part of its work processes and is therefore an example of a company that has realised these ambitions for automation.

HandsON – A multitude of humans and robots

The fieldwork at HandsON contributes to an interesting understanding of how the workplace culture and the fact that employees are used to robots, causes them to not be afraid, and therefore they further develop the solutions they receive.

An important insight from the fieldwork is that those who work with the AGVs and the other robots can quickly develop a mindset where they trust their own abilities and the robots, they work with so much that they end up sabotaging them or forgetting to show consideration.

In addition, the environment the AGVs operate in is a dynamic one, which is interesting to compare with other places where we carry out fieldwork.

Introduction

The factory is located in an industrial area in a larger European city, where the company's production takes place. We stepped into the building, which has high ceilings and is nicely lit. The receptionist is sitting in a small glass-fronted room, working. We sit down on a sofa and wait for our contact person Patrick to arrive.

When he arrives, we greet each other, and we are allowed to put our bags in the reception area, and then Patrick starts giving us a tour of their entire production. We leave the reception and walk through a warehouse, where we slowly begin to see the many robots, they use in their company. We were still looking for the AGVs we had actually come to see, but first we were shown into a building where they had robot arms for grinding, washing and packing various products.

Here we were told that they had acquired one robot to take over half of another robot's tasks because the job was too complex for one robot alone. Patrick explains that they have around 35–40 robots in total, and that they have had robots since the 1980s. Many of the robots they have are used directly in production.

After this we were shown into a larger warehouse where, besides a huge robot arm behind some fence – whose function we were not told – there was a robot that was supposed to laser-weld metal parts. We positioned ourselves behind the employee who was sitting in front of a screen from which she started the robot. The robot itself carried out the laser welding behind a protective door with a small window where you could look at the robot.

The door slid shut when the employee started the cycle, and then the robot was supposed to begin its work. But suddenly the door shot open, and a red error message appeared on the screen. The employee ignored the message and commanded it to weld again.

We were standing in front of the door to film the robot welding, and we were also surprised and a bit nervous when it opened again, because we thought it might be our fault. The door closed, the laser moved closer, but again it shot open, and Patrick asked: *"Is it because we're standing too close?"* The employee denied this and dismissed the error message again. She started it once more, and this time it was able to do its job. She said this is something it does now and then, and that it has nothing to do with where we are standing.

After this little demonstration we moved on and came to a large gate that opened as we approached, and we entered a bigger warehouse where the AGVs we had come for were moving around between different stations. We went up a set of stairs from where we could look down on all the AGVs driving around, and from there we conducted a longer semi-structured interview with Patrick.

This was also the location from which, on our second visit, we stood and observed for about an hour. Unfortunately, we were not allowed to film or take photos in there out of consideration for the employees.

It is a large warehouse with very high ceilings. The platform we stood on is a small extension that contains some offices and toilets for the staff. We stood above these, from where we could see the entire production. Some employees also come up here now and then to do stretching exercises.

The work of the AGVs

Right below the platform there was a conveyor belt that led into another warehouse where products are stored before being shipped to customers. This is also where the AGVs drive to deliver the products once they have collected them from the workstations. There is also a conveyor from which the AGVs pick up products and deliver them to the workstations. These products are delivered on a roller table by an employee.

There are about 30 workstations arranged in 3 rows of roughly 10 each. The workstations stand freely so employees can move between them. On one side, the conveyor with finished products runs past various inspection points, and on the other side there are tall racks with goods on them, stretching all the way down past the rows of workstations and along the end of the warehouse.

On the far side of the racks, employees drive forklifts, moving various goods and products into the shelves. On the inner side of the shelves, a panel has been installed along the bottom 20 cm so that the AGVs can navigate close to them more easily.

The AGVs move around between the stations on marked "roads". The racks form a physical wall delimiting the area where the AGVs drive. A barrier cord is also strung between the shelves and the conveyor in the AGV area, so employees don't walk in there without reason.

Markings are painted as yellow dashed lines on the floor, indicating where employees must yield to the AGVs. The AGVs move freely among the employees and their stations. On the floor, there are marked routes where they drive, and there is also a small R2D2 figure – the iconic robot from the Star Wars films – printed on the floor to show where the robots travel.

The robots have a load-carrying platform on top where they carry the different packages that need to go on in the production flow. The robots zip around but reduce speed if they see an employee step out into their path, or if they detect a robot coming towards them.

Patrick explains that they had to cover the bottom 20 cm from the floor on all shelves and similar structures in the work area so that the robot doesn't "see" a box and mistake it for a person or a palette, and slow down or fully stop.

Employees' roles in the collaboration

The employees in the company come from many different backgrounds, and there is roughly a 50/50 gender split. However, it seems that those who work in technical support are primarily men. The people who work in the warehouse with the AGVs are not specially trained for that role, but they have all received a short introduction to them, and some have been trained to be responsible for them.

Those responsible are the ones who must help the AGVs if they stop or run into other problems. If they cannot fix them, they have to contact technical support. We were told that the robots navigate using a map of the warehouse, which they created by driving the robot around the warehouse while it scanned its surroundings.

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 - HandsON

In my conversation with Mike, their in-house expert on AGVs, he described how the AGVs can end up in *"unfortunate situations"*. The AGVs have three safety scanners which, in his words, are only used to avoid these "unfortunate situations". This covers people walking out in front of them, or something sticking out from a rack that they really should not drive into.

Mike also thought this mostly happened in the beginning – these "unfortunate situations". He also explained that it was often something they themselves had caused, because they had explored how fast they could drive them, how close they could run them to walls, and so on. They ignored the safety clearance that is programmed into the robots – how close they are allowed to drive to walls and people – so they could, for example, make them drive along a wall where they knew no one would step out.

As mentioned earlier, they have traffic rules in the warehouse. This means employees must yield for the robots and preferably avoid disturbing them too much, as this slows down the robots' working speed and can disrupt the work in the warehouse. That is also why the AGV routes are marked, so that employees are aware of them and do not just step out in front of the robots.

Shark's teeth markings are also painted on the floor to indicate that you should look before stepping out, but they also make it clear that the AGVs will brake before running a person over.

Another interesting point Mike mentioned was a wish that they could run the AGVs faster. However, that would mean they would no longer be as safe for employees to be around. So, this was a trade-off, where more speed would mean less safety.

Beyond that, there is also safety in relation to the AGVs themselves, because they are quite expensive, and HandsON would very much like to avoid them crashing into things and being damaged. So many of the safety measures they have are also there to prevent the AGVs from being broken.

HandsON has robots in almost every department of the company, and this first and foremost means that they are very ready to set up more robots and to handle what the different robots require in terms of space, etc. This also includes any safety requirements the robots might have, and how these can affect safety measures in the workplace. The entire company is built to use and host robots almost everywhere.

