An overview of

Artificial Intelligence (AI)

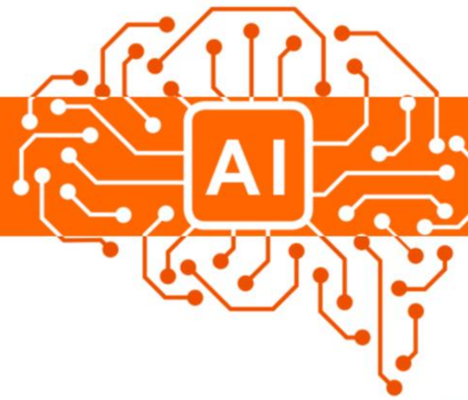

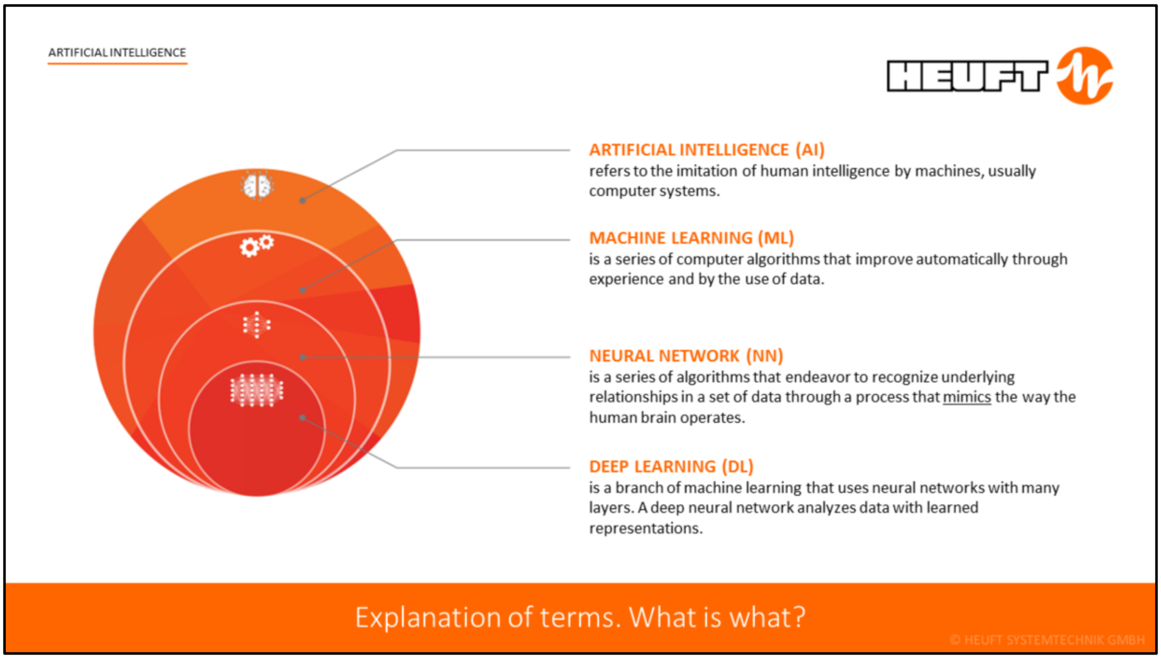

First of all it is important to understand what the terms stand for as they are constantly mixed up

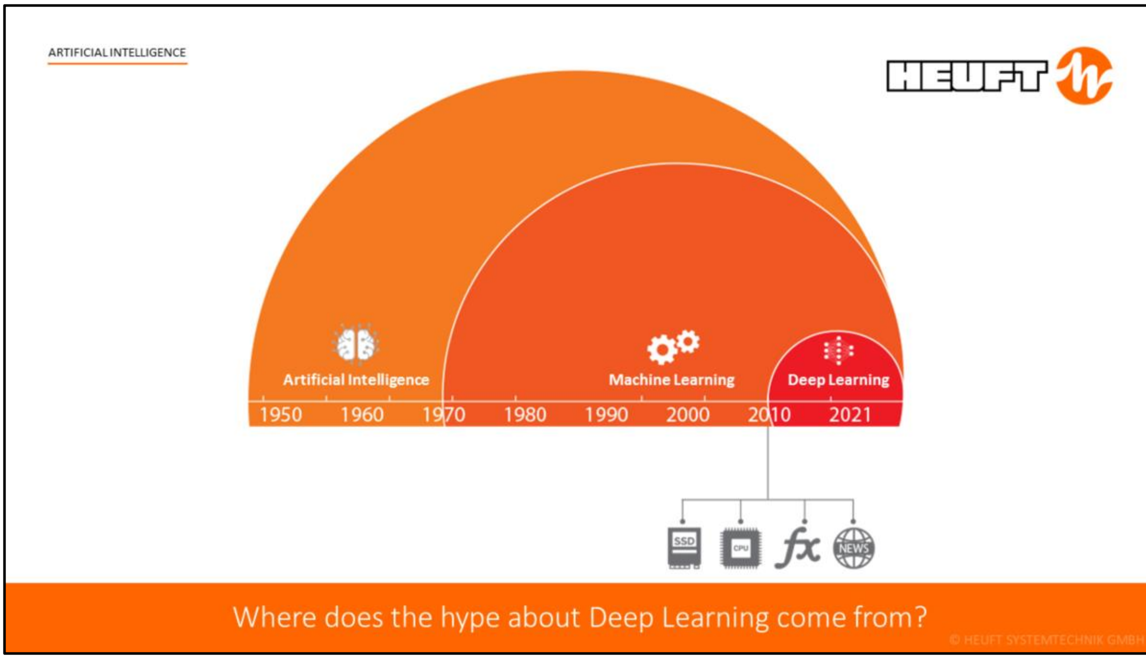

The Idea of artificial intelligence is not new. In 1943 Warren McCulloch and Walter Pitts created a computational model for neural networks.
Deep learning started to become popular in the early/mid 2010s due to a rise in computational power and data storage availability, advances that have been made in the algorithms and marketing and news media awareness.

