imec
embracing a better life

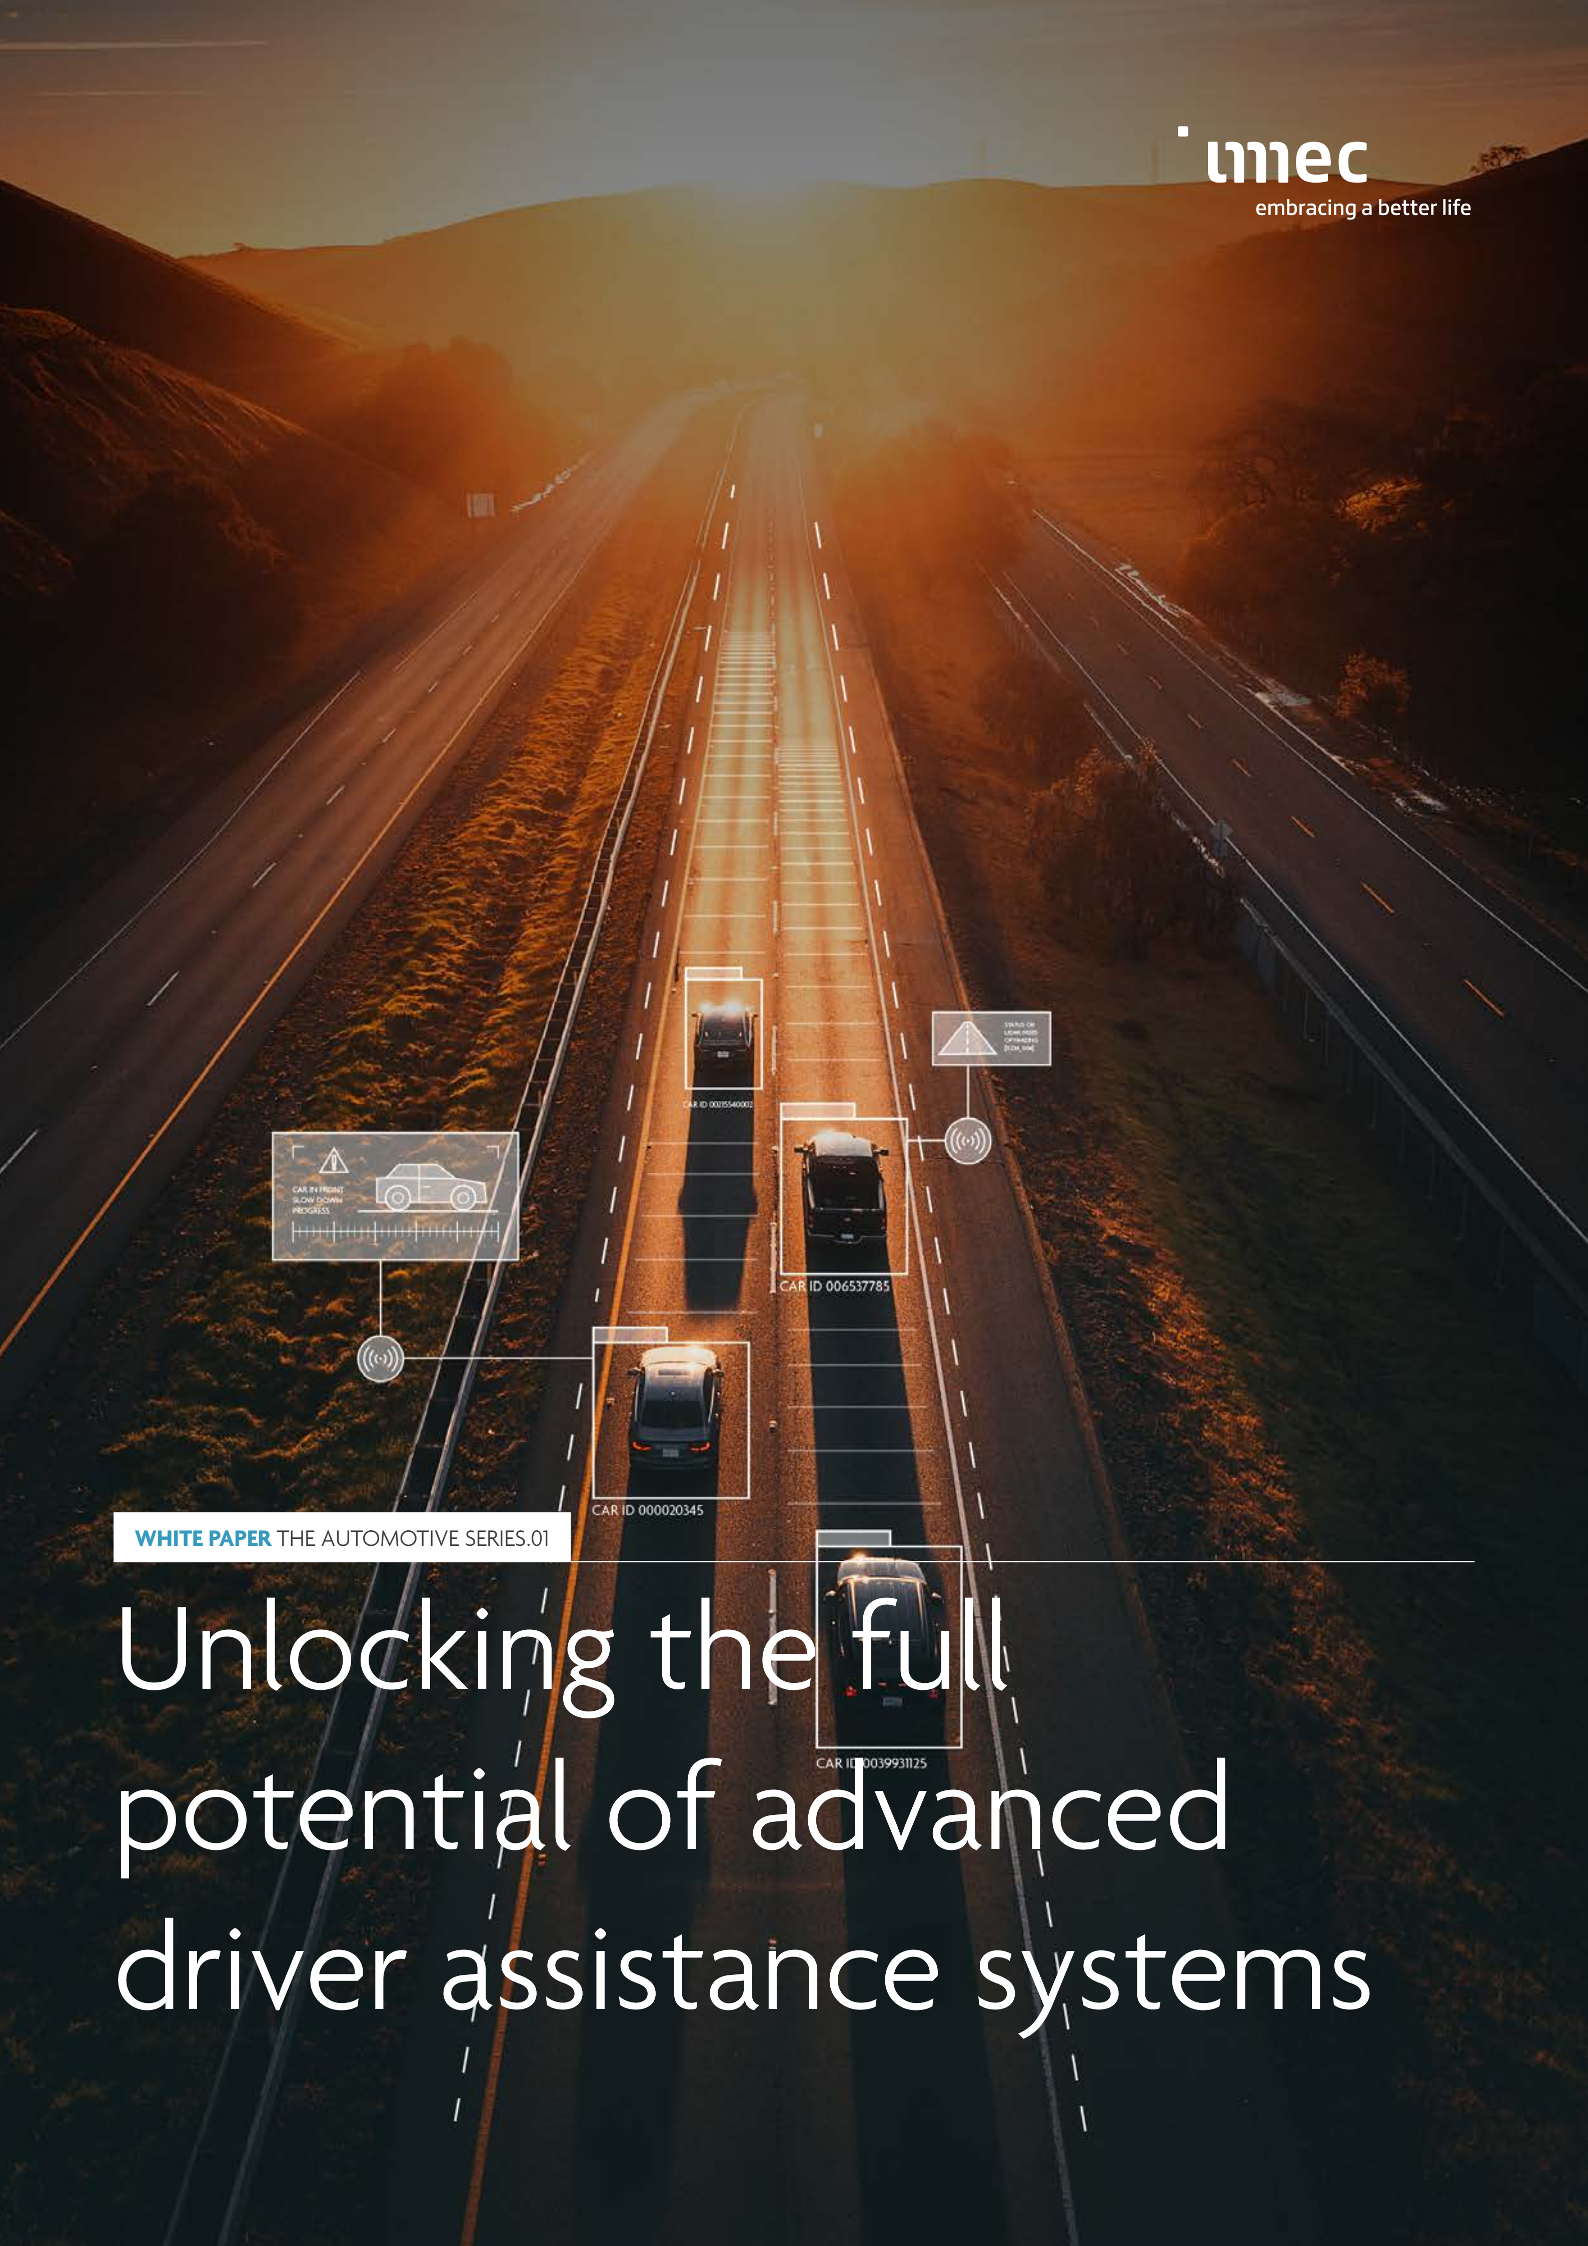

WHITE PAPER THE AUTOMOTIVE SERIES.01

Unlocking the full potential of advanced driver assistance systems

Contents

Who should read this white paper?

Imec's automotive white paper series is aimed at technical and business professionals from:

- original equipment manufacturers (OEM)
- car system developers (Tier 1)
- car component manufacturers (Tier 2)
- hardware and software companies with an interest in entering the automotive market

It offers them **state-of-the-art and objective insights** into the pathways and roadblocks that automotive companies are facing today. And allows them to take informed technological and business decisions.

Introduction

No one can say when autonomous driving will become a reality. One thing is certain: when it does, it will not be come from of one game-changing breakthrough.

The autonomous car will result from of a **series of improvements in advanced driver-assistance systems (ADAS)** that mitigate or prevent human driving errors.

This first instalment of imec's automotive white paper series gives you a thorough overview of the **current and future state of core ADAS technologies**. You'll find answers to questions such as:

- What are the benefits and weaknesses of **sensor technologies** such as cameras, radar and lidar systems?
- Which form of **sensor fusion** is most suited for mapping a car's surroundings?
- Which **AI developments** will make cars truly 'smart'?

Our goal is to offer you an objective and up-to-date overview of several nano- and digital technologies for the automotive sector. Of course, this is a fast-moving world, where new technological advances are continually made. We therefore urge you to **check the latest developments** by visiting www.imec-int.com/press and filter for 'automotive'.

If you want more information about our technologies, or if you are interested in working with us, feel free to contact Philip Pieters, business development director at philip.pieters@imec.be and he will soon be in touch.