They also have plans for the AGVs to drive around in more departments, where all the gates between departments/warehouses will be programmed to open automatically when a person or an AGV approaches them. This will, of course, make the work environment the AGVs move in even more dynamic.

This points to a culture where you must be ready for change and prepared to work with robots almost no matter where in the company you work, and regardless of your educational background.

This can be related to Stian Antonsen's 2009 article on the connection between culture and safety on ships. He describes how the culture on board a ship, and their view of safety, can affect both perceived safety – whether they *feel* safe – and how it can actually affect safety on board in practice.

In relation to a company like HandsON, their culture also matters for both perceived safety and safety. HandsON places a very high value on safety, and so it is something their employees receive training in and something they focus on a lot in the implementation of AGVs as well as their other robots. This can also help employees feel safe when they walk around the warehouses.

For example, they do not feel that their freedom of movement has been taken away by the implementation of AGVs. Some even expressed that they could move more freely now than before, when there were conveyor belts between the stations.

However, this can also lead to them trusting that the AGVs will always brake in time before hitting them – which can lead to them sometimes ignoring the traffic rules of the warehouse.

Trust - HandsON

When you work in a warehouse with robots driving around, especially in a place where there are safety markings on the floor, you might think you need to be careful and that the robots are not entirely to be trusted. But that is not what we were told by the warehouse employees. On the contrary, they all expressed that they trust the AGVs and their ability to stop.

According to Fiona, this can also lead to problems, because people do not "respect" the AGVs. They just walk out in front of them because they assume the robots will brake in time, and they ignore the traffic rules that say the AGVs have right of way. She says it can be annoying to watch when, like her, you know it makes the AGVs work more slowly, which can slow down the workflow in the warehouse.

This is thus an example of employees trusting the AGVs *too much*, and therefore not thinking about the fact that they are affecting the robots' efficiency. She describes this respect as something that can appear quickly when employees are reminded of it – but that also disappears quickly again.

Anomalies - HandsON

On our second visit to HandsON, we were standing with Patrick talking while we observed the AGVs driving around the warehouse. One of them was on its way to deliver some items when it suddenly stopped – about 1–2 meters before the conveyor it was heading towards.

It then stood there for several minutes. None of the employees noticed at first, as it was a fairly quiet day and everyone was focused on their own tasks. But it quickly became a problem, because more and more AGVs – and eventually seemingly *all* the AGVs in the warehouse – lined up neatly behind the one that had stopped, waiting to deliver to the conveyor.

It was then spotted by an employee, who checked a board to see who was responsible that day, then went to find her and make her aware of the problem. It turned out to be a very simple fix. All she did was press a button, pull the AGV about one meter backwards, and then start it again.

After that, it drove over to the conveyor, finished its task, and went off to begin the next one, followed by all the other AGVs that then delivered their items and went back to driving around the warehouse.

In our interview with Fiona, we asked whether she is ever positively surprised by the AGVs and how they work. She answered: *"We're surprised when they work."* She was the only one of the warehouse employees we spoke to who expressed this view. The statement almost makes it sound like it is an *anomaly* when the AGVs run as they should and don't stop.

She also mentioned that they stop almost every day, and that they often have to go out and help them. In contrast, in the conversations we had with employees while Patrick was present, they expressed that they did not feel they experienced many problems with the AGVs.

The vast majority of anomalies they experience with the AGVs seem to be situations where their internal map does not match their actual physical position. We were told this is typically because the wheels do not have enough grip on the floor, or because a pallet or something else makes their internal map no longer match reality. So, this can also be a kind of sabotage – or simply not enough maintenance of the warehouse and the AGVs.

Sim2Real - HandsON

The AGVs in the warehouse have a map of the warehouse, which is generated when they are set up. This works by a technician taking them for a walk around the warehouse while the AGV scans its surroundings and builds the map from that.

From there, it sets up a "clearance", so it knows how close it is allowed to drive to various walls and so on. This clearance can be adjusted, which HandsON has done so that the AGVs can drive closer to walls where they know no one will step out.

The AGVs plan their routes based on this map. They always choose the "cheapest" route:

"There's an algorithm for that, which means it has to… When you calculate it, you calculate it in cost. And it must always choose the cheapest route – the one where it travels the shortest distance, while still keeping its clearance in order, so it doesn't have to slow down."

Mike explains here what is meant by the cheapest route: it is a combination of the fastest and the shortest route. In addition, direction of travel matters, and the AGVs always keep to the right so it is easier to see where they are coming from.

The map also has various zones encoded, where only one AGV is allowed at a time, to avoid congestion and other problems that can arise if too many AGVs meet in the same place. There are also parking spots encoded in the map so that AGVs are available at both ends of the warehouse.

The map is supposed to function as a 1:1 simulation of the warehouse, and the AGVs navigate solely based on their position on the map. This can lead to problems if the AGVs think they are in one place but are actually somewhere else.

For example, this can happen if the floor is too slippery or the wheels are too worn, so that their braking distance becomes longer than normal. Over time, this can cause them to drift out of sync with the map and eventually cause them to stop driving at all.

Therefore, the maps are also synchronised and updated continuously to avoid such errors.

Sabotage and Tinkering - HandsON

When we asked Mike and Patrick whether they experience sabotage, they both denied that this was something that happens. Mike pointed out how expensive the AGVs are, and that this is therefore not something they want people to do.

However, Patrick also showed us during the warehouse tour how they have developed and worked on improving the AGVs since they got them. They also told us that they had changed the "clearance" so the AGVs could drive faster and closer to some of the racks in the warehouse.

This may be because their understanding of sabotage is that it can only be *negative* sabotage – and that is something they neither experience nor tolerate.

Another interesting thing to mention here is a conversation with Fiona, one of the warehouse employees who is responsible for the AGVs. It is her job to make sure that if something goes wrong with them, she goes and fixes it, and she is also the one who must contact the technicians if there is a problem with the AGVs she cannot fix herself.

Fiona explained that there is a hierarchy in the warehouse when it comes to the AGVs. At the top of the hierarchy is "facility": the technicians who repair the AGVs, update them, and this is also where Mike worked. They have not hired anyone new to take over his job with programming the AGVs since he left.

Below them, there is the person responsible for the AGVs on the warehouse floor, for example Fiona. Responsibility rotates weekly within a small group of employees. They are the ones who must be the first to help the AGVs and contact facility if they cannot solve the problem themselves.

At the bottom are all the other employees, who know nothing more than the traffic rules. If they notice that one of the AGVs has stopped, they are supposed to find the responsible person and let them handle it.