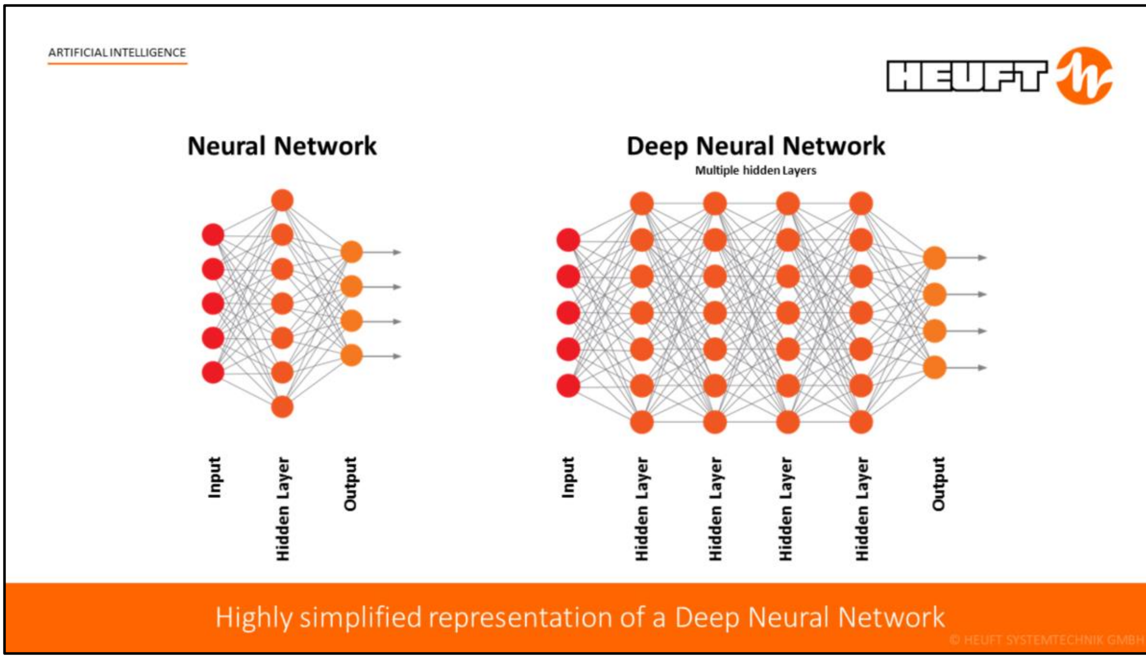

Here you can see highly simplified representations of neural networks. The basic idea is to mimic the operation of the human brain bei interconnecting artificial neurons (the dots). By weighting the connections, outputs are then generated from the inputs as they traverse the network.
What makes a Neural Neetwork "Deep"? Simply said, the number of hidden layers makes it a Deep Neural Network. If there are many of them, more complex tasks can be processed. On the other hand, the demand for computing power also increases rapidly.