Setting the scene

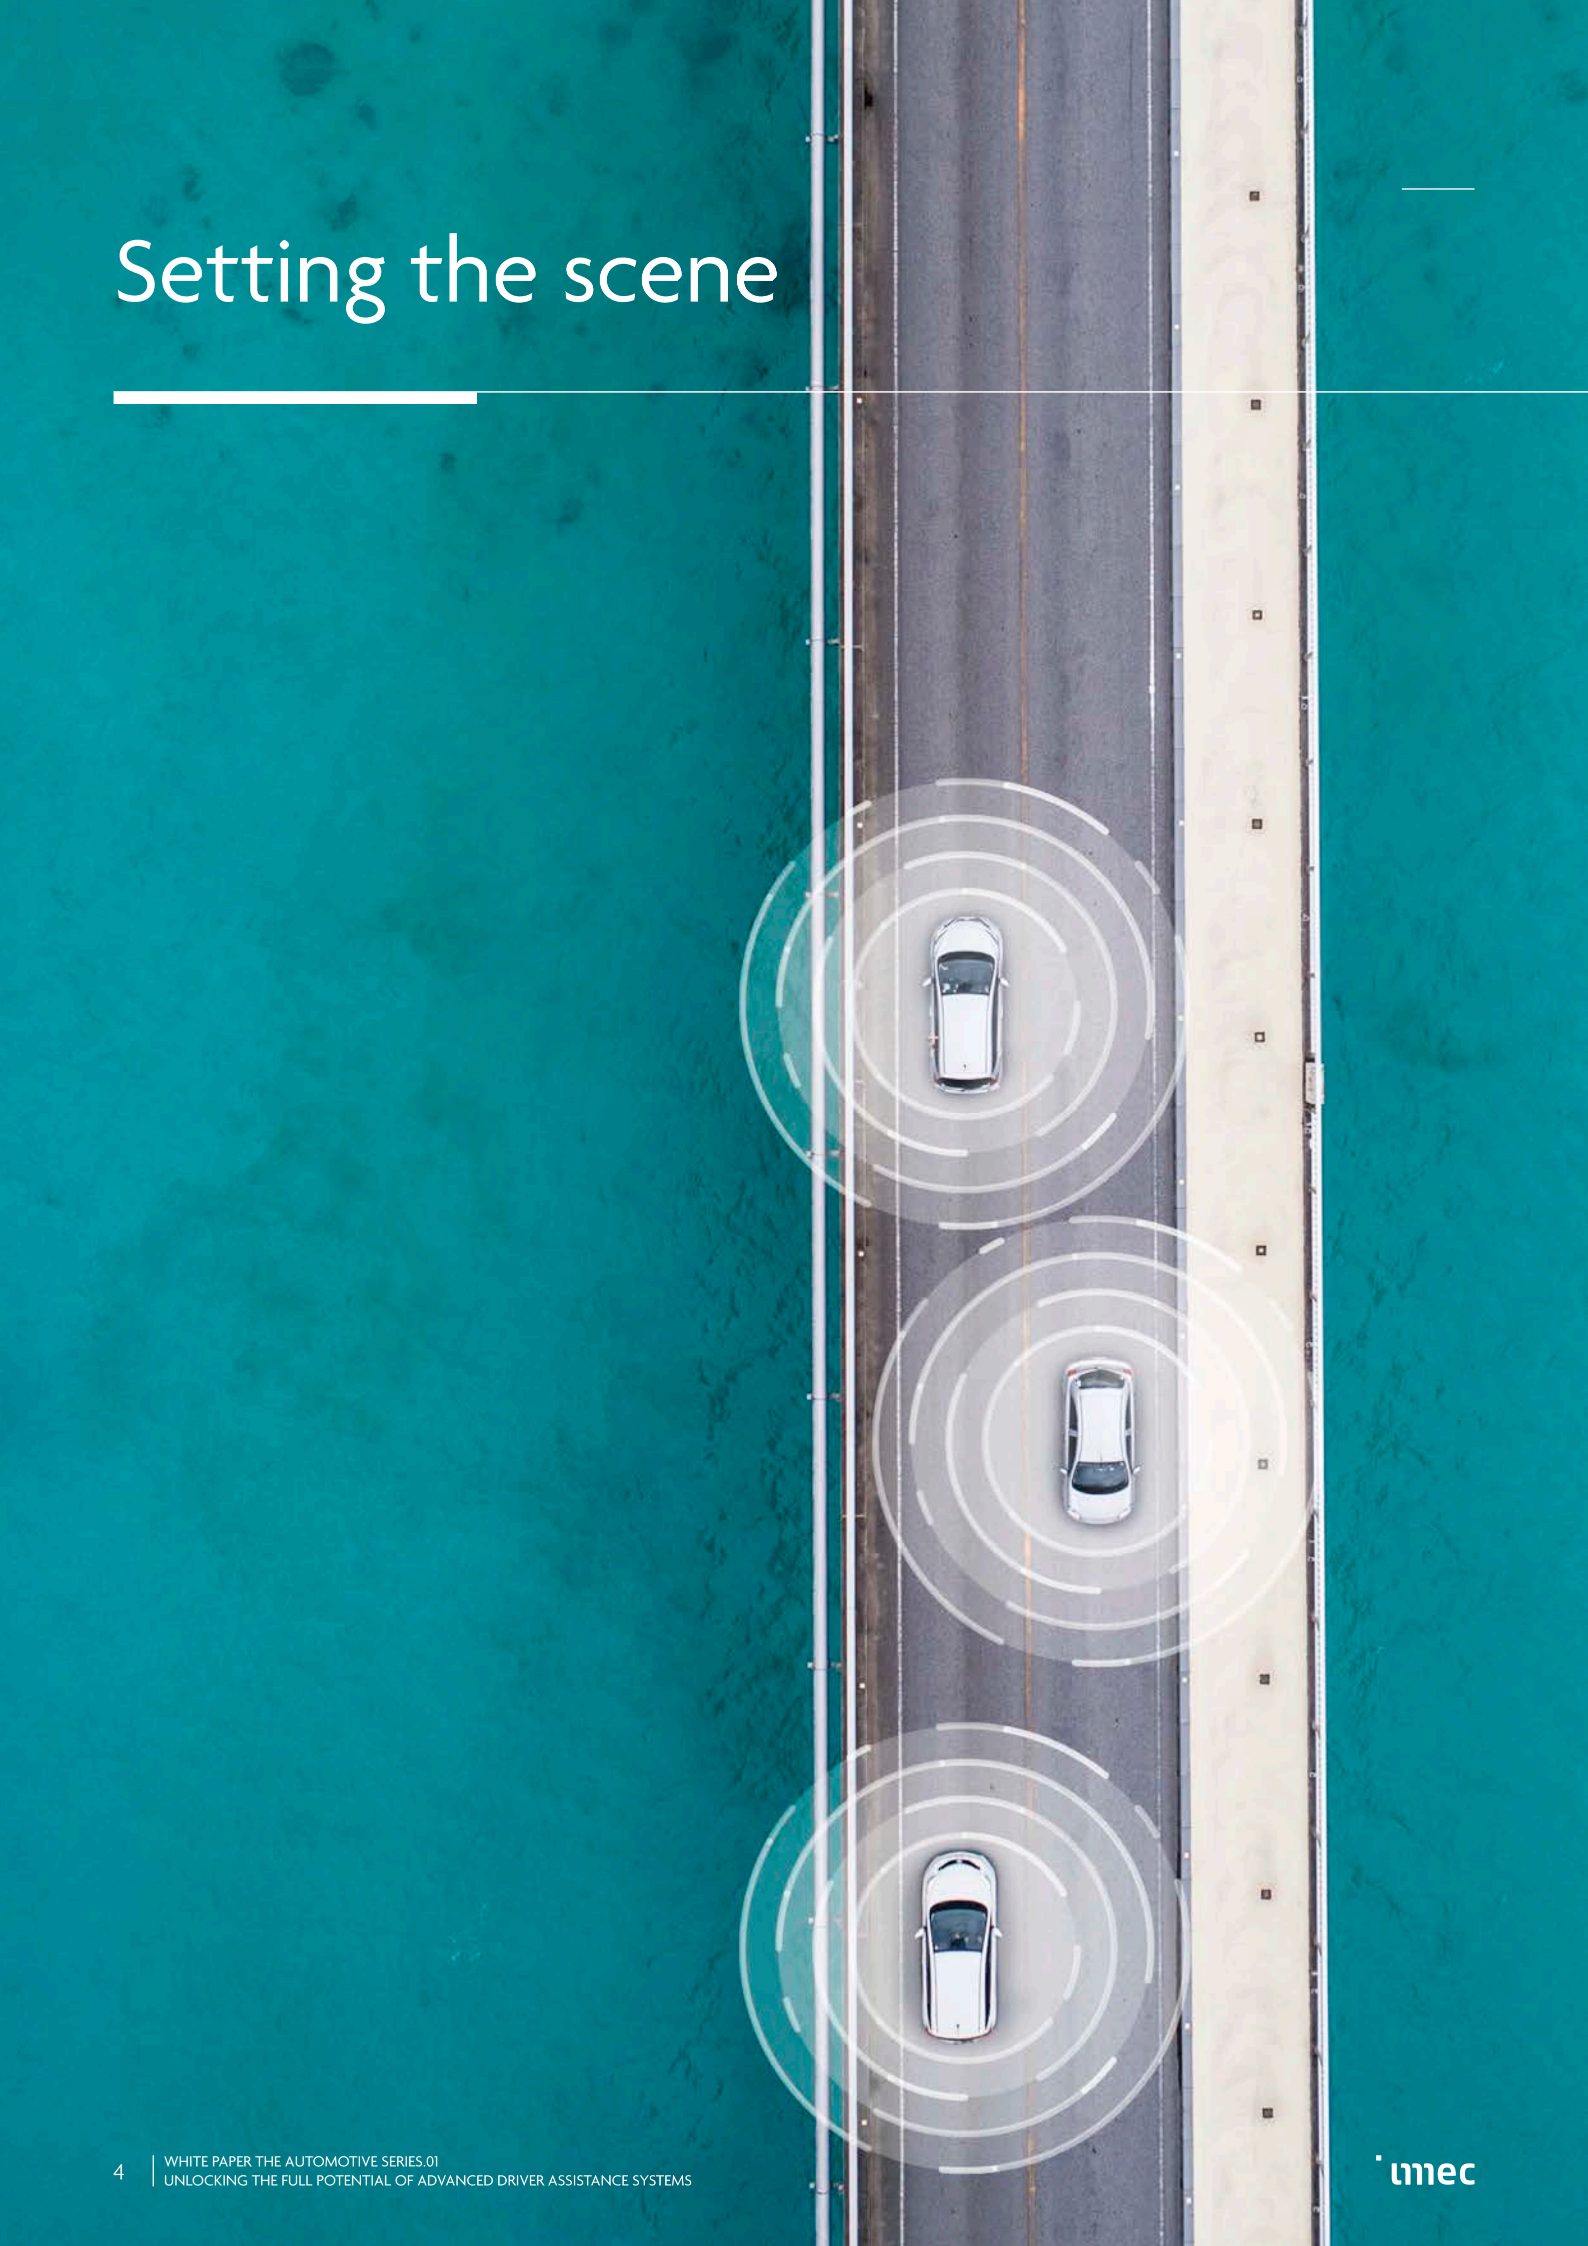

According to the World Health Organization, road traffic injuries caused an estimated 1.35 million deaths worldwide in the year 2016[1]. That is, one person dying in traffic every 25 seconds.

In response, a number of regions have started to make road safety a strategic priority. This, in turn, has allowed them to initiate a steady decrease in the number of casualties. In the EU, for instance, the European Transport Safety Council (ETSC) counted approximately 25,000 road casualties in 2018 – down from 55,000 in the year 2001[2].

It is clear, however, that road safety continues to be a problem child – with the latest figures seeming to indicate that some of our roads are becoming less safe again[3,4].

Hence, policymakers around the world have started to present even more ambitious road safety programs. The European Union's '**Vision Zero Initiative**', for example, aims at achieving a near zero fatality rate on EU roads by 2050[5].