And this is where problems arise – and a kind of sabotage – because sometimes they skip finding the responsible person and instead try to help a stopped AGV themselves, which can end up doing more harm than good. Fiona says:

"Yes, there are some who want a bit more responsibility. And it's again about patience. Sometimes, if they just left it standing, it might have sorted itself out."

So, these employees end up *sabotaging* the robot, either because they are impatient and cannot wait for it to sort itself out, or because they want to be "the responsible one" who goes in and fixes the robot.

In a similar way to HandsON, the next field site also has extensive experience with implementing robots (here especially AGVs) in the workplace and in work processes. However, the tendencies identified at MedCare are more concerned with the alternative and individual solutions they develop in collaboration with the robots.

MedCare – Cleaning Robots

This fieldwork contributes with the valuable insight, that automation *must* be a learning-by-doing process, since they do not have other examples to orient themselves toward. For example, if the same error keeps happening with a particular employee, they can ask about that person's routine. But if the errors occur across several employees, they know they have to look for the fault somewhere else in the system. In other words, deviations are used to improve the overall system, and only through deviations can they learn this.

Introduction

Many large companies use robots for cleaning tasks – cleaning floor, windows or some kind of objects. The MedCare Company clean professional equipment for hospitals in an edifice based in the outskirts of a large European city. At our visit we arrive at a huge white building where we are greeted by the secretary. She tells us that this branch of MedCare opened in the beginning of the 2020s, and that she has worked here ever since. We are also told that several hundred employees are attached to MedCare. As we understand it, they are primarily IT staff and cleaning assistants.

We meet our two guides: one has previously worked as a nurse and cleaning assistant but is now a project coordinator (Laura), and the other works in IT (Andy). They take us into a meeting room where Andy starts screen sharing and shows us videos from the parts of the process, we won't be able to see that day. Some areas are off-limits to us due to hygiene rules, and other areas are blocked off for humans (because of robots), so it's not possible to get close. He also shows us the system they use to monitor the process and to see when errors occur.

Both guides - Laura, who has been here since the project started four years ago, and Andy, who joined later to help solve some of the IT problems - explain that in the beginning management had quite naïve ideas about what a robot is and how much it can do. The assumption was that full automation would be a smooth process that would work without human involvement.

They soon discovered, however, that full automation could not meet the required standards for sterility and cleanliness, so they had to hire human cleaning assistants to handle parts of the process.

The Robots

It is a large workplace with many safety procedures and hygiene rules, which points to how many human systems must be built up in automated environments, especially when handling sensitive products (such as instruments for the health sector). All robot arms are caged, so there is no direct human–robot contact while the robot is active. Andy shows that the door into the robot cell cannot be opened at all before the robot has been switched out of operation.

The robot has no understanding of its environment beyond what it is programmed to do and therefore would not notice if a human were standing in front of it. There are other robots on site, such as AGVs, which stop if a person stands in front of them, but they cannot sense humans from every angle, which Andy demonstrates by putting his foot in front of one at a point where it does not stop. So, there are different robots present which are able to interact safely with humans to varying degrees.

The robots are also guided or specifically programmed for the functions they are to perform. In other words, there is still a great need for human control and verification. That is an indispensable part of the daily operation.

Laura explains that the story behind setting up a central cleaning facility goes back to the early 2000s, when it was discovered that the rate of post-operative infections was rising and had actually reached about 10%. This was far too high and ended up costing a lot of money, because more and more patients were getting sick after surgery.

A committee was therefore formed to investigate the state of the different cleaning facilities and whether there was something that could be improved there. They found that the quality across facilities was not standardised and that everyone was doing things differently, based on different local conditions. They also discovered that the work was physically very heavy and that the work environment was poor, as employees were being worn out physically.

They therefore concluded that common standards were needed to secure better quality, and they settled on standardisation and work environment as key goals. On this basis they began building the new central cleaning facilities.

No new measurements of post-operative infection rates have yet been carried out, however, because it has taken much longer to get everything up and running than expected. The project was so delayed in its start-up phase that the guarantees on the robots and systems were almost about to expire before they came into use.

Their work at MedCare basically consists of cleaning and delivering cleaned items to their customers. They emphasise the word "customers", as they want to be regarded as a production company. Treatment of the items is part of a long process that involves many different steps of cleaning and care before they are finally packed.

There is a "recipe" for how they are packed into trays, so that they are placed in exactly the same way every time. After the items have gone through treatment and cleaning, they are packed into special containers designed for this purpose, which fit their facility and the robots. They are then sent to a department where they are packed into trolleys and sent out to the relevant sites.

Over time they have discovered that the items are fragile in ways they had not previously been aware of. Because of the many steps in the process, the items sometimes have to stand for a long time, and they found that if they are not cleaned properly, they rust.

Laura and Andy explain that they have had to constantly learn new things, and because the system keeps changing and new updates arrive, they constantly have to learn new things and then forget the old ones. Their learning with the robots is therefore a dynamic process in which they continually learn and relearn the technology's functions.

They actually say that it is physically impossible to learn everything about the machines, and so they use various pop-up messages on the machine interfaces to continuously teach employees about the technology. This applies not only to training new staff but also to employees who have been there a long time, since it is, as they say, impossible for them to have all the needed knowledge about the machines on top of everything they must know about their daily tasks.

The facility's distributed systems

The robots at the hospital come from a company that makes various robots for hospitals and the healthcare sector. The software that coordinates communication between the robots is MES, which stands for Manufacturing Execution System, and it is a system many companies use for their robots. They get MES from an external company, which is therefore responsible for handling breakdowns or problems in the system. In addition, they have an internal system where orders for the instruments are created. Roughly sketched, the technical system works like this:

A surgical nurse creates an order in the internal system → the internal system passes this order on to MES → MES sends commands to the robots to pick this order. The process can be as fast as five minutes, for example if they receive an urgent order from a customer.

The human role in this technical process is actually quite minimal, but it is still humans who initiate it. As mentioned, an employee has to go in and place an order, which is then sent on to the internal system, which understands that there is now a need for cleaned instruments. An order comes in from the internal system; the crane retrieves the units from the high-bay warehouse and sends them on to a conveyor. The conveyor carries them to the department where the robots are ready to receive the trolleys, which they then unpack and send on to the emptying area. During this process, the boxes get a packing slip which ensures that they are registered and can continue through the workflow.

This entire process above is completely fully automated and controlled by MES and therefore does not require human intervention.

All steps in the process before the picking job is created are referred to as an "object". Andy says: "For each step there is basically an object that triggers something." But once the robots' own software takes over, it is no longer called an object, because the robots do not use that syntax. This is yet another example of how the distributed nature of their system can sometimes do more harm than good, because it can be, as Andy says, like looking for a needle in a haystack when you have to spot where in the process the error occurred.