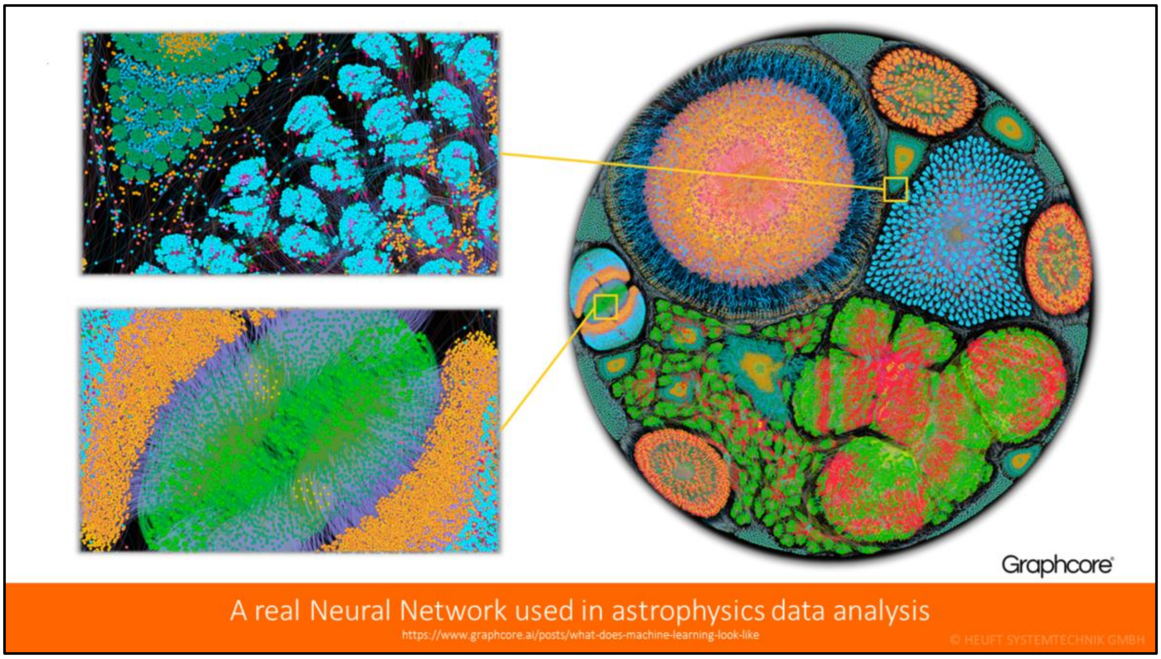

To give you an impression of the real complexity of real Neural Networks. Here you can see a Neural Network used in astrophysics data analysis. The fancy display has a purely aesthetic purpose. If you look at the enlargement, you will see that there are countless artificial neurons that are connected to each other.

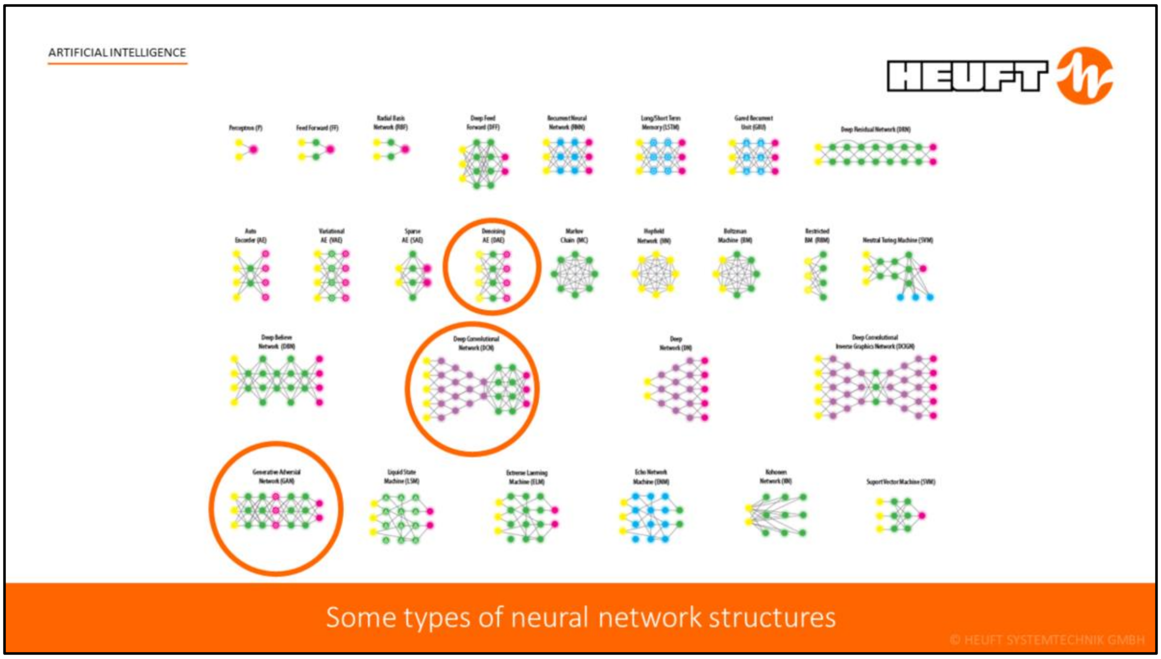

Designing and training a neural network is far more effort than the marketing hype implies. Even though the features no longer need to be extracted by humans, but the network must be designed to fit the task and the desired result. A Deep Learning Neural Network leads to better results for appropriate tasks. But it is not the optimal answer to all tasks. If I want to distinguish two labels based on color, a simple thresholding method is much faster and easier to set up and more efficient and reliable in operation. If, on the other hand, I want to find or distinguish things for which it is very difficult or impossible for humans to find features, a neural network may be a better choice. For example distinguishing small foreign objects from grain from image noise.