Meanwhile, studies from the American National Highway Traffic Safety Administration (NHTSA) show that human error (e.g. speeding, fatigue and drunk or distracted driving) is at the basis of 94 to 96 percent of all motor vehicle accidents[6]. As a result – in order to meet their ambitious (zero fatality) road safety targets – governments, insurance companies, the medical community and other interest groups are eagerly awaiting the introduction of solutions that minimize or prevent the chance of human (driving) errors. Case in point: **autonomous vehicles**.

Yet, it might take a while before fully autonomous vehicles become an integral part of everyday life/traffic: they still require years of development and thorough testing. As a stepping-stone, though, car manufacturers have started to equip their cars

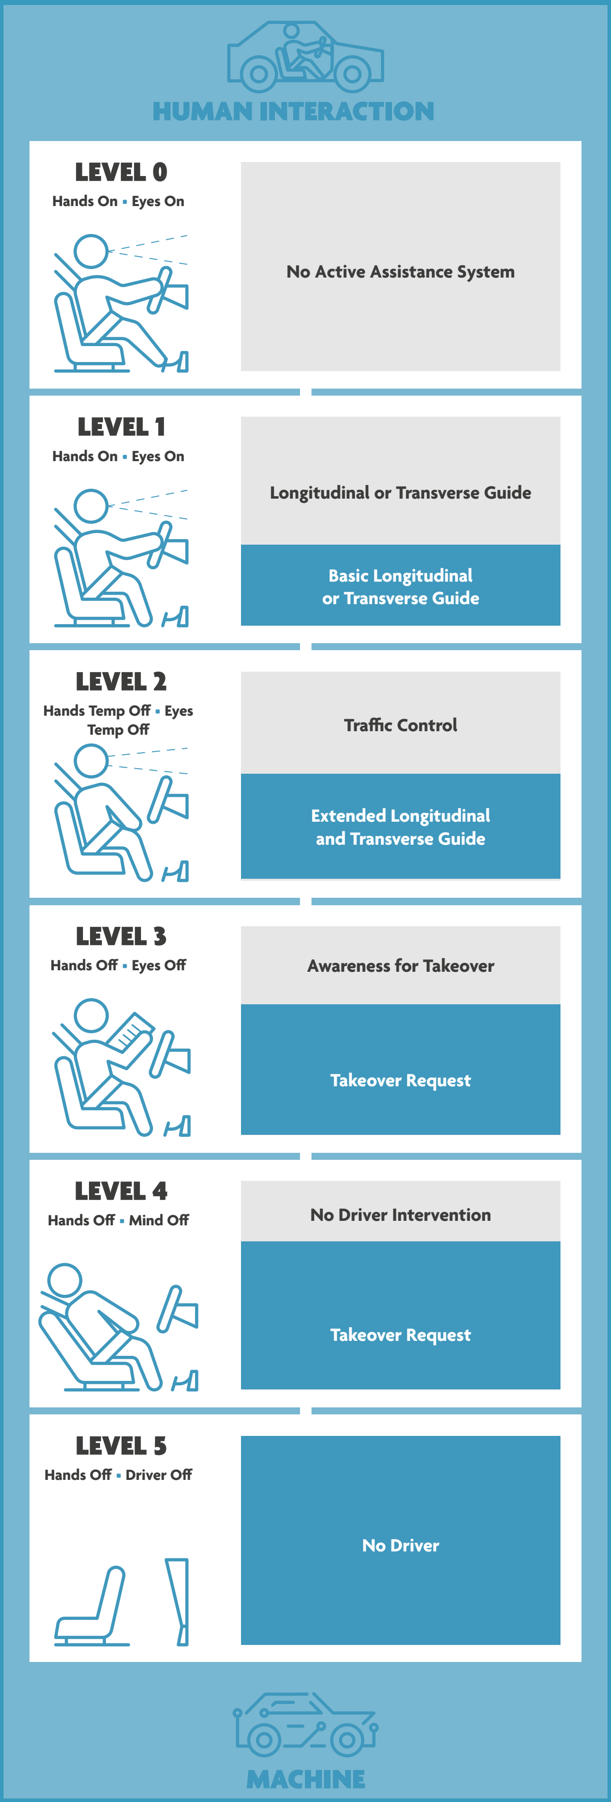

with the first generation(s) of **advanced driver assistance systems (ADAS)**.