The robots and MES each run on their own "language", and during the process they have to switch between these two languages several times, which can also bring its own challenges. From Andy and Laura's perspective, this only really becomes a problem when something actually goes wrong in the systems. They themselves who have chosen such a large variation in systems – it is actually a legal requirement, as lawmakers want to avoid a situation where a single developer or company has a monopoly on the systems.

Beyond avoiding monopolies, this also ultimately comes down to money. If a new company comes along and says they can do a part of the system more cheaply, then the system has to be split into modules so different companies have the opportunity to bid on parts of it.

Andy notes, however, that the cheapest solution is not always the cheapest in the long run, as other costs suddenly appear – especially in the start-up phase, where they need a huge amount of technical support.

The system's distributed character therefore ends up being a hindrance for the employees, who would get more out of being able to troubleshoot and solve the problems themselves. During the tour of the facility, everything suddenly stops – there has been a breakdown. Andy casually says: "Looks like an MES problem."

Despite MES's central role in this process, it is not a kind of "super brain". As mentioned, MES simply passes information along – it just delegates, without knowing anything more about what it is delegating. Andy explains that if something goes wrong and the process stops somewhere, MES still continues to send out commands as if nothing has happened, because it is not "smart" enough to know that anything is wrong. This means that after a breakdown they have often been able to see, inside the directory system, that MES has queued up so many commands that the system is completely full, and they first have to clear the queue before they can start up again.

The distributed technology and the many different suppliers make their work more difficult, especially when something goes wrong. Andy explains, however, that one of their colleagues, who is very skilled in technology and software, has actually developed a system specifically for these fault situations.

He has made a screen they can pull up in the system, showing an overview of all order numbers, picking jobs, trolley numbers and other types of information that are relevant for troubleshooting. Andy says:

"And you can basically click into these. And then you get a whole lot of information that, to the untrained eye, looks boring. But for those of us who know what's what in all this, we can see the different processes and what might have gone wrong."

This colleague has therefore managed to create an overview of systems that were otherwise impossible to navigate, which makes their troubleshooting process much more manageable. In this case, it is not the companies that delivered these systems who can take credit for troubleshooting working better, but rather the fact that MedCare happens to have an IT-competent employee.

Even though this is obviously a huge help in the department, one could argue that it might have been better if they *hadn't* solved the problem themselves, so that the integrator companies would be forced to take responsibility for their systems and the problems with them.

The system architecture is extremely complex, and it therefore takes quite a bit of SQL knowledge to understand how the system's information is structured and organized. It is with this knowledge that their colleague Mark has managed to create the overview screen with all the commands and orders across the systems.

A further problem with finding employees with the necessary IT competences is that those people could earn more if they went to work in the private sector. Another problem when errors occur is that they call the relevant supplier, who then starts investigating the error until they hit a boundary that turns out to be another supplier's system. Then the second supplier has to be contacted, who in turn begins troubleshooting until *they* hit a boundary that belongs to the next supplier's system, and so on, until they hopefully find where the error originates.

If the robots stand still for more than a minute, they light up red, which means employees can quickly see when there is a breakdown and fix it as fast as possible. But this red light only applies when there is a fault with the robots themselves. If there is a problem in the internal system or in MES, they technically have no way of knowing until the error propagates all the way out to the robots. With the work Mark has done, however, they can catch some of these errors earlier, so they do not have to wait until they show up at robot level.

Even though errors are cumbersome, they actually end up being a learning process for them, which can potentially have a positive effect. They also pay attention to whether a problem is collective or individual, so they can see, for example, if a specific employee is stressed, unhappy or unfocused for one reason or another.

If the problem is with the technology, they pass it on to the relevant supplier. They tell us that something goes wrong *every single day*, so this process is continuous and never-ending.

Automation potentials

We ask the two guides whether this automation has created fewer jobs for humans than at non-automated cleaning facilities. They say that there are more employees here than at another cleaning facility in the country which is less fully automated — but in terms of size and capacity, this place is also significantly larger.

The employees are divided into departments such as technicians, customer support, reprocessing, cleaning team, instrument team. Although they are split into all these different teams, they are constantly dependent on one another, as a good flow between departments is crucial to making the system work.

Since we visited the facility, a lot has happened, and through their website we can see that training has begun in a new shared system.

MedCare receive visitors from medical professionals from around the globe, who are impressed by their work with automation. This is also because there is enormous potential for further development and digitalisation in the medical field.

For now, it is primarily humans who handle the objects to be hand cleaned and Andy explains that there are no plans to fully automate this part of the work in the near future. It is clear from both Laura and Andy that this is a task that probably cannot be fully automated for many years, as it would require a much more advanced technological development.

They say they do strive toward the goal of being fully automated, but they understand there will obviously be areas where human involvement is necessary. So even though they have a vision of full automation, they are not naïve about how achievable that is.

This is because the cleaning process demands such a high level of detail that full automation would require more scanners, sensors and cameras for the robot to be able to pick up the same things as the employees. In addition, there are aspects of the items that employees notice as a result of experience and close familiarity over time.

This includes, for example, noticing when some parts are rusty, have been replaced, or have expired. All these variations would require programming the robot to be able to distinguish between them and understand what they are. The employees thus have a large amount of knowledge and familiarity with the items that is both difficult to transfer and would require substantial programming and adjustment of the robots.

Furthermore, the automation consists of several overlapping systems — for example for communication, for the robot itself, and for its basic software — which means that many things *can* go wrong, and many things *do* go wrong.

Overlapping IT systems are a result of market competition, but in everyday work they make troubleshooting complex, since it becomes easier to disclaim responsibility if the error could just as well lie with one of the other suppliers. Many systems create more vulnerabilities.

They say that when it works, the automation is fantastic - but when it doesn't work, it can be a complex process to get things running again. Among other things, this has led to a larger IT department.

Both emphasise that there have been learning processes in rolling out this facility: *"We've had to adjust both the people and the technology*," they say. It is, however, a good thing when things go wrong, because that is where much of the learning happens.

Andy explains that there can be quite a lot of waiting time if the robot makes an error: *"We're not allowed to touch anything called programming,"* because of the warranty. In other words, they are only allowed to restart the system. But external IT specialists are hard to get hold of, and it would be much more efficient to have programmers on-site.

Laura says she has worked at the facility for about four years, and that she herself started out doing cleaning work. When the robots were first introduced, she was not very impressed by them, as they were nowhere near as smart as she had expected. She had been left with the impression that they could do much more, and that everything would be running faultlessly from day one (see Report I for how overselling robots raise expectations).

