Artificial intelligence in camera systems for areas at risk of explosion

Figure 1:
Artificial intelligence and visual analysis systems increase occupational health and safety in areas at risk of explosion.

A system supervisor detects a fire in the Ex area via camera surveillance. They identify two injured people, start evacuation and prevent others entering the danger area. Do you think that's a good reaction? That's right. Only that the system supervisor is not a person, but rather a deep-learning algorithm!

Our impression of artificial intelligence is shaped by science fiction classics such as Stanley Kubrick's "2001: A Space Odyssey" or James Cameron's "Terminator". In both films, the artificial intelligence created by humans spells doom for their creators. The HAL-9000 computer in Kubrick's masterpiece trusts its own capabilities more than those of its creators. It begins to sabotage them and finally becomes their murderer. "Skynet" from "Terminator" declares war on its human creators. But what's the story with these apocalyptic visions of "self-determined machines"? How far are we in the development of artificial intelligence? And what does all of this have to do with camera systems, modern industrial systems and areas at risk of explosion?

Artificial intelligence

In order to introduce artificial intelligence, it first makes sense to look at natural intelligence. You can define this as the "capability to think abstractly and to devise actions from this". Humans are intelligent according to this definition: they think and react. This is not the case for a normal computer program. It neither thinks nor learns and it does not draw experience from its actions. It does not matter how complex the classic algorithms are: they can only implement what their programmers specified.

Let's play "GO"...

...is what the developers at the Google company DeepMind thought and developed AlphaGo. "GO" is a game that is more complicated than chess and AlphaGo was programmed so that it could learn "GO". The algorithm improved with every game and won the "GO" world championships in China in 2017. The algorithm does not just have a neural network, it also has short-term memory. This corresponds to a "memory" as far as possible. Even if that event in China in 2017 did not signal the

birth of artificial intelligence, it was still a significant breakthrough. What does this all mean with regard to cameras?

Seeing "remotely"

This desire is probably not much younger than humanity itself: to be able to see something outside your range of vision in order to then be able to react to it adequately. This desire was first overcome in 1942. Siemens developed the first surveillance camera on behalf of the German army in order to be able to monitor the rocket test beds in Peenemünde, on

which the V2 rockets were developed, constantly. Cameras could be used to see "remotely", analyse (and) act.

Electronic networks

In 1996, the Swedish company AXIS introduced the first network camera to the market. In comparison to analogue systems, digital cameras have the significant advantage that the on-screen information is available electronically and can also be analysed electronically. Cameras haven't just been used simply for imaging surveillance for a long time. Thanks to ever-improving processors, the decentralised task of image evaluation can be performed directly on the camera, known as "On-Edge" solutions. Benefit: lower bandwidth use, as the actual raw data analysis is performed on the camera and only metadata has to be transferred. Events can therefore be reacted to almost in real time. As computer performance improves, the analysis functions have become more and more powerful. Algorithms can read number plates or detect movements for example. You could say that, in the Seeing "remotely", analysing, acting" process chain, systems are taking over the analysing part more and more skilfully.

Machine vision...

...has become a sector in its own right thanks to the advances in classic analysis algorithms. Special cameras with internal or external CPUs can distinguish between "good screws" and "bad screws" with a frequency of which neither the human eye nor the human brain would be capable. Faces are recognised and compared, and biometric data is becoming ever more important. No matter how complex classic analysis algorithms are: they only process what their programmers specify.

Deep learning...

...is a decisive step towards automating camera systems and machine learning. The aim is clear:

the system supervisor should not see that there is a fire in the system; the camera itself should detect this.

It is not the operator who should see that employees are in danger. The camera system itself should detect this. And the camera system should obtain experience and learn from this.

From a technological perspective, there have been significance advances in recent years but yet we are still just at the beginning. Intelligent cameras such as our ExCam IPQ1615 are equipped with a double CPU. While the main CPU is occupied with video streaming, the Deep Learning Processing Unit (DLPU) is only intended to execute intelligent analysis algorithms. Our process chain has become "Seeing, analysing, acting, learning, seeing, analysing, etc...". It no longer ends with the action, but is closer to the processes that run in our brains with "learning and obtaining experience". The British company KCA DEUTAG is providing the first look into the future.

Intelligent safety on drilling systems

Are workers moving in the danger zone? If yes, how dangerous is this for the worker? How can the danger be counteracted? Standard analysis algorithms reach their limits here.

"The current algorithms that we use are able to differentiate between moving people from moving objects", says Mr Schevel, project engineer at KCA DEUTAG. Furthermore, the ExCam IPQ1615 can even differentiate between cars, lorries and motorbikes thanks to new technology for object analysis. The enormous improvement to the surveillance functions and the precision of the data collected result in fewer accidents on the system. "Of course we would like to find out even more from our camera systems in the future, such as whether a person is standing or lying down and whether the person has appropriate protective equipment. We decided on the SAMCON ExCam IPQ1615 specifically because it has its own DLPU that provides an adequate hardware environment for future algorithms."