ADAS have the potential to avoid or mitigate the severity of car crashes – through features such as **adaptive cruise control**, **lane keeping assistance**, **blind spot warning** and **pedestrian detection**.

ADAS (level 2) technology – such as lane keeping or semi-automated parking assistance – is becoming increasingly prevalent on new vehicles. In fact, according to data from the American Automobile Association (AAA), at least one ADAS feature is present in 92.7% of new car models available in the US as of May 2018 . And that is just a start: over the next decade, the demand for ADAS is expected to increase exponentially – not in the least since (as of 2020) both the EU and the US require all new vehicles to be equipped with safety features such as automatic emergency braking and forward collision warning.

In pursuit of the next level of car automation

While a number of basic building blocks are in place already, today's driving assistance technology continues to suffer from some important shortcomings – both from a hardware and a software perspective. For one, current automotive sensors lack accuracy (especially at longer distances). Secondly, sensors work in silos and have no real 'understanding' of a vehicle's surroundings. Finally, today's ADAS merely provide input: any action is still to be taken by a human driver.

In this first white paper of the automotive series – which highlights how the next generation(s) of ADAS can be brought within reach of the automotive industry – key advances in three intertwined ADAS-enabling technology domains will be discussed: **sensor systems**, **sensor fusion** and **artificial intelligence (AI)**. More specifically, this section describes how:

- Anti-collision sensor technology – from cameras to radar and lidar systems – can support higher resolutions and longer-range detection;
- A 3D mapping of a car's environment will be crucial to accommodate autonomous vehicles;
- Autonomous systems can be equipped with new forms of artificial intelligence to assess – and adequately react to – the intentions of other road users.

Next-generation (anti-collision) sensor technology: making the invisible visible

To do their magic, advanced driver assistance systems heavily rely on a combination of technologies – from processors and image processing algorithms to actuators. But everything starts at the level of a vehicle's built-in **sensors**, gathering information on road signs, pedestrians, cyclists, oncoming cars, etc.

Choosing the right sensor technology

In pursuit of ever more powerful and increasingly differentiating ADAS features, original-equipment manufacturers (OEMs) and tier-1 suppliers can opt for multiple sensor technologies and sensor suite configurations. Technologies range from camera / vision systems to ultrasound, radar and lidar – each supplying unique information on a vehicle's surroundings, yet each coming with its pros and cons.

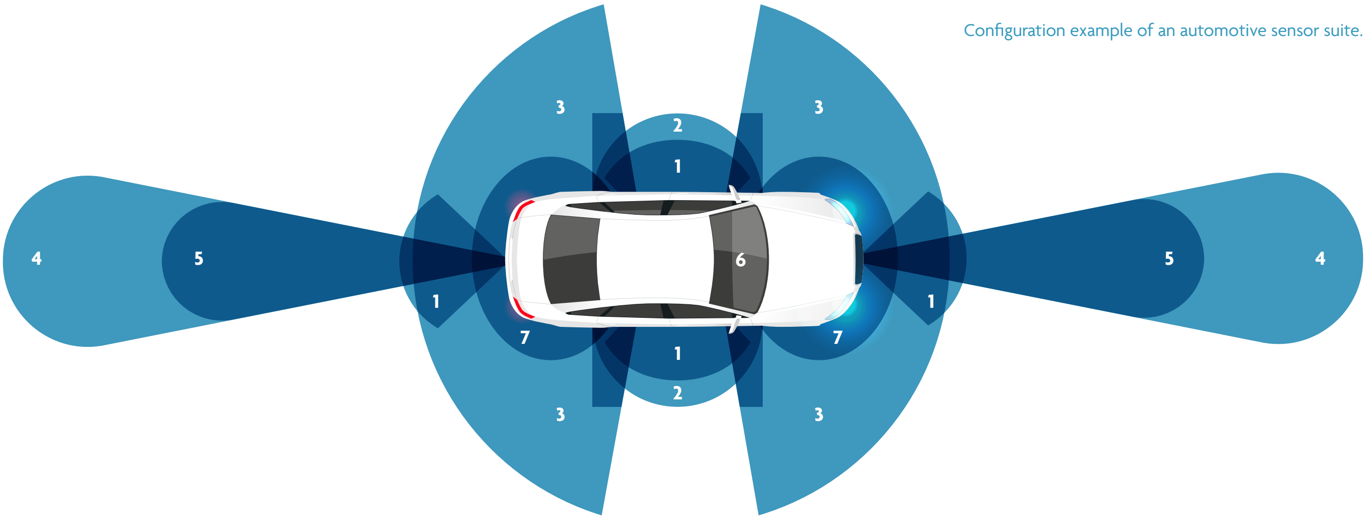

Configuration example of an automotive sensor suite.

1 CAMERA SYSTEMS

Using visual input to monitor a vehicle's surroundings; supporting ADAS features such as lane departure warning, surround view and (semi- or fully-) automated parking assistance.

Making use of the visible light spectrum (wavelengths from 380 to 740 nm).

- (+) Proven, reliable and widespread technology
- (+) Comes with mature (pixel-based) image processing algorithms

---

- (−) Suffering from poor lighting/ weather conditions (heavy rain, fog, twilight, facing the sun)
- (−) 3D imaging requires use of multiple cameras, in combination with simultaneous localization & mapping (SLAM) algorithms

2 - 3 - 4 RADAR SYSTEMS

Detecting and calculating the range, velocity and position of approaching objects; supporting ADAS features such as blind spot warning and adaptive cruise control.

Using the radio spectrum around a vehicle for object detection; typically operating in the 76 GHz to 81 GHz bands.

- (+) Robust (indifferent to lighting, only mildly affected by weather conditions)
- (+) Becoming cost-efficient (available as a silicon-based chip technology)
- (+) Working across all lighting (and most weather) conditions
- (+) Long range (up to 300 m)
- (+) Instant measurement of target speed and direction
- (+) Radar antennae can be integrated in prescribed car shapes
- (+) (Planned) support for 3D imaging

- (−) Limited angular resolution

5 LIDAR SYSTEMS

Using a laser scanner to generate a 3D image of a car's environment and to provide the exact location of objects in its vicinity; accommodating ADAS features such as dynamic driving assistance.

- (+) Excellent range (up to 150 m)
- (+) Excellent resolution
- (+) Best technology option to create detailed 3D maps of the environment

---

- (−) Expensive
- (−) Difficult to calibrate
- (−) Bulky, requiring external mounting
- (−) Limited performance in case of bad weather conditions (rain, snow, fog, dust)
- (−) Large amount of data to be processed
- (−) Full silicon-based solution not yet available

With so many options to choose from, it is impossible to decisively say whether the automotive industry will continue to rely on each of these sensor technologies (for safety/redundancy reasons, for instance), or whether some sort of consolidation/cost-optimization will ultimately take place. That being said, some clear trends can be discerned already:

- No one system covers all needs, scenarios and conditions. As a result, the power of next-generation ADAS will likely stem from a combination of sensor technologies and the inherent need for **sensor fusion** (using the input from two or more sensor systems – and overlaying that info for every scene).

- Within this ecosystem, camera systems will continue to play an important role – thanks to their versatility and proven accuracy in good weather conditions.

- Radar technology is not likely to disappear either, since it is commonly used in automotive already (to enable features such as automatic emergency braking, blind spot warning, adaptive cruise control, etc.).

- Main question for now is whether the automotive industry will resort to cost-optimized lidar technology to offset some of radar's current limitations (by moving to all-silicon lidar, for instance). Or will next-generation (silicon-based, on-chip) radar systems be able to cover all future automotive requirements and approach lidar angular resolution in both azimuth and elevation of less than 1 degree?

As of today, no consensus has been reached on this topic. As a result, different technology options – and mitigation scenarios – continue to be explored.

> "The power of next-generation ADAS will likely stem from a combination of sensor technologies."

Sensor fusion

Sensor fusion - the process of merging and analyzing data from multiple sensors – will be crucial for developing next-generation ADAS. A sensor fusion engine uses the images and point clouds – a set of data points in space, generated by a car's sensors – to create a perceptive (3D) model of the world surrounding that car. Based on that info – and leveraging **deep learning** approaches – detected objects are classified into categories (e.g. cars, pedestrians, cyclists, buildings, sidewalks). Based on that 'knowledge', intelligent (driving and anti-collision) decisions can be taken.