Instead, they required an enormous amount of manual work, as they were very "finely sensitive" and constantly needed a cleaning assistant to come over and check on them. She does think that when they work, they work really well - but when they don't, they require much more work than what they were used to before.

Andy makes an interesting point here: it is difficult to create manual emergency solutions that are as easy as the earlier purely manual solutions. In other words, when you have to step in and help the robot with an error, that solution will rarely be as simple as if you did not have a robot and just did things the way you used to.

The "clean-up work" is therefore the biggest culprit in human–robot collaboration. One of the classic problems they experience with the robots is that employees step in front of the robot's sensors, which makes the robots stop all the time.

Manual emergency procedures and everyday fixes

One of the biggest challenges when robots or systems fail is that they still have to deliver instruments to their customers on time – and in some cases, that can literally be a matter of life or death. That is also why they themselves have developed so many emergency procedures and contingency plans for when the system or machines break down. As they say, they always have an ace up their sleeve – they have to, because patients' lives may depend on it.

Their collection of emergency procedures explains how they can handle things manually when something goes wrong. They have had to develop these themselves, because the emergency plans, they were given only described what to do if, for example, there was a train accident or other external crisis. So, these internal emergency procedures are something they have created in practice, based on what is relevant in *their* specific context.

When the robots and the associated systems were first rolled out, the plan was actually that there would not be a technical department at all, because the idea was that the technology would run without problems and without human intervention. They quickly discovered that this was impossible: problems constantly arose that they could not solve themselves because they did not have any technical staff on-site.

They also discovered that there was an ongoing need for humans to be involved in the processes, which they had previously not counted on. Andy tells us that when all these problems arose in the start-up phase, they called the different suppliers, because no one was sure who was actually responsible for fixing which issues. This resulted in a lot of wasted time and some hefty invoices from suppliers – and in the end, they decided they needed a technical department.

Looking back now, people in the department cannot understand how anyone ever thought this could be done without a technical department. Andy and Laura say that the people who ordered the system had quite naive ideas about robots, and that this is partly because the concept was not developed by technical people. Laura says that many non-technical managers were involved, and that they had not thought of the concept as a production system, because that is not their everyday frame of reference.

However, Laura and Andy think this is the biggest mistake in the project: there should have been IT and technical people involved who could design the concept as a factory. As Andy sums it up: "We are not an operating theatre. We are a production."

Those leaders had imagined the robots as far too agile and simply assumed that they could follow the same workflow as humans. Laura puts it like this: "There are some very specific rules and steps they use out in the departments, and you can't just bring that straight into a production facility."

At the beginning of the process, they actually had consultants who were supposed to help distribute tasks between humans and robots. This division was based on work environment considerations, with the aim of minimizing heavy lifting and physical strain for humans, and to minimize the need for human involvement wherever possible.

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 - MedCare

The robots are actually quite well equipped for safety issues. At one point, Laura shows that if you step in front of the robots, they simply drive around you. The programming on these is not as "locked" as on the other robots that are set to follow a fixed route and therefore can only drive that route. The others are set to drive from A to B, and if something gets in the way, they are able to drive around it.

But they do have some AGVs that are supposed to fetch containers which *cannot* drive around obstacles if someone stands in front of them. If they detect a person in front of them, they light up red and then turn green again when the obstacle disappears, If the person or obstacle remains there long enough, the robot says, *"please move out of the way".*

Andy explains that even though these AGVs cannot drive around obstacles, they are still smarter than the old AGVs they used to have, because they are capable of moving again once the obstacle is gone. The AGV does, however, have a challenge when it comes to angles: if you approach it from the side, you have to get quite close before it detects you and stops.

The problem is that if you are caught in its blind spot and up against a wall or something similar, you could actually be crushed because the robot does not detect you and therefore does not stop. They have experienced, once at another facility, that an employee was run halfway over because of this blind spot. The biggest safety challenge here is therefore the quality (and perhaps also the number) of sensors and cameras.

They have tried to anticipate some of these failures, and they themselves are aware of the challenge that you can never secure *everything* when you are dealing with humans. This is also one of the reasons why employees are not allowed to enter certain robot areas. In addition, it is another reason why they are not allowed to "mess with" the robots' programming, as the company that built them must be able to take responsibility for any accidents or injuries.

The robots are programmed to withstand a certain amount of resistance, which is true of all robots and is something that is often assessed in a risk assessment. This is not something they explicitly talk about at MedCare, but we assume that this resistance setting is part of the robot's programming that they are not allowed to change, since it is based on specific safety decisions.

In practice, however, this means that the robots can end up damaging the items, as these are very fragile and the robots do not have the same kind of sensitivity that humans have. Moreover, the robots are built on the assumption that everything is done correctly and that the procedure always runs as it should, and therefore they are not very sensitive to deviations or errors.

One could imagine that this has been considered in a possible risk assessment, but we have not heard anything about that from Andy and Laura during our visit. It may be because they themselves have not been involved in that work, as it may have been done in collaboration between the robot developers and MedCare's senior management.

However, it is clear that the challenges Andy and Laura continuously point to could likely help nuance the risk assessments of the robots.

One of the tasks the robot has is to go and fetch containers. It has to lift up a net, pull out the container, and then lower the net again. This is one of the tasks where work environment considerations have clearly been central, since the containers are very heavy and employees could potentially get their fingers caught.

In this case, the robots help reduce work-related injuries, and that is why this particular task has been designed without any human involvement at all. Employee safety has thus repeatedly been used as a basis for deciding the robots' tasks, rather than simply thinking in terms of replacing workers.

There are, however, still a number of tasks at MedCare that could advantageously be automated, as they are physically demanding for employees. The staff who send off the boxes also apply labels with the packing date. This task is very manual and quite wearing, as they often have to bend down or stand in awkward positions to reach the boxes.

Trust - MedCare

Questions about employees' trust in the robots show up most clearly in their use of sabotage or "everyday fixes" around the robots. This will therefore be unfolded more in the sections on Sabotage and Anomalies.

In general, it seems that employees at MedCare have more mistrust towards the *suppliers* of the systems than towards the *systems* themselves. They put a lot of emphasis on how difficult it is to get hold of the right person when they experience problems or errors. Their description of being passed from support person to support person, with no one taking ownership, creates a sense of responsibility being pushed away, which in turn creates mistrust among the staff.

Andy and Laura do not express direct mistrust towards the robots, despite mentioning examples of employees who have sabotaged them. They mention the example of the AGV's blind spot and the colleague who was run over, but this is told more as an observation of the robot's limitations than as an expression of mistrust towards the robot.