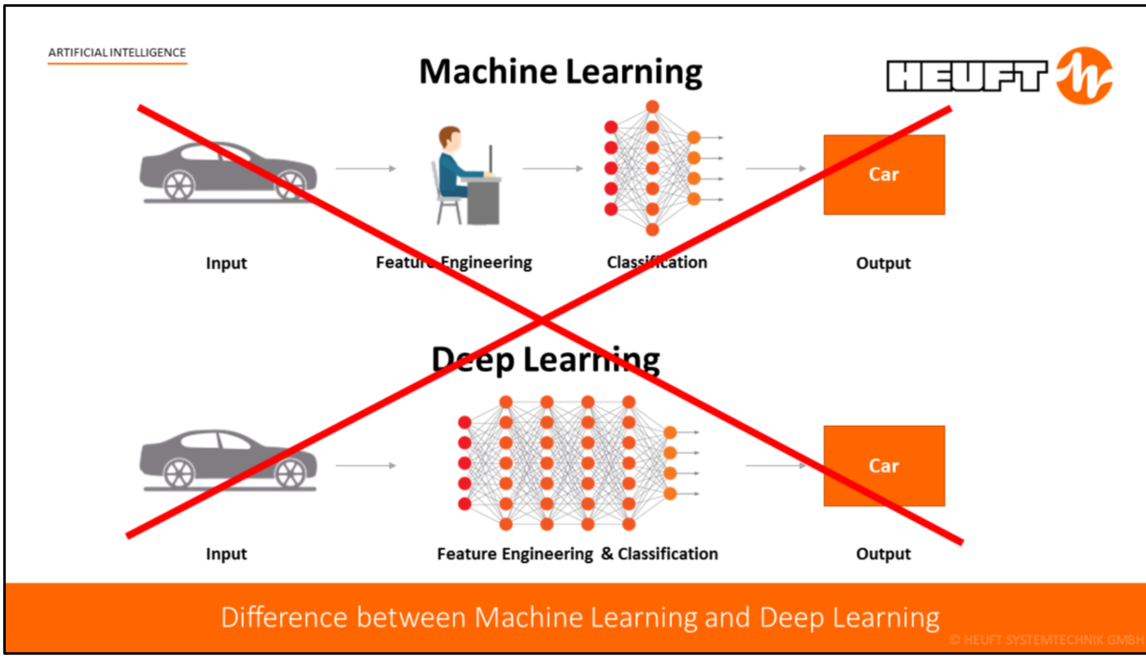

You can often find the following graphic to differentiate between Machine Learning and Deep Learning. But the illustration is misleading and suggests that the advantage of a neural network is that it works on its own without the intervention of a human.

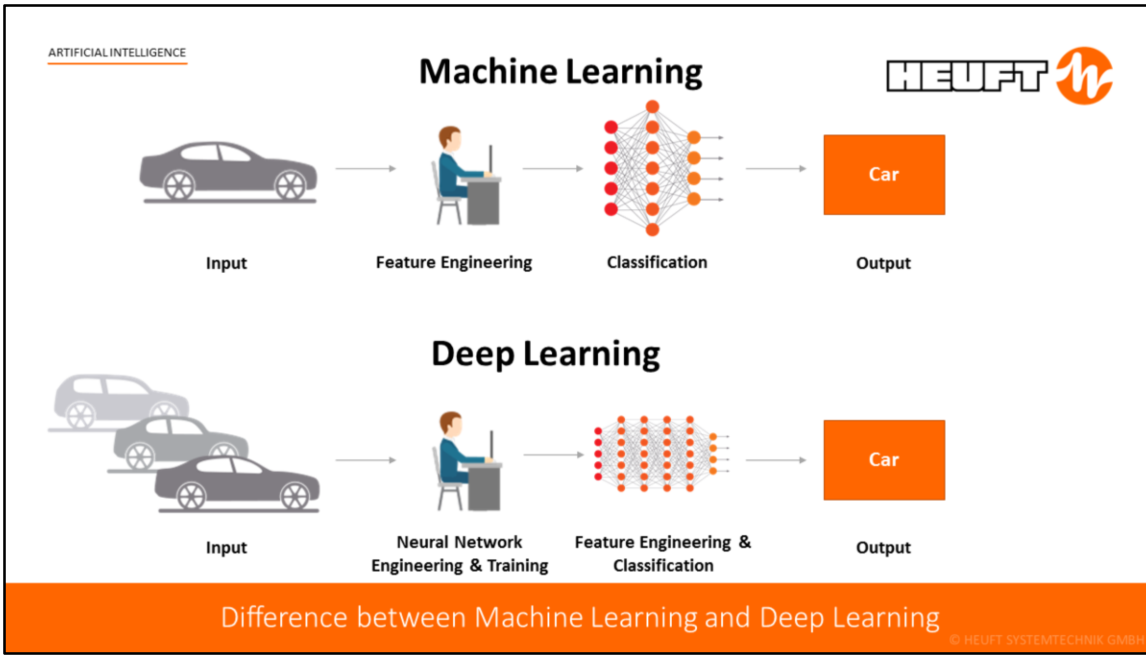

This illustration shows the engineering and training process for both techniques. During the use in production both technologies don’t need human intervention to do their job. What this illustration shows better is that Deep Learning needs way more training data (in this case car pictures) and engineering in the neural network which can be quiet complex. These disadvantages of deep learning are only outweighed in appropriate use cases.

They are not immune to mistakes! Not for nothing is the autopilot not yet a true autopilot. The driver is still responsible and must keep an eye on what is happening in order to intervene if necessary. When accidents happen the driver is responsible in court and not the artificial intelligence.

I know. The way of looking at it is drastic. On the other hand, the promises are tempting. "The device trains itself", "You don't have to do anything", "It sees everything". We wish it so much, but we have to remain critical. What is the point of an investment that does everything by itself, but may also do things we don't want unsupervised? Do you let your employees work completely unsupervised?

Would you drive in a driverless car today without steering wheel, and breaking pedal?! Without a way to intervene?

Even with the leaps and bounds of deep learning, AI is not immune to error. Deep Learning is a part of Machine Learning. It's one more tool in the toolset that when used for the right task is very powerful. A driverless car has to be supervised by the driver and so has an inspection device to be supervised by the operator. The judge will charge you as the driver or operator in the event of an accident, not the car or the device.