Sensor fusion goes hand in hand with increasingly stringent hardware requirements. For one, sensors need to be made more **energy-efficient** – a key requirement for integrating larger quantities of them into moving vehicles. In fact, this is one of the hot topics being addressed by the semiconductor industry. Leveraging advances in silicon technology and complementary metal-oxide-semiconductor (CMOS) fabrication processes, on-chip sensors can be created that feature a small form factor, physical robustness and resistance to electrostatic discharge – while being energy and cost-efficient (to mass-produce).

From an 'intelligence' perspective, sensors will need to be equipped with local **(edge) computing power** as to avoid that each piece of data needs to be sent to the cloud (and back) for processing and analysis before any action can be taken. In other words: embedding computing power on an (on-chip) sensor will allow tomorrow's smart vehicles to react to dangerous situations much more quickly. In this domain as well, the semiconductor industry – and more specifically the work that is ongoing in the domain of memory chips – is bound to play an important role. Today, experiments are already ongoing with various form factors, from graphical processing units (GPUs) and field-programmable gate arrays (FPGAs) to application-specific integrated circuits (ASICs).

The economics: new technologies come with a new business case

Equipping a vehicle with next-generation ADAS comes with a considerable price tag. Will consumers actually be prepared to spend several thousands of dollars/euros on the latest safety features? Possibly, the next generation of ADAS – and autonomous vehicles in the longer run – will be linked to a whole new business case, i.e. **mobility-as-a-service**.

Today, cars are standing idle on people's driveways for approximately 95% of their lifetime[8]. As a result, any investment in better equipped (and thus more expensive) vehicles will have to be compensated for by an increased utilization rate. Enters **car sharing**.

One possible scenario is that – in the future – people will no longer acquire their own (autonomous) vehicles but will rather opt for a (subscription-based) mobility-as-a-service package, with the right type of (self-driving) car picking them up whenever they need one. In such a scenario, the cost equation of building and maintaining cars will change completely: the investment will be spread over a large number of subscribers, meaning that the cost of innovation might no longer be a prohibitive factor.

Where today's sensors fall short

Camera, radar or lidar systems alike, each sensor technology comes with its limitations; limitations that need to be addressed in pursuit of the next level(s) of car automation.

As illustrated by means of the following examples, sensors' basic functionalities as well as their robustness need further refinement.

Camera systems

Challenge	Example	Mitigation scenario(s)
Easily affected by poor lighting and weather conditions	Camera systems based on visible light cannot be used at night and in harsh weather conditions (heavy rain, snow, etc.)	Start using **high-definition (HD) cameras** Investigate the use of short-wave infrared (SWIR) frequency bands that support longer-distance monitoring, and that can see through smoke or rain (e.g. the 1450 nm and 1550 nm bands) Develop more powerful image processing algorithms to extract more information from the available imagery Complement with radar systems

Radar systems

Challenge	Example	Mitigation scenario(s)
Suffering from low angular resolution and limited precision	Radar systems only have a basic idea of a vehicle's surroundings: they have been programmed to focus on the limited number of moving objects on a road, and to ignore stationary ones on the roadside Adaptive cruise control systems, for instance, will typically ignore stationary objects – such as a lamp post accidently lying on the road; this makes for potentially dangerous situations	Complement the use of 77/79 GHz radar with radar systems operating in higher frequency bands (e.g. 140 GHz or 300 GHz radar) to increase radar's angular resolution and/or form factor Investigate the use of imaging radar to increase resolution Develop new radar algorithms that are able to process doppler info (instead of the pixels generated by camera systems)

Bike and deer at 30 m range
79 GHz, BW = 750 MHz, nFreq = 161

15cm x 15cm

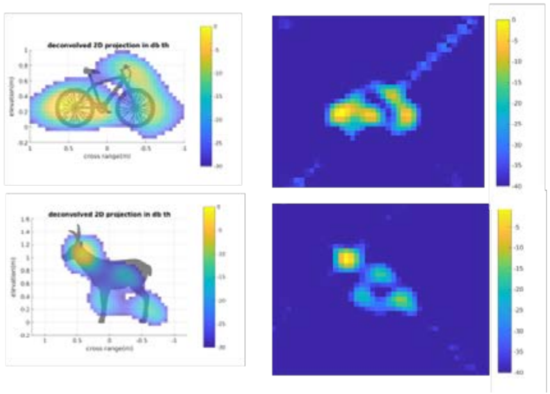

Bike and deer at 30 m range
140 GHz, BW = 750 MHz, nFreq = 161

15cm x 15cm

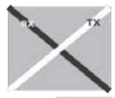

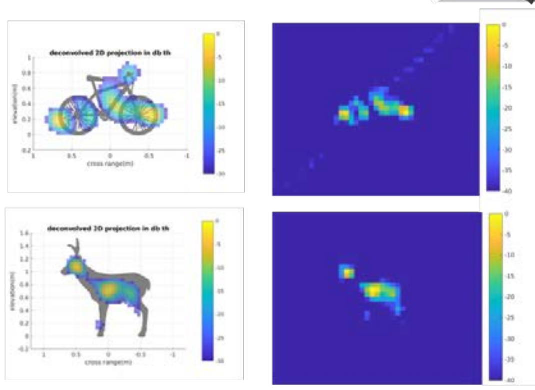

Comparing the resolution of 79 GHz and 140 GHz radar at a 30 m range (with a radar aperture of 15 cm x 15 cm). Source: imec

Lidar systems

Challenge	Example	Mitigation scenario(s)
Cost	The number of high-quality components that go into a lidar system makes these devices bulky, fragile and expensive	Integrate lidar technology into a silicon platform (i.e. solid-state lidar)

A closer look at three promising research tracks

Lots of research is currently ongoing in an effort to mitigate the limitations of today's sensor systems. As it is impossible to focus on all that is happening, let us have a look at three promising research tracks that are particularly relevant to the automotive industry: **SWIR camera systems**, **imaging radar** and **solid-state lidar**.

Short-wave infrared (SWIR) camera systems

Cameras making use of the short-wave infrared (SWIR) spectrum could offset some of the limitations of today's camera systems. Cameras operating in the 1450 nm band, for instance, support longer-range scanning while cameras operating in the 1550 nm band can see through mist, smoke and water vapor.

So far, however, these wavelengths have been transparent to silicon-based, on-chip imagers (a must-have when equipping vehicles with larger quantities of sensor systems) – since silicon photodiodes cannot detect light with wavelengths above one micrometer. Approaches that use III-V materials can overcome this detection barrier but are not available for consumer devices at an acceptable price point.

Today – thanks to the combined efforts of experts in material sciences, semiconductor process engineering and system-level design – some important breakthroughs are laying the foundation of silicon-based SWIR sensors that come at the price level of conventional CMOS imagers. Key enabler to achieve this are thin-film photodetectors (TFPDs): multilayer stacks with an overall thickness of a few hundred nanometers, with one of those layers being sensitive to the short-wave infrared spectrum. By post-processing these onto a Si-CMOS read-out circuit, the best of both worlds is combined: infrared detection via a CMOS-compatible process flow.

In the automotive domain, specifically, researchers are investigating the use of monolithically integrated TFPD stacks into the RGB pixel composition on the CMOS imager. In such a design, pixel-level multispectral sensors in the near-infrared and SWIR ranges can be added to the conventional red, green and blue photodiodes. This means a separate sensor for IR detection would not be required, reducing the system's price, footprint and power consumption.

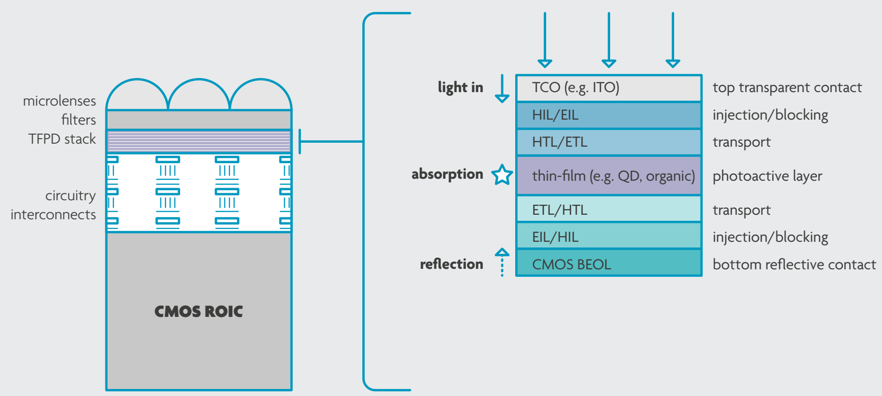

By processing a thin-film multilayer stack with photoactive layer that is sensitive in the infrared range (right), on top of a silicon readout circuitry (ROIC), an IR-sensitive CMOS imager (left) is created that is suited for mass-production.

Imaging radar

Imaging radar can be used to create a high-resolution 4D image of a vehicle's surroundings – at a higher angular resolution than what can be achieved by current automotive radar systems. As such, it can detect roadside obstacles as well as smaller targets (e.g. a person or a bike), even if these are partially hidden by a larger object.

Imaging radar holds great potential, not in the least since many questions remain on the economic feasibility of using (costly) lidar systems in automotive. That being said, imaging radar still comes with its share of challenges as well.

One such challenge relates to the use of the right beamforming approach – a technique to focus the radar's transmitter and receiver in a particular direction. Over the years, the industry has moved from analog to digital beamforming techniques. Imaging radar for automotive applications, however, might rather benefit from a **hybrid beamforming architecture**, combining the best of both worlds.