They repeatedly say that the robot stops if you step in front of it and they demonstrate this during the tour. Andy says that he thinks the high-bay warehouse is the "most fun" place, because that is where the most dangerous robots are and where it is the most off-limits for employees. This is because the largest and heaviest robots (2–3 tonnes), which are not intended to work together with humans (industrial robots), operate down there.

Despite the robots' limitations and "dangerousness", there is more fascination and interest in them than fear and mistrust. Laura and Andy say that errors occur in the department every day, and again this is narrated as disappointment towards the *system* and software, rather than towards the robots as hardware.

Anomalies - MedCare

These interactions with the robots are not necessarily only negative. Another thing we observed in the fieldwork is how employees invent small, creative everyday fixes to help the technology along.

The robots are triggered by so-called “objects” — small clicks or vibrations that tell the technology that it should move on to the next step, for example start a movement or move something.

When we were at the goods reception, we were shown how racks with many boxes of instruments are pushed into a lane. Three minutes after the last trolley has been placed in position, the lane is supposed to start and transport them down to the end, to the elevators where they continue onward.

But the lane does not always register that the trolleys are ready, the transport worker explains. “It doesn’t always start, so we give it a little kick,” he says, showing how he gently kicks the area near the sensor, so it detects the vibration and starts up. “We don’t know why it happens,” he says, “but if it works, we do it.”

This kick is, in principle, an anomaly from the robot’s perspective, since the system is probably not designed with the idea that someone should kick it. But it ends up being a helping hand for the robot.

It is on the basis of this example that we have developed the concept of “positive anomalies”, since cases like this show that a deviation from the robot’s perspective does not necessarily have to be negative. These positive anomalies mostly arise in practice, as employees learn together with the robots what works and what doesn’t.

These everyday fixes move from being spontaneous, random nudges or kicks to becoming necessary, embedded practices that arise from the collaboration between robot and worker.

Another example of employees’ everyday solutions is when Laura describes how the containers have barcodes on them, which are used to scan and register the boxes. These barcodes must, however, be washed off when the box is finished being used, as it then needs a new one for the next order.

Although the barcodes are made of paper and mostly come off in the washing process, there is always some glue left, which is very hard to remove. During the tour, Laura instinctively walks over and starts peeling off these barcodes from nearby boxes. She explains that this has become a kind of “occupational habit” - they just go around removing barcodes when they see them, because then they know the work will be easier later on.

This kind of human intervention has no formal procedure and does not directly affect the robots, but it still provides significant help to them indirectly by making later work less complicated. This intervention is based on experience, habits and intuition and is not something they can easily teach the robots to do.

It shows that although there are a number of tasks at MedCare that can undoubtedly be automated to improve employees' work, there are still many small tasks that they perform without really noticing. It is precisely these tasks that are so difficult to teach robots and therefore so difficult to automate.

Even though this is not an anomaly *for the robot*, it can potentially become a bottleneck for production if this glue is not removed. So, in the same way as positive anomalies, it is an example of human intervention that has a positive effect on the robot's work.

There are a lot of different systems, and it is therefore necessary that everything runs correctly for the whole system to function. They have considered hiring interns to help handle trolleys that cause errors or go wrong.

The containers mentioned earlier run along a conveyor, and to make this procedure completely free of anomalies, it is designed so that the container runs up against a perfectly straight edge so it can never enter crooked. Because there are so many systems and thus many possible bottlenecks, they simply do not have "room" for too many anomalies — and this is precisely why employees' everyday fixes are so critical for keeping operations going.

Sim2Real - MedCare

The employees do not say anything about simulation at MedCare. The visit to the facility also took place before this analytical concept was selected, and therefore we did not actively ask about it during fieldwork.

One can, however, assume that the robot developers use simulations in some form, as this is something we have seen at all the robot manufacturers we have encountered. Since the employees do not bring it up themselves during the visit, we can assume that it is not something they experience as a major source of problems in their everyday work.

Sabotage and Tinkering - MedCare

The aforementioned AGVs that were less intelligent than the ones they otherwise use were not dropped only because of this lack of intelligence. Andy tells an anecdote about how, when these AGVs first arrived at MedCare, the employees were quite annoyed with them because they feared being replaced.

These employees quickly discovered that the robot could not move on if it detected something in front of it, and so they began sabotaging the AGVs by placing objects in front of them, stopping them from doing their work. The robots therefore turned out not to be as effective and productive as hoped, and the project was scrapped.

Andy does, however, believe that sabotage or not, the AGV project would probably have been cancelled eventually anyway, as they were simply not intelligent enough. The sabotage in this case stems from a fear of being replaced, and Andy explains that this fear is understandable in some ways, since it is often those employees who have fewer tasks and are more dependent on their current tasks who feel the most threatened.

Another example of an employee disrupting the robot's work is where a worker "tricks the robot". Andy tells the story like this:

> *"Somehow we had a colleague who tricked it into thinking that it had already... or that it hadn't placed a container where it had placed a container. Then it tried to push another container into the same place. And as you see, it doesn't stop for anything. So, it pretty much knocked the trolley right off."*

It is unclear whether this employee did it on purpose, or whether it was a mistake that they ended up "tricking" it. In any case, it again illustrates the complex relationship between robots and employees, as the work at the facility depends on them being able to cooperate — yet the examples above show the opposite.

Whether this sabotage stems from mistrust, fear, or simple dissatisfaction is not entirely clear, but ultimately it shows that humans are constantly interacting and negotiating with the robots.

The example also shows that these interactions can lead to accidents. Andy explains that the trolley almost got completely knocked off, and whether or not that was the employee's intention, the situation creates a risk that someone or something could be harmed.

To understand how these fields of research together shape our understanding of robots in organizational contexts, we will now compare them.

Understanding Robots in Controlled Environments

The field sites in this report are, as described, characterised by robots that operate in daily collaboration with people who are trained to work with them, and in controlled environments where they do not encounter new humans.

Here Human–Robot Interaction (HRI) becomes a central focal point for how technology and humans interact, as Tom Sorrell addresses in his article on cobots from 2022 (Sorrell 2022). A recurring theme for Sorrell is that cobots are often marketed as collaborative helpers, but in practice gradually move in the direction of replacing human labour. This can lead to an ethical dilemma where employees both lose tasks and are used to train their technological replacements. This becomes especially problematic when the "collaboration" between human and cobot does not involve mutual benefit, but an asymmetrical relation where the human ends up servicing the robot, for example by constantly restarting it, cleaning up after it or compensating for its shortcomings.

This insight can help to understand the caution and carefulness that ElectroFlight saw from their customers and their employees. This could also be a potential problem for GreenTech, if they implement a robot, and the employees have to help "fill the empty AI brain" with their extensive knowledge of the products.

