Suppliers forge the relationship between AI and sensors

Deep learning technologies will play key roles in safely guiding AVs – but some, such as AEye, are considering alternative approaches to handling data for computer vision. By Xavier Boucherat

Technology based on artificial intelligence (AI), specifically deep learning, is only as good as the data that powers it. For the self-driving car, that data will arrive in part via sensor suites featuring any combination of cameras, radars and LiDAR. These are the inputs for any computer vision system, which in the future will help vehicles make the decisions normally made by humans, and trained by vast amounts of data.

“Based on what we saw at CES, most solutions on the market today are still based on massive, high-end hardware. We saw a lot of GPUs and servers on display. Our hardware is small, cost-effective and power-saving. Leaner, less expensive solutions are more likely to better meet automaker’s needs

At least, that’s one theory, and the industry appears to be pursuing it with vigour. “There is no better solution than deep learning to tackle computer vision problems,” says Adi Pinhas, Chief Executive and Co-Founder of Tier 2 start-up Brodmann17 which has developed deep-learning algorithms for use on both high-power and low-power CPUs. The upshot of this is the potential to do away with bulky graphic processing unit (GPU) systems, an alluring prospect for automakers already facing tight weight and packaging requirements.

“Based on what we saw at CES, most solutions on the market today are still based on massive, high-end hardware,” says Pinhas. “We saw a lot of GPUs and servers on display. Our hardware is small, cost-effective and power-saving. Leaner, less expensive solutions are more likely to better meet automaker’s needs.”

> "If you imagine a pedestrian holding a bag, that bag may cover up parts of his body required by the computer to properly classify him. There are limitations to classification, and it makes the deep-learning approach to detection problematic

For certain tasks, argues Pinhas, deep learning can mean computers achieving higher accuracy than the human eye. Some have voiced concerns about the efficacy of deep-learning to handle the sheer scale of the sensing challenge. Deep learning requires vast amounts of training for successful image recognition, and is extremely poor at filling in the gaps in its own knowledge. As The Verge reported in July 2018, studies have shown that even small changes in angle or lighting may throw off a computer vision system, resulting in incorrect classification of objects.

Pinhas is confident that with time, and input, the technology will keep pace. "Deep learning solutions are only familiar with the scenarios they are exposed to in the training process," he suggests. "In most cases, the training data involves broad daylight with clear vision, and so deep learning algorithm accuracy is limited to optimal conditions." However, he continues, this only represents 'Stage 1' deep learning, and the industry is rapidly advancing into 'Stage 2'. "We believe that with time and the right training," says Pinhas, "it will prove sufficient for computer vision tasks in autonomous vehicles."

Doran Cohadier, VP Business Development at Foresight Automotive, a stereo / quad camera vision system, is less optimistic. "If you imagine a pedestrian holding a bag, that bag may cover up parts of their body required by the computer to properly classify them," he says. "There are limitations to classification, and it makes the deep-learning approach to detection problematic." Furthermore, he adds, following the classification approach raises another issue: you cannot classify everything in the world. If, for example, a certain animal was not registered in the system, the computer might fail to acknowledge its existence entirely.

"There is interest from the industry in non-deep-learning algorithmic detection approaches," he says. "Our approach is multi-layered: we need to be sure we can detect everything, and this is the role of the first layer, which does not use classification. However, a second layer might then be added which does use classification. This can be used by an AV to make decisions, for example by determining whether an upcoming obstacle is human or not."

Humans do it better

Then there are those who believe that in order to bring sensing up to the remarkably high standard required for fully self-driving vehicles, certain assumptions at the heart of the AV sector must be questioned: namely, that more data is better, and that all of it can be effectively fused and processed at a core. Blair LaCorte is the President of AEye, an intelligent detecting and ranging (iDAR) developer exploring an alternative. The question, he asks, is what makes a car smart, and why can't today's self-driving vehicles travel over certain speeds, or outside of geo-fenced areas?

"We attribute intelligence to brain processing power," he says, "but in reality, intelligence normally starts with how you collect the data."

> "The industry is waking up to the idea that total AI, providing a full solution for autonomous driving, is far harder than previously thought. AI stacks can't be built deep enough whilst maintaining sufficient latency

Humans process some 70% of the spatial data they receive in the visual cortex, rather than sending the data to a central processing unit like a computer would - in our case, the brain. In effect, intelligence is pushed to the edge, and that is what AEye wants to offer via its system: a move away from post-processing orthogonal data. What success Mobileye has enjoyed, says LaCorte, has been down to its pre-processing capabilities, which allows for unnecessary data to be quickly discarded, instead of brought back to a core and processed. Today's computer vision systems, he suggests, dispose of as much as 90% of the sensor data collected.

"That means we've collected 90% of the wrong data, and that requires considerable power - up to ten times more than you would if there was processing at the edge," he says. "When you bring back masses of data, the computer has to know what that is. If you don't pay attention to that unnecessary data, you don't have to organise it, you don't have to encapsulate it, you don't have to recreate it."

> "Systems are collecting 90% of the wrong data, and that requires a lot of power – up to ten times more then you would if there was processing at the edge. When you bring back masses of data, the computer has to know what that is

As such, AEye is encouraging a new, selective approach to data. Instead of deep stacks for machine learning, LaCorte explains AEye talks in terms of shallow stacks, capable of quickly making simple decisions, such as ruling out non-essential objects such as trees and leaves. AEye's system therefore sees interplay between AI at the level of perception, using edge computing, and AI at the core, which performs classification tasks.

The approach is known as bio-mimicry: "People look at the spatial and temporal aspects of their situation simultaneously," he explains, "and they are constantly re-interrogating their environment. If you look at how the visual cortex has developed, the first thing someone does when they detect a threat is focus on the edges, vector and velocity of the object. They don't need to classify it - if it's big and moving towards me at speed, it's dangerous."

"The industry is waking up to the idea that total AI, providing a full solution for autonomous driving, is far harder than previously thought," he concludes. "AI stacks can't be built deep enough whilst maintaining sufficient latency. And so there are two avenues people are exploring. One is more intelligent sensors, such as our solutions. The other is to break the self-driving problem into pieces, namely specific applications. As such, advanced driver assistance solutions (ADAS) have become popular."