A second challenge (and opportunity) relates to the size of imaging radars, which can become large and difficult to integrate behind vehicle bumpers and fascia. A possible solution makes use of **structural electronics**, i.e. integrating electronic functions – such as radar antennae – into molded plastic structures (e.g. a car's bodywork). After all, design and aesthetics remain one of the primary concerns for car manufacturers: they require any (sensing) technology to be hidden. Within a highly tiered supply (and thus design) chain, new technologies and close collaborations are needed to come up with cost-effective and aesthetically satisfying solutions, such as the structural integration of (imaging) radar systems into a car's fascia or other subsystems, e.g. the headlights.

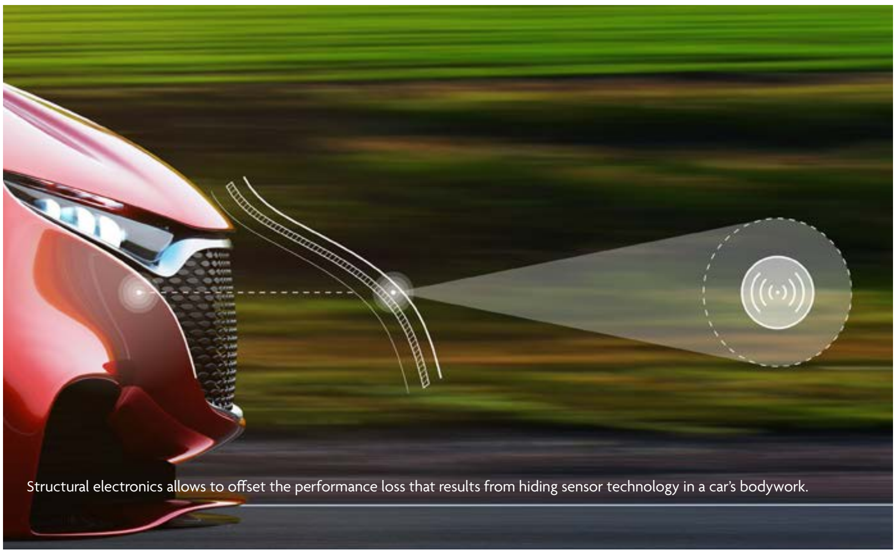

Structural electronics allows to offset the performance loss that results from hiding sensor technology in a car's bodywork.

“The lidar market remains very fragmented. This makes it difficult to objectively evaluate the performance of lidar systems and to make the right technology choices.”

Solid-state lidar

Using laser-based lidar technology, very detailed 3D maps of a car’s environment can be created. A major drawback, however, is that lidar technology builds on some very bulky and expensive pieces of equipment (costing several tens of thousands of dollars). On top, its moving parts make for a complex calibration exercise – with a laser scanner moving the (laser) light around and a detector capturing the light that bounces back.

Enters solid-state lidar: a lidar system built entirely on a silicon chip. No moving parts are involved, which makes the system more robust and allows it to be made smaller and cheaper (with predictions that the price could eventually fall below $200 per unit).

Research into solid-state lidar is well underway. Initial focus has been on the integration of the lasers and on the related optical beamforming approaches. In the next phase, the integration of the detectors needs to be tackled, while more focused work on the lidar chip’s output power will be required as well.

Yet, in the meantime, the lidar market remains very fragmented. Each vendor typically focuses on improving the performance of one specific component (source, MEMS detector, scanner, etc.). This, in turn, makes it difficult for OEMs and lidar systems suppliers to objectively evaluate the performance of lidar systems and to make the right technology choices.

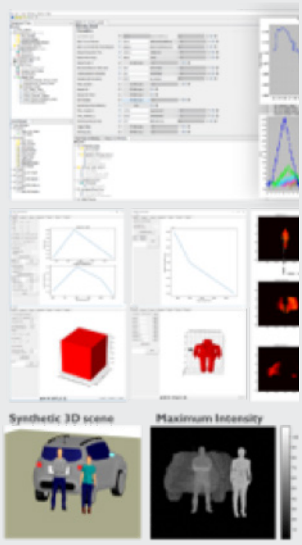

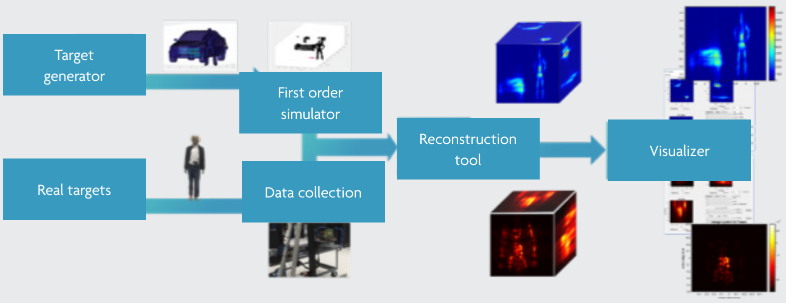

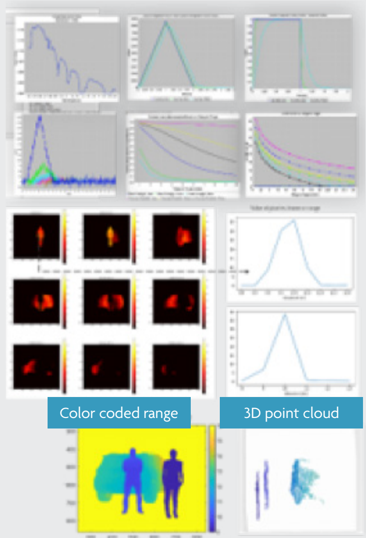

Modelling and simulation tools help automotive players evaluate a lidar system’s optical components (and their performance) and allow them to translate device-level requirements to a system and application level; this, in turn, enables them to decide objectively which system best fits their needs.

Imec covers the entire sensor systems research chain – from material science, photonics and electronics to chip and systems design, sensor fusion and machine learning.

Via imec, (automotive) companies can access the latest sensor technologies, put these to the test, check whether they comply with their specific requirements and use them to develop their own sensors for specific applications.

Imec can help companies explore and evaluate all technology options – based on objective evaluation criteria – and works with its industry partners to verify the potential of a given use-case. If this is of interest, companies can license the underlying imec technologies and building blocks afterwards.

Imec has pioneered the development of the digital radar sensor chip.

While most of today's wireless communication products use digital modulation to deal with interference and provide higher data rate, radar products still use analog modulation. The imec 79 GHz radar was the first prototype to support phase modulation and other digital wave forms. That opened many new dimensions to increase radar resolution and improve robustness to interference. Moreover, it can be implemented in standard CMOS technology, whereas radar products – so far – rely on more expensive specialty processes.

CMOS-based radar chips are:

- Compact
- Low-power
- Low-cost
- Supporting high resolutions

As a next step, imec researchers are aiming to develop a radar sensing solution that focuses on pedestrian detection – featuring a:

- 120° field-of-view
- Range of 30 meters
- Range resolution of 7.5 cm
- Maximum velocity of 50 km/h
- Velocity resolution of 1 km/h
- Angular resolution of 10 degrees
- Latency of 10 milliseconds