Oliver Bendel elaborates in an article from 2022 that precisely this aspect of the cobot collaboration is often overlooked in ethical regulation, and that the use of cobots can lead to a de-skilling of employees, where human professional judgement and situated experience are marginalised in favour of standardised processes. He also points out that the human ability to make decisions in complex, context-dependent situations is difficult to translate into algorithms. This can be linked to the two brothers Hubert and Stuart Dreyfus' theory of human skills, which, according to them, develop through stages from novice to expert (Dreyfus 2004). They define a novice as someone who follows fixed rules, and an expert as someone who acts intuitively on the basis of embodied and experience-based understanding. Cobots can be said to typically be at a "novice-like" level, as they can only act on the basis of explicitly programmed rules and lack the contextual understanding and intuitive judgement that characterise human experts. This creates limitations in the HRI collaboration, because even though cobots can perform a range of predictable tasks, collaboration in practice often requires improvisation, flexibility and a sense of the situation.

This adaptability and flexibility is especially prominent in the case of MedCare, where employees develop their own solutions, in order to adapt to the robots' fixed programming. In the article *"Workarounds and Shadow IT – Balancing Innovation and Risk"*, Mark S. White (2023) investigates how employees in complex organizations develop informal solutions to bypass systems and processes they experience as impractical or too slow. White describes these actions as "workarounds" and notes that such practices often emerge in technologically complex environments where the organization's official IT structures do not adequately support employees' everyday needs.

As a result, employees resort to alternative programs, unofficial data channels or manual procedures in order to solve problems more efficiently. According to White, these "workarounds" are paradoxical: on the one hand they create significant risks, since they take place outside formal IT frameworks and can lead to data loss, lack of traceability and weakened information security. On the other hand, they can also be seen as sources of innovation, because they reveal where the official systems do not fit users' actual work practices. As White points out, organizations can benefit from viewing "workarounds" as signals of necessary improvements rather than simply as rule breaking.

"Workarounds", and thus everyday fixes, become expressions of employees' adaptability and ingenuity and can provide valuable insight into how organizational processes can be improved. Laura and Andy explained how they are constantly learning and unlearning in relation to the technology, which highlights a high degree of adaptation and creativity among employees. They very rarely cancel an order; they usually manage to get things out on time, but sometimes they are forced to step in and fix the problem manually. Many small adjustments are constantly being made; it is all about "learning and unlearning" and constantly updating procedures for the cleaning staff.

This issue is supported by Dreyfus' point that one can only move from competence to expertise through emotional engagement and responsibility for one's own decisions. HRI can therefore never be completely symmetric, because only the human element in the collaboration has the capacity to learn, improvise and take moral responsibility. In summary, the texts show that a well-functioning HRI cannot be reduced to functional efficiency or superficial collaboration. To be meaningful, the interaction must respect and preserve human judgement, experience and professional expertise. Without this, cobots risk not only replacing human labour, but also eroding the expertise and flexibility that precisely make humans indispensable in complex work contexts.

Compared to industrial robots, where the division of roles is clear and physically separated, the above text shows that a kind of grey zone arises in the cobot collaboration where the robot gradually takes over those parts of the work that previously required human decision-making ability, but without being able to match the professional level. At the same time, the employee's role changes from active performer to passive monitor, which is a development that can undermine both motivation and job satisfaction. This change of role is clear in HandsON's development of a "hierarchy", where certain employees are trained in aiding the robot if a breakdown occurs. Where this gives these employees the opportunity to expand their technical knowledge of the robot, it also creates a difference in employees where some possess skills that other don't. As explained in the fieldwork, some employees actually wish to have this technical knowledge, and therefore attempt to aid the robot, even though they are not the responsible one. This creates a trouble dynamic between employees, but it also results in unintended sabotage and therefore delays the robot's work process.

At ElectroFlight, Nick showed us a video of a robot that has to move some plates into a box with grooves, and we notice that the robot makes a small notch to one side before placing the plates in. Nick explains that this is because the plates bend slightly, and therefore they are hard to get into the grooves if you do not start with a notch in one corner. He says that they learned the importance of this notch by observing the people who did it first and then asking them about it. We therefore ask whether they always talk to the employees who carry out the tasks that are to be automated with the robots, to which he replies that they have to, because the bosses do not know what it is like to work with the tasks every day. It is noteworthy here that those who are actually to be "replaced" possess embodied competencies and expertise that are necessary for them to develop their robotic "replacement".

Once again, this is an example of how employees' workarounds, knowledge and role as experts are absolutely necessary in robot development in order not to create useless and unnecessary industrial robots. All of these examples above thus make clear the need for human involvement in robot systems, which supports Nick's point that people should not fear the implementation of robots and the automation of industrial tasks. It is therefore very important for future research and development in the field to be aware of not seeking to create systems that are to directly replace humans.

Another challenge with creating these autonomous and elaborate robotic systems, is that if often entails a degree of distribution of the parts (whether it is hardware of software). As explained with ElectroFlight and MakeBot, this also complicates the understanding of responsibility, since it becomes harder for the developers and manufacturers to control the end users.

In the article *"Practical Accountability for Distributed Systems"* by Haeberlen, Kouznetsov and Druschel (2007), the concept of accountability is presented as a central property of reliable and complex IT systems. They describe how distributed systems often consist of several independent domains that must work together, while no single actor has full insight into or control over the entire system. This creates major challenges for monitoring, fault identification and the allocation of responsibility, especially when errors occur across organizational or technical boundaries.

This is exactly the problem we see at MedCare, where the many different systems for warehouse management and robot communication make it difficult to maintain an overview of possible errors. The various systems are developed by separate vendors and managed by different departments, which means that responsibility for operations and error correction is fragmented.

When a robot stops working, it is often unclear whether the fault lies in the robot software, the communication link or an external control module. This uncertainty leads to lengthy troubleshooting sessions where employees have to spend time tracking down the correct contact person or supplier. As Haeberlen et al. describe, such situations resemble faults in distributed network systems, where lack of overview and coordination means that problems are diagnosed "ad hoc", which often extends downtime and increases operating costs. We see exactly this at MedCare, where the employee Mark has used his IT skills to create a more manageable troubleshooting platform for himself and his colleagues. As explained, these "ad hoc" or everyday solutions are central to the employees' daily work, whether they appear as positive anomalies or as in the example of Mark's troubleshooting screen

We see this, for example, in the employee who gives the trolleys a little kick to get them moving – even though he does not know *why* it works, he does it anyway. In our view, this example lies right on the border between positive and negative anomalies, and equally on the border between risk and innovation.