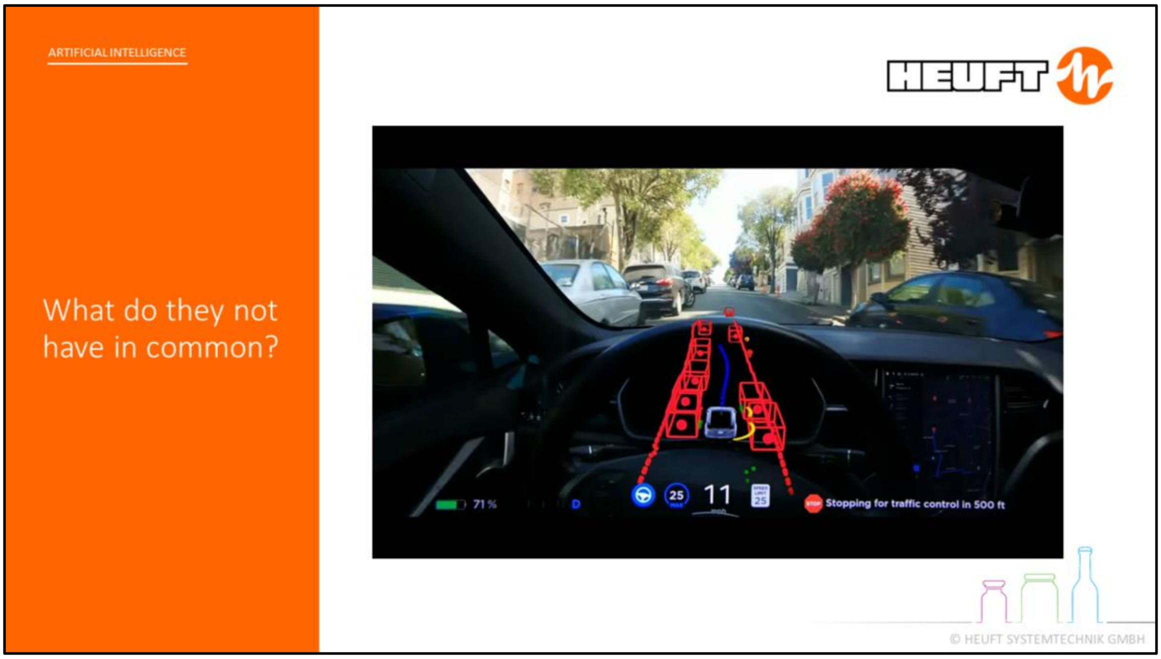

During the inspection we have exactly one shot of the container per perspective. The self-driving car reanalyzes the environment every few milliseconds and can correct previously made decisions without endangering the safety of the passengers.
It's easy to see how this works. When we as humans inspect a container for faults we pick it up, rotate it, look at it from many angles and make a decision from the sum of the views we see. We don't go blink once briefly and try to make a decision from the memory of what we saw. But this is what our inspection does because there is simply not enough time for a detailed analysis.
Seen in this light, the comparison between a self-driving car and an inspection machine is very misleading.

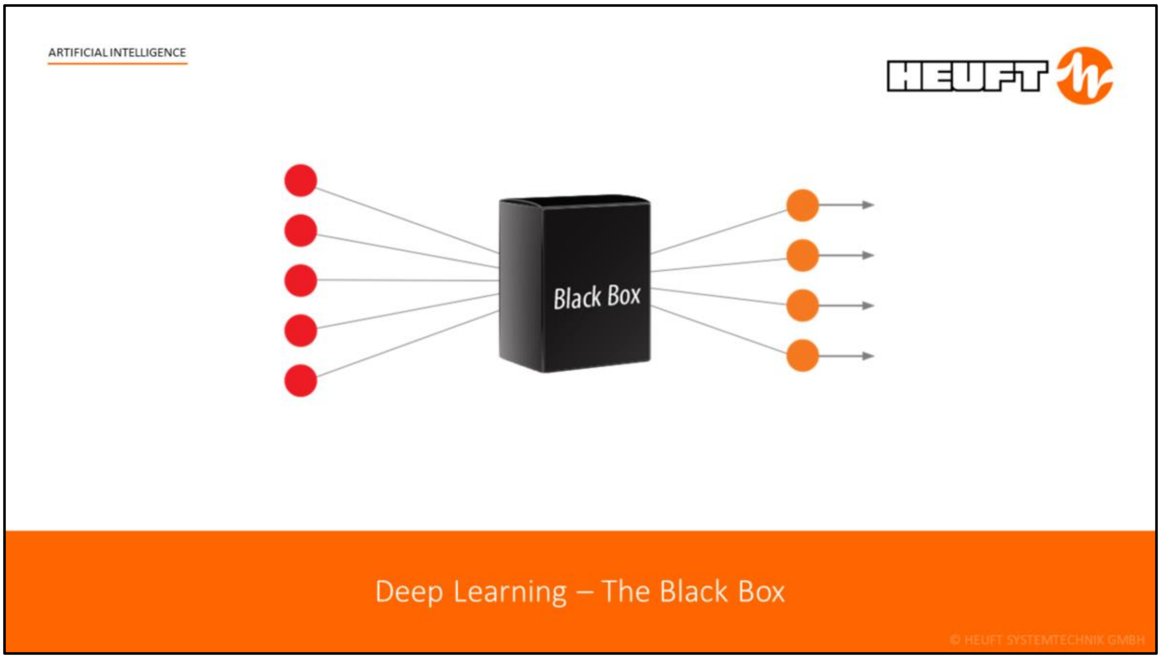

Now we come to one of the main problems of Deep Learning. It is a black box. So even when it has very high correct recognition rates most of the time, it can suddenly out of nothing surprise us with a false recognition result on an image that should be clearly recognizable. And then it gets tricky. An army of scientists is working on this problem. Basically one can say, the great advantage of the deep learning neural network to extract its own features from the image for its decision is its biggest curse due to its compexity when it comes to explain why it decided this or that way. This becomes a problem especially when you want to correct the fault. Since you can't tell exactly which nodes are responsible for the result, you can't simply adjust them. HEUFT is much more advanced in this respect with the HEUFT reflexx² image processing classifier.
Let me give you some examples.