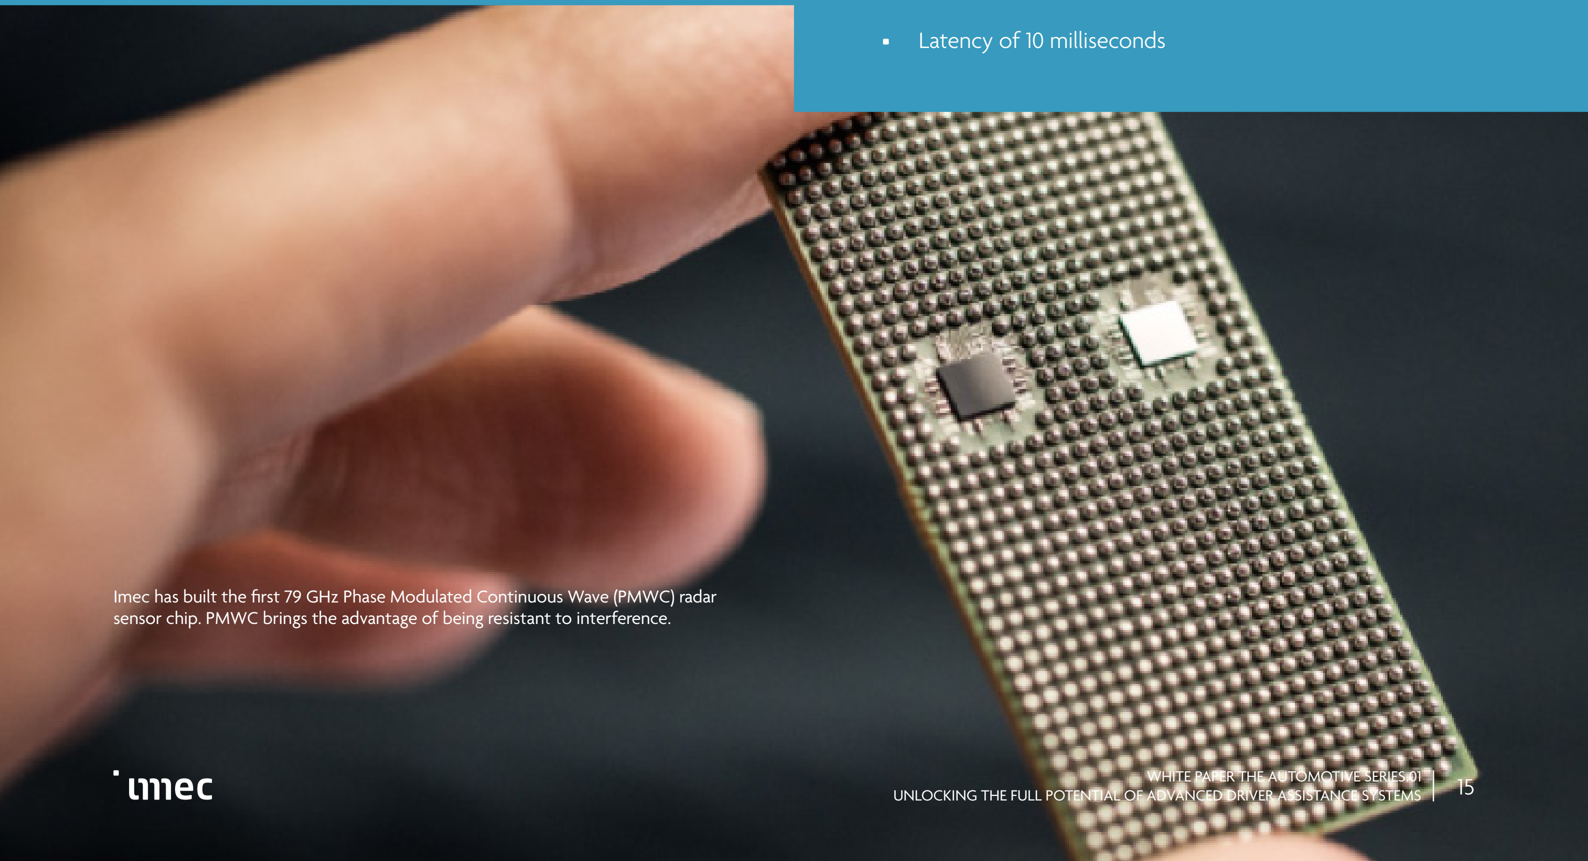

Imec has built the first 79 GHz Phase Modulated Continuous Wave (PMWC) radar sensor chip. PMWC brings the advantage of being resistant to interference.

Imec's solid-state lidar research aims to reduce the size of lidar systems by a factor of 20.

Imec's solid-state lidar research focuses on the efficient integration of semiconductor materials and devices into a silicon chip, without reducing the technology's efficiency or accuracy. Concretely, imec aims to reduce the size of conventional lidar systems by a factor of 20 (or more) and to reduce its cost from tens of thousands of dollars to less than 200 dollars.

Imec is developing on-chip optical beamforming technology that is up to three orders of magnitude more powerful than the milliwatt beams of existing solutions.

As the automotive industry has started to look into the potential of solid-state lidar for detection and ranging, optical beamforming technologies are needed that enable the modulation of the beam emitted by the on-chip antennae to form one precise, steerable beam.

Imec's optical beamforming program investigates the use of on-chip waveguides and phased antennae to manipulate the laser inputs and outputs – meaning that only a tiny integrated circuit is needed.

One of the key requirements to make this work, is for the (small) chip to emit a sufficiently powerful beam. While this is a problem for conventional silicon photonics systems, imec researchers are leveraging new materials and advanced patterning techniques that enable the formation of beams that are up to three orders of magnitude more powerful than the milliwatt beams of existing solutions (without any trade-off in resolution).

Advantages of imec's optical beamforming technology include:

- It is a full photonics solid-state solution, lacking any mechanical parts; as a result, it is much easier to assemble – while featuring a higher reliability;
- It supports a high optical power (30 W) with low beam divergence (<1°), suitable for long-range measurements;
- As it builds on semiconductor technologies, the system can be fully integrated with all electronics into a system-on-chip and/or a system-in-a-package;
- Leveraging semiconductor technology, its cost is drastically reduced;
- Imec can manufacture beamformers tailored to a company's specific needs.

imec

Sensor fusion and AI: making a vehicle 'understand' what its surroundings look like

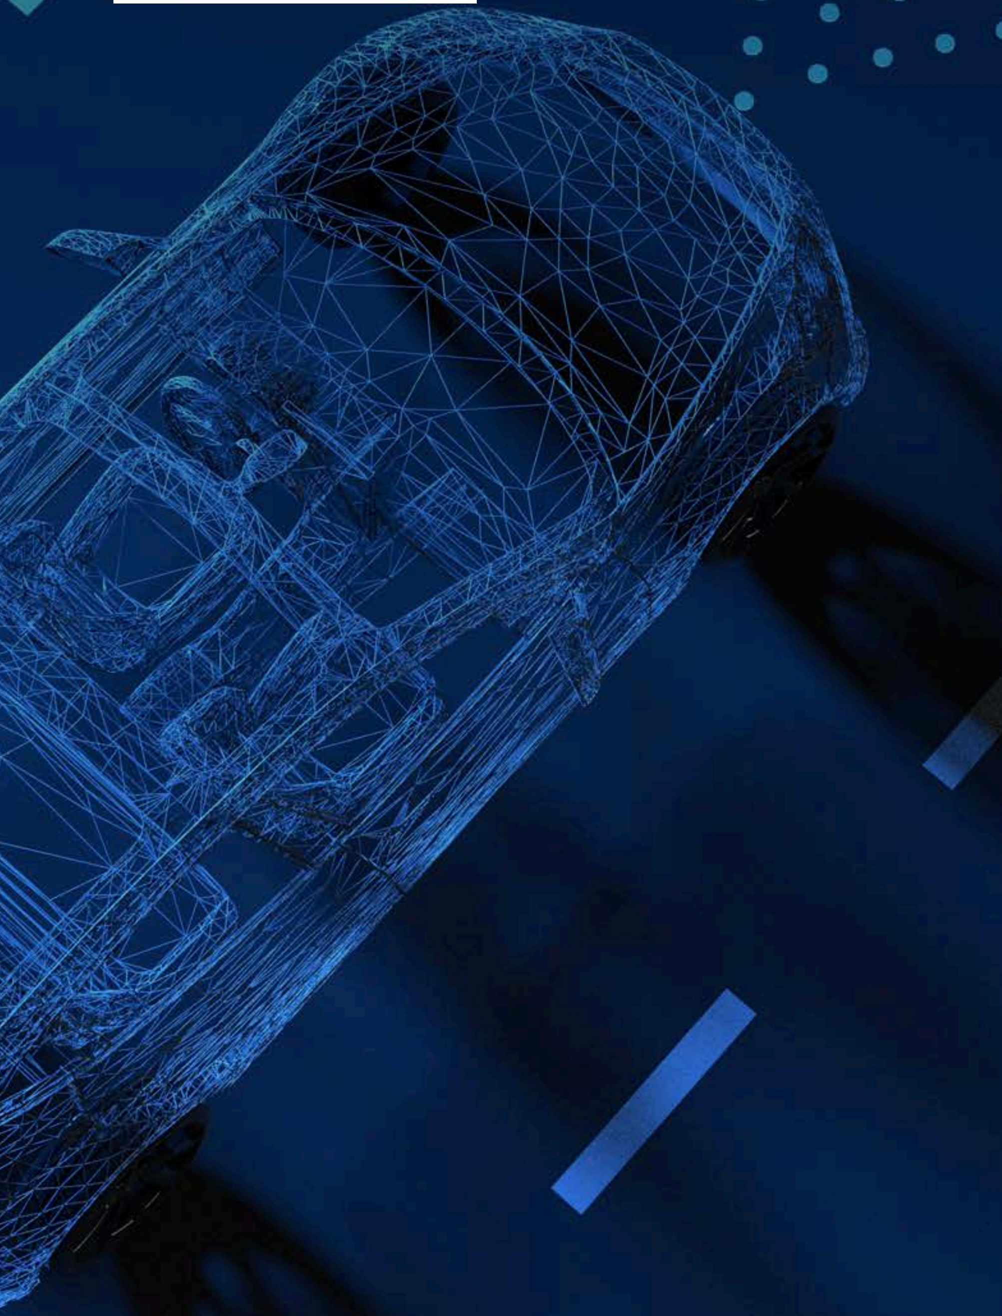

Accurate detection and tracking of road users and obstacles is essential to the introduction of driverless cars and many other smart mobility applications. As no single sensor can provide the required accuracy and robustness to cover all circumstances and scenarios, the output from several sensors will need to be combined. Enters **sensor fusion**.

When driving in poor lighting conditions, for instance, a vehicle's regular cameras only see what is being illuminated by the car's headlights. In such circumstances, it would be great to have the camera switch to the infrared (IR) or near-infrared (NIR) bands to improve its sensing capabilities.

Another advantage brought by sensor fusion is a much higher level of **redundancy**. As long as a human driver is still in control of the car, a false-positive (pedestrian) warning is not so much a problem; but that changes when self-driving vehicles appear on our roads. In such a scenario, sensor fusion allows the validation of various sensors' data to achieve the right level of **precision** and **confidence**.

Actually, the concept of sensor fusion is no longer a groundbreaking idea: virtually every automotive OEM has started to investigate its potential. But sensor fusion comes in various flavors – and the industry might very well be betting on the wrong one...

The inextricable link between sensors, sensor fusion and AI

Sensors, sensor fusion and artificial intelligence (AI) are inextricably linked. It is a combination of technologies that can help us overcome people's physical and cognitive limitations. While we suffer from limited night vision capabilities, for instance, an AI engine that gets its input from a combined camera and radar sensor system can analyze and interpret its surroundings in much greater detail – no matter if it is day or night. In a next step, leveraging those sensor data, the AI engine should be able to adequately respond to whatever it observes.