For MedCare, these informal solutions can be seen both as a risk and as a sign of organizational innovation. The risk lies in the fact that unauthorized actions can cause faults in other systems or compromise data security, which in the end can be very costly. At the same time, employees' everyday fixes are an important form of in-house innovation, as they show where systems do not work optimally and where the company should focus its development efforts.

In general, we see at the different field sites, that employees adapt themselves to the robots, in order to avoid errors given the complicated troubleshooting process. What we *do not* see is robots adapting to the employees, the environment or the work routines, which seems to be a general problem.

This is evident, for example, in the story of a facility, familiar to MedCare, where the physical environment literally could not withstand the robots. They have not had that particular problem at MedCare, as they tell us that the place was specifically built with future automation in mind. Still, there have been modifications along the way. Laura explains that the AGVs depend on being guided through the environment, so they had to put up markers and move equipment around so the AGVs could get through. Same is evident at HandsON, where they have made safety markings, and adjusted the heigh of shelves, in order to accommodate the robots.

Once again, these are examples of the environment having to adapt to the robot, because the robot cannot adapt to the environment. Lack of robot autonomy and adjustability therefore creates a lot of extra cost in the industry. These robots are programmed very precisely for the functions they perform, and they cannot themselves resolve faults or anomalies, but must simply stop and start blinking red.

In the article *"Laws on Robots, Laws by Robots, Laws in Robots: Regulating Robot Behaviour by Design"*, Leenes and Lucivero (2014) describe the concept of "techno regulation", which concerns how the design of technology can function as a form of regulation of human behaviour. According to them, norms and rules are embedded in technological artefacts so that, in practice, they guide or limit how people can act.

For example, a robot may be designed to only carry out certain procedures or respond to specific commands, which means the user must adapt to these built-in rules instead of deciding freely how the work should be done. In this way, technology becomes a kind of regulator that shapes the behaviour of those who interact with it.

This form of techno regulation has important implications for the fields sites in this report. The robots used in production are developed based on technical and safety standards which, in practice, dictate how the work must be performed (as explained by Rosa at MakeBot). Because the many different systems are not flexible or adapted to local work processes, it is the employees who must change their routines to make the technology work in practice.

According to Leenes and Lucivero, this entails a loss of autonomy, in the sense that decision-making power and freedom of action gradually move from the employees to the machine's built-in logic. This technological steering is closely connected to the overarching problem described at the beginning of this section regarding workarounds, ethical problems and distributed technologies.

When the robots and systems restrict employees' options for action, and when error-reporting processes are at the same time slow and fragmented, this creates a need for informal solutions. The everyday fixes that employees develop can be understood as a direct reaction to techno regulation, as they represent attempts to regain control and autonomy in a work environment where technology otherwise dictates the boundaries for action.

Taken together, the report shows that the implementation of cobots is not just a technological project, but a social, organisational, and practical interplay between expectations, experiences, and concrete work practices. From the companies that mainly see possibilities, to those that make a living by selling the technology, to those that struggle to make it work in everyday life, and finally to those that have made it a natural part of their working life – together they sketch a picture of robotic technology as something dynamic and changeable. Robots do not only create efficiency; they also create new relationships, new routines, and new understandings of work, which are constantly evolving as people use, challenge, and adapt the technology. Companies are especially faced with the challenge of creating work processes that facilitate collaboration between robots and humans. MakeBot and ElectroFlight explain why the process of ensuring physical safety has to be in place before robots become part of an actual work environment. The other field sites are examples of what it actually looks like when a collaboration between humans and robots has to be established. Here it becomes clear that employees create their own practices to make the robots more relevant to their work, because the robots cannot adapt to their work – it has to be the other way around.

For even though the physical parts of the workplace are static, the actual work dynamics are constantly up for negotiation and in flux. In a warehouse where the shelves do not move, it is not necessarily the technical or programming aspects of the robots that are the most difficult. Instead, it may be a new colleague who does not understand why the robot stops at a corner, or an experienced worker who was there before the robots and is therefore dissatisfied with them. The robot cannot be programmed to handle such situations; instead, people have to adapt to the robot's fixed programming. And unlike robots, people can be driven by feelings rather than reason. If an employee feels that the robot is working against them or is there to replace them, then instead of thinking about what is best for the workplace or the robot, they do what they feel like doing. It is in these situations that sabotage and other forms of negative interaction with the robots arise. On the other hand, we also see employees creating alternative work processes to accommodate the robots' fixed programming. Everyday fixes, internal hierarchies, positive anomalies and home-made IT tools are examples of employees who see the robots' shortcomings and, instead of resisting, create innovative solutions. It is also an example of how the unrealistic expectations of robots presented in Sub-Report I show up in practice, and how they are handled by the employees who interact with robots every day.

The report makes it clear that, even though the robots operate in a physically static environment, the work processes themselves are dynamic and can change depending on the social context. Even when employees in many places manage to relate to the new work processes and routines, the smallest changes can turn those processes upside down and people must adjust. Where they understand how to adjust, the robots instead stop, throw error messages, or do something wrong. The question then is: if the robot were autonomous and could adjust to the humans, would the two end up getting so much in each other's way that it would all go wrong? This question of autonomy will be explored further in the final report, Sub-report 5. The next report, however, examines what it looks like when robots move in a far more dynamic environment, where the surroundings around them are in constant change.

References

Antonsen, Stian. "The Relationship between Culture and Safety on Offshore Supply Vessels." Safety Science 47, no. 8 (2009): 1118–28. doi:10.1016/j.ssci.2008.12.006.

Bendel, Oliver (2022). Cobots from an Ethical Perspective

Dreyfus, S. E. (2004). The Five-Stage Model of Adult Skill Acquisition. Bulletin of Science, Technology & Society, 24(3)

Haeberlen, A., Kouznetsov, P., & Druschel, P. (2007). Practical accountability for distributed systems. Max Planck Institute for Software Systems.

Leenes, R., & Lucivero, F. (2014). Laws on robots, laws by robots, laws in robots: Regulating robot behaviour by design. Law, Innovation and Technology, 6(2)

Mol, Annemarie (2010) Care in Practice: On Tinkering in Clinics, Homes and Farms

Sherson, J., Rabecq, B. T. M., Dellermann, D., & Rafner, J. F. (2023). A Multi-Dimensional Development and Deployment Framework for Hybrid Intelligence. I P. Lukowicz, S. Mayer, J. Koch, J. Shawe-Taylor, & I. Tiddi (red.), *HHAI 2023:* Augmenting Human Intellect

Sorell, Tom (2022). Cobots, 'co-operation' and the replacement of human skill

White, M. S. (2023). Workarounds and shadow IT – Balancing innovation and risk. Business Information Review, 40(3).

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.