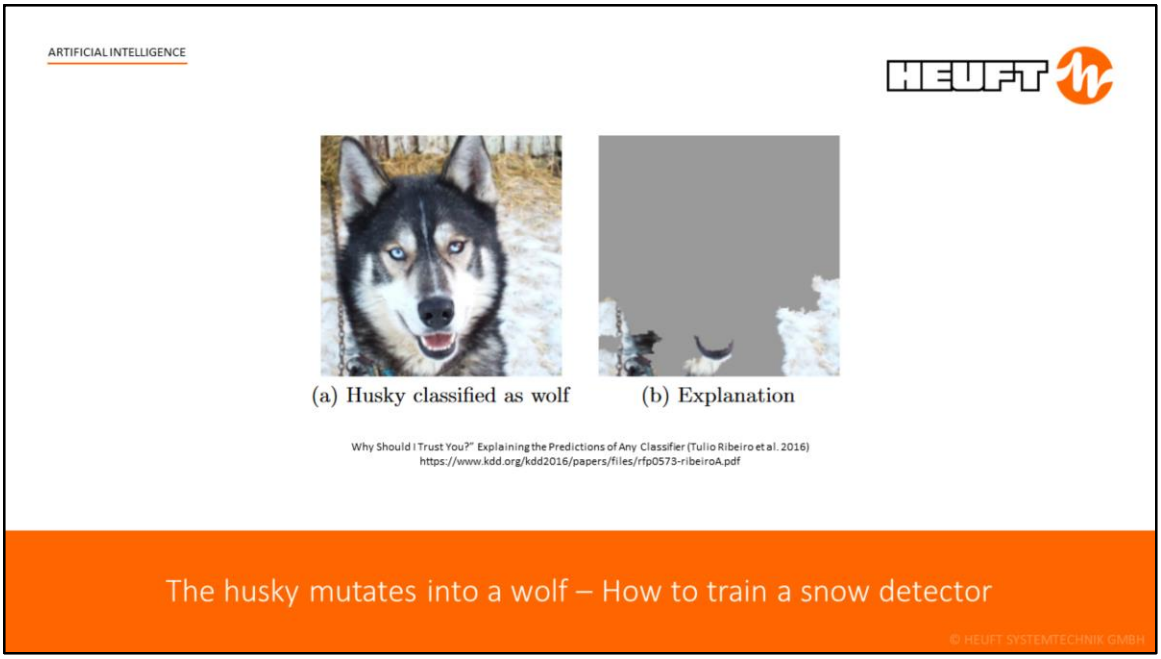

Here you can see an example of a neural network that was trained to distinguish huskies from wolves. The hit rate was very high and only one husky was recognized as a wolf. But why? As it turned out, the AI learned the snow in the background as a crucial distinguishing feature. Not the features (eyes, ears, fur, etc.) of the husky as we would think.

This sort of thing can happen quickly if you don't pay attention to such clusters in the data you learn with. If you look at the training images used, wolves are always shown in the snow. Since the AI does not know what the image is about (it is supposed to teach itself), it takes the background into account as much as the husky or wolf in the foreground.

Such errors can only be avoided by providing the AI with an extensive variance of learning data. Ideally, it has "seen" everything and learned what to judge. But this is a big problem. One would have to have recorded all huskies and wolves of the world at all places of the world for 100% recognition rate. For reliable detection rates you need way less but still a large number of combinations.

Transferred to our industry, this means that we need a large number of combinations of containers with faults to be able to perform reliable detection.