Learning from mistakes

Let's imagine the scene that is shown in Figure 2. What does a human controller's "manual brain" think?

Five people are securing an unknown load with a rope. Two people are active, three are passive. All five have hard hats and life jackets on, and the person holding the rope is "loaded". One person has their hand on the first person's life jacket and

FIGURE 2:
Example of automatic danger detection with an AI approach.

another, female, person is securing the rope in the second row. She is close to the water's edge that is safeguarded with railings.

In the second step, the experienced operator interprets the scene: The person holding the rope can lose their balance so that the person at the water's edge can fall over the railing if there is a sudden jerk.

The following occurs neurally: the operator compares the situation with their experience. Without "knowing" that a sudden jerk may occur and without "knowing" that it is life-threatening to fall from the drilling platform into the open sea, they cannot evaluate the scene correctly. And if they cannot analyse the scene correctly, they cannot act accordingly.

What can artificial intelligence do?

As of today, using standard algorithms to interpret the scene described above is difficult to impossible. Most analysis algorithms interpret images, image changes (movement) and/or contrasts. They neither detect how the scene is established, i.e. its history, nor do they have the experience to evaluate what "could" happen.

AI algorithms that will take on these tasks in the future are first "fed with basic knowledge": E.g: Working without a hard hat = danger. Person holding the rope = danger. AI algorithms then learn independently, also with the help of experienced operators. Just as parents do with their children, they will pass on their experience to the AI algorithms. If an experienced operator evaluates similar situations as dangerous situations four times, they teach the AI algorithm to judge the situation to be dangerous itself when it occurs for the 5th time. Two points are important here:

An AI algorithm does not just identify the scene as an image, i.e. as a snapshot. Thanks to its short-term memory, it can

also interpret a scene with its historical context. For example, if a worker was wearing their hard hat 2 seconds before and has not taken it off, it is probably that they are still wearing it even if they are obscured by another person. An algorithm without a "short-term memory" would not be able to do this. The algorithm becomes better with each analysis. If people have fallen into the water twice from near the water's edge, its "sensitivity" towards people standing close to the water's edge increases.

Acting actively vs. suppressing actions

AI algorithms acting actively still remains hard to imagine. What happens if the AI algorithm makes the wrong decision? Who takes responsibility for that? This is the exact question being asked in the quest for autonomous driving. Must we or can we impart moral and ethical principles to the AI algorithms? What should an autopilot decide if it can only prevent running over a group of school children by driving on the wrong side of the road and therefore endangering the life of the vehicle occupants? As with autonomous driving, there will be intermediate steps with AI algorithms before systems act autonomously. The two most important steps are WARNING and BARRING, and these are often found in the form of driver assistance systems in modern vehicles today.

Warning

An algorithm can warn the system operator. This works like a child who wants to learn asking its father:
"Look: that boy is walking on the frozen lake. Isn't that dangerous?".
The child is not acting independently yet in this situation. It is using its father's confirmation to obtain experience. If dad calls the fire brigade and the boy can be saved after falling through the ice, the child learns: that was dangerous!

It learns without having acted itself.

Barring

Barring means preventing actions that could lead to (further) danger. In modern vehicles, the accelerator is locked when you brake, which seems sensible. With regard to our scene in Figure 2, the system could prevent ship traffic under the platform for example. The tragic accident at BASF in autumn 2016 can be used as a concrete example for potential loss prevention. An initial fire followed by a series of explosions occurred, probably due to an angle grinder cutting into a live product wire. A tanker full of ethylhexanol was at the quayside to discharge its load at the time of the explosion. Luckily, the fire brigade was able to stop the fires spreading to the ship!

Future AI algorithms would be able to recognise that working with an orbital grinder on a live production wire is dangerous. They would therefore be able to act quicker to prevent this than a trained system operator and they would guide ships out of the danger zone, as well as issue guidance for how to avoid further dangers.

Conclusion

AI will become more and more important in increasing people's safety in systems at risk of explosion.

We are at the start of a technological development where the final consequences are still hard to envisage.

Experience feedback makes algorithms able to learn and they can apply what they have learned to complex situations. Whether their actions will be better than those of humans in the distant future is an ethical question.

While AI algorithms will be able to not get distracted by a tea break and therefore make fewer mistakes, will they ever be able to develop an instinct for how to act?

Own work by Altera levatur, CC BY-SA 4.0, https://commons.wikimedia.org/w/index.php?curid=62951837