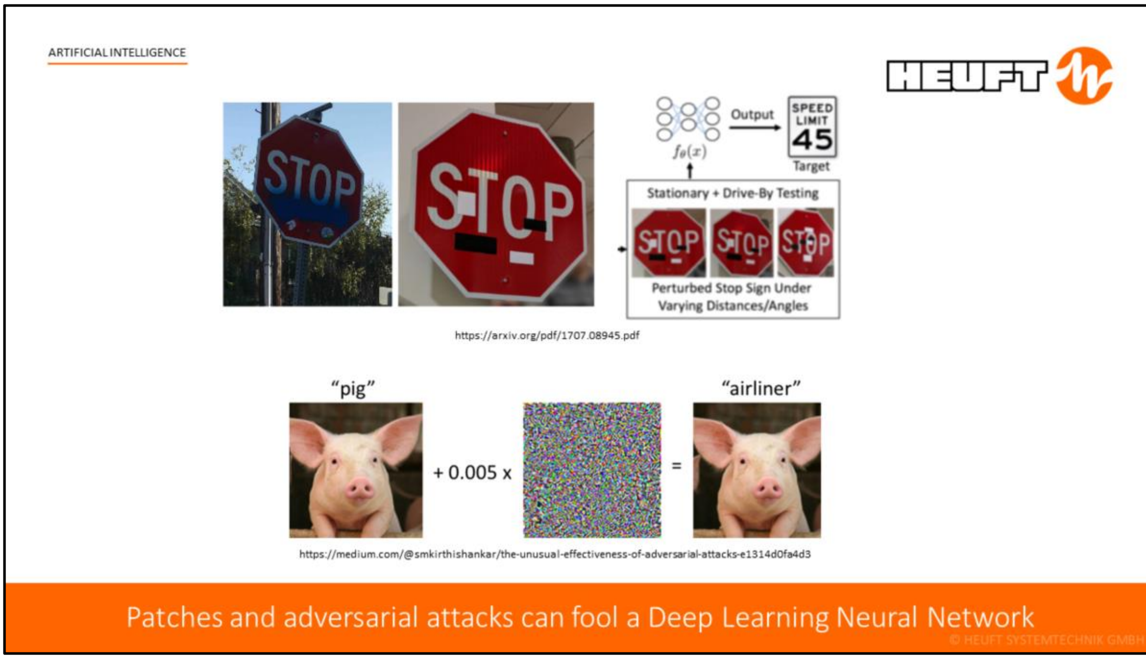

Here is another example. Spray paint on a stop sign makes it hard for the neral network to detect it reliably. During the analysis of this fault ways were found to trick the neural network in thinking it sees a speed limit sign simply by putting some black and white patches on the stop sign in the right positions. This shows that a neural network is „just" a bunch of numbers and weights that, when „stimulated" correctly lead to wrong results. It is all about statistics.

While the spray paint and patch method can actually become a danger in real life the shown adversarial attacks are NOT! You see a picture of a pig and by calculating a very special noise pattern and injecting it into the image the neural network can be tricked to think that the pig is an airliner. Again, this shall only show, that the recognition of an object in an image is not really comparable between a neural network and a human brain. It basically is a large statistics machine.

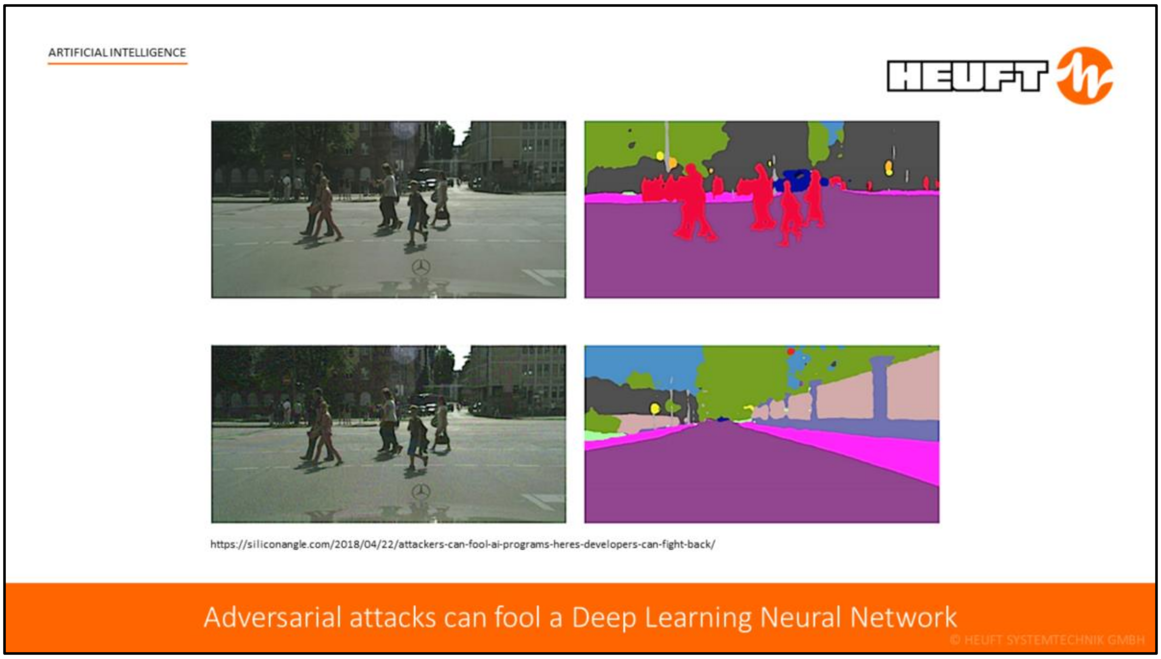

Here I show you an even more impressive adversarial attack. And again, this is NO real threat in the real life inspection process! You see a driverless car seeing a crossing. In the bottom view a special noise pattern again is injected into the image and tricks the neural network into thinking it looks at an open street and can drive.

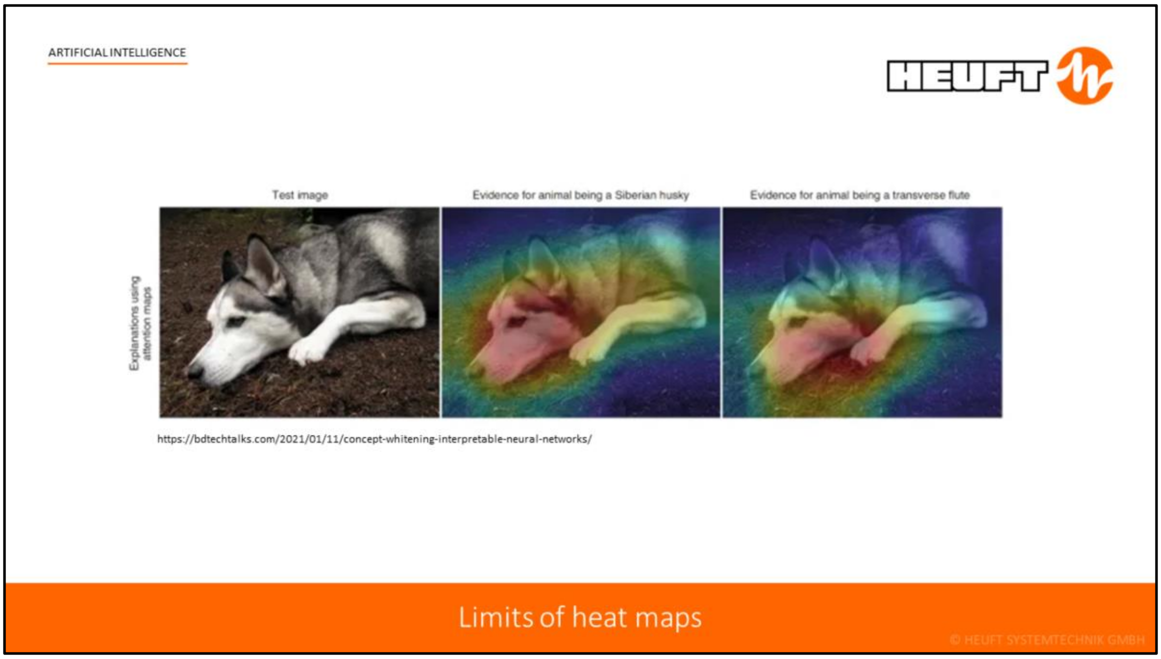

Heat maps do not provide accurate representations of how black-box AI models work. As you can see in the image the highlighted areas overlap significantly but are responsible for different results.

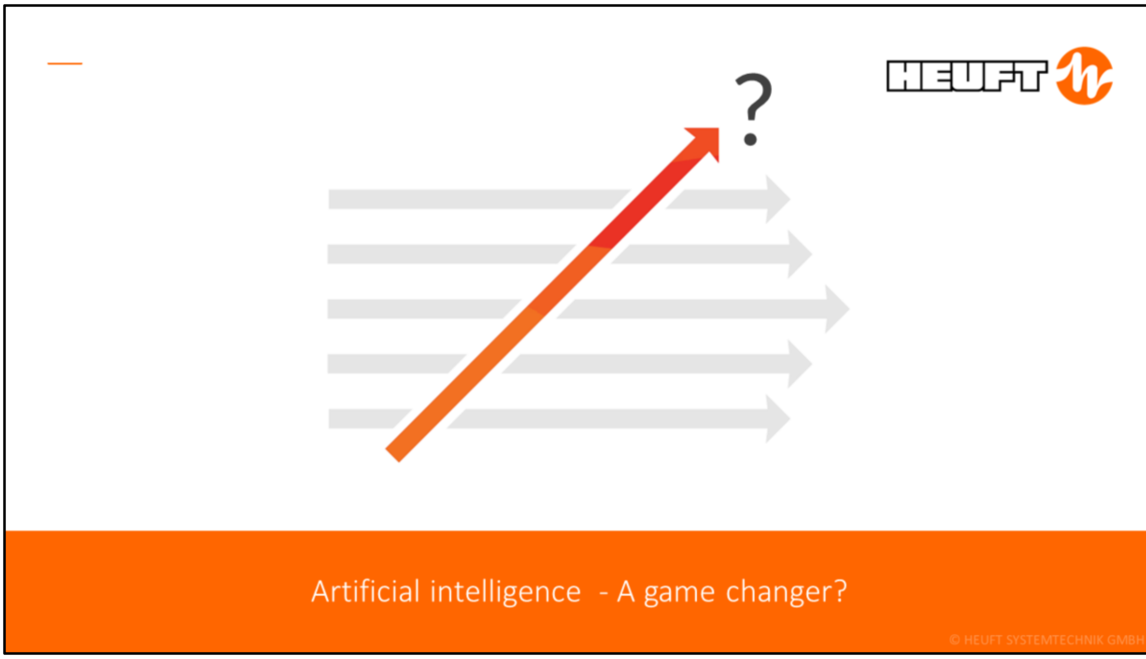

So, is A.I. and Deep Learning a game changer? Artificial intelligence, and in particular Deep Learning, brings many new opportunities. Nevertheless, many of the machine learning methods are well known, tried and tested, and in no way obsolete. The fact that driverless cars drive autonomously is due to the state of research of all machine learning methods and not to one single method. The real game changer is the availability of computing power. Only the development of the last 5-10 years makes such complex methods like Deep Learning practically usable. And only this development makes the methods usable to this extent in real time. This is where our own hardware and the know-how it contains show their strengths. Since 2010, we have been able to offer artificial intelligence based on machine learning methods on the HVPC2 in real time for inspection even in fast lines. The latest addition the HEUFT reflexx² camera is a game changer in terms of computing power, resolution and flexibility. More about this in the HEUFT reflexx² presentation. The next big step will be our extremely powerful hardware especially for the use of Deep Learning but also to further accelerate our Machine Learning methods in general. More on this in the HEUFT reflex A.I. presentation.

To be able to assess the situation further, we absolutely need continuous feedback from the market on the subject. Especially feedback from test installations as well as commercial and technical information is very interesting. Only in this way can we further investigate how capable the approaches of the competition are and whether the promises are kept in reality. This way we can then compile further material and information for you to support you in the discussion with the customer and to assess the situation.