Actually, the concepts of sensor fusion and AI are no longer groundbreaking ideas: virtually every automotive OEM has started to investigate their potential. But sensor fusion and AI come in various flavors. That, by itself, warrants some further discussion.

Early fusion vs. late fusion: neither approach suits the automotive case

Today's most popular type of sensor fusion is called late fusion (as opposed to early sensor fusion).

Late fusion implies that sensor data are fused after each individual sensor has performed object detection and has taken its own 'decisions' (based on its own, limited collection of sensor data). Late fusion comes with the main drawback that each sensor throws away all the data it deems irrelevant. As such, a lot of sensor fusion potential is lost. In practice, this might cause a self-driving car to run into an object that has remained under the sensors' detection threshold.

Early fusion (or **low-level data fusion**) builds on all low-level data from every sensor – and combines those in one intelligent system that sees everything. Academically speaking, early sensor fusion makes for a very popular research path because it works with each and every bit of info that is available Yet, in practice, it requires high amounts of computing power and massive bandwidths – featuring high-bandwidth links from every sensor to the system's central processing platform. Moreover, practically speaking, it typically requires all equipment to be procured with one and the same vendor (causing vendor lock-in). So, in the end, early fusion does not really suit the automotive use-case either.

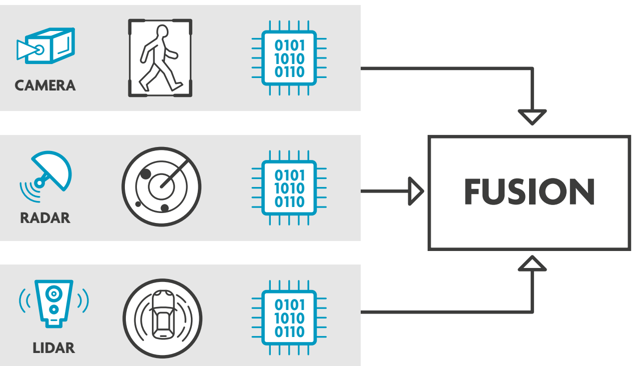

Late fusion: sensor data are fused after each individual sensor has performed object detection and has drawn its own 'conclusions'. Source: imec.

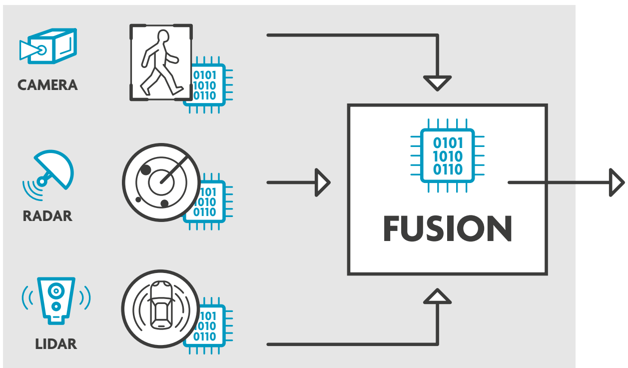

Early fusion builds on all low-level data from every sensor – and combines those in one intelligent system that sees everything. Source: imec.

Cooperative sensor fusion: the new kid on the block

In response to the shortcomings of traditional sensor fusion approaches, researchers have developed the concept of **cooperative sensor fusion**. It features an extra feedback loop, with different sensors exchanging low-level or middle-level information to influence each other's detection processing. If a car's radar system suddenly experiences a higher degree of reflection, for instance, the threshold of the on-board cameras will automatically be adjusted to compensate for this – thus improving overall tracking performance.

First results show that this method is much more powerful than the late fusion approach that is commonly used today. Moreover, it is easier to implement than early fusion since it does not come with the same bandwidth issues and practical implementation limitations.

Cooperative sensor fusion closely couples the detector and tracker. This allows to exploit the benefits of early fusion of the rich sensor data without the need for high-throughput data links to the fusion center.

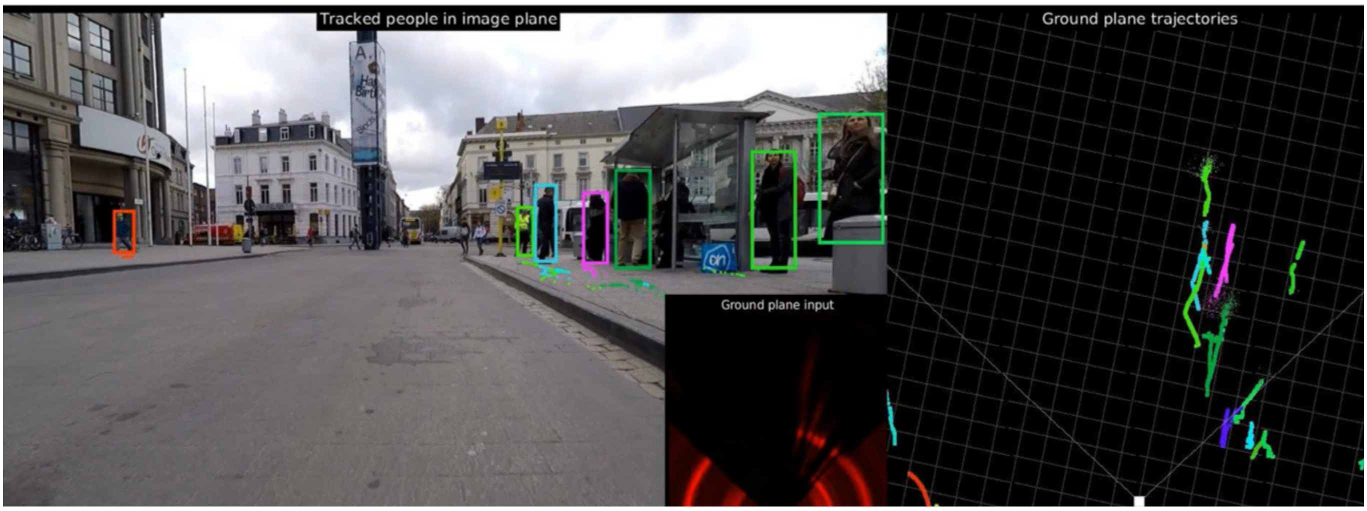

Field tests in the City of Ghent (Belgium), 2019 – featuring cooperative (radar and video) sensor fusion. Source: imec.

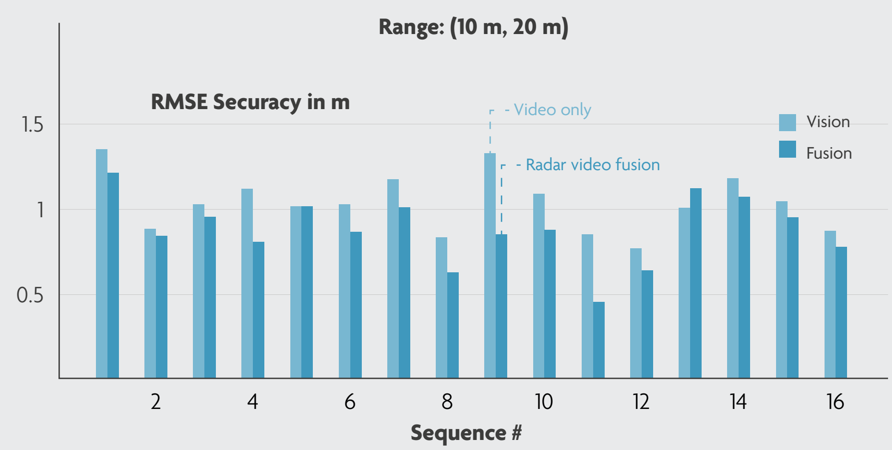

Demonstrating cooperative sensor fusion's accuracy, researchers captured a dataset consisting of various complex traffic scenarios in a European city center. For multiple targets up to 10 meters in front of the vehicle, they measured accuracy improvements of 20% over a camera-only system, with only a fractional load to the network compared to low-level fusion. Source: imec.

Video and radar make for a good sensor fusion match

As discussed above, every sensor technology comes with its pros and cons. For instance, cameras don't work well at night or in dazzling sunlight. And radar can get confused by reflective metal objects. But combining the two technologies makes for a good match, because their strengths and weaknesses complement one another.

It is expected that the use of camera systems as a sensor technology will remain fundamental. Ideally, however, they should evolve to a single box that features high dynamic range (HDR) capabilities (supporting a wider range of contrasts) as well as a night vision component (IR or NIR – which is less expensive than its IR counterpart). Those technologies could then be fused on a low level to end up with a single image.

Secondly, camera technology will need to be combined with a distance/depth perception system. Today, the only technology for depth sensing that is widely accepted for automotive use cases is radar – as it is a mature technology and well understood. However, it suffers from low angular resolution and limited precision.

As a result, efforts are ongoing to either squeeze more out of today's (relatively cheap) radar systems – or to accommodate new capabilities using different frequency bands (think of pedestrian detection). Apart from innovations from a hardware perspective, this goes hand in hand with the development of **machine learning technology** that does a better job at distinguishing objects from one another (e.g. pedestrians from road signs or other metallic objects with a lot of reflection).

In parallel, powerful lidar systems for distance/depth detection are getting cheaper. Research indicates that **flash lidar,** for instance, makes for a more cost-effective and robust alternative to present-day lidar. Flash lidar replaces rotating lasers with an array of light emitting diodes and uses beam steering approaches to scan a vehicle's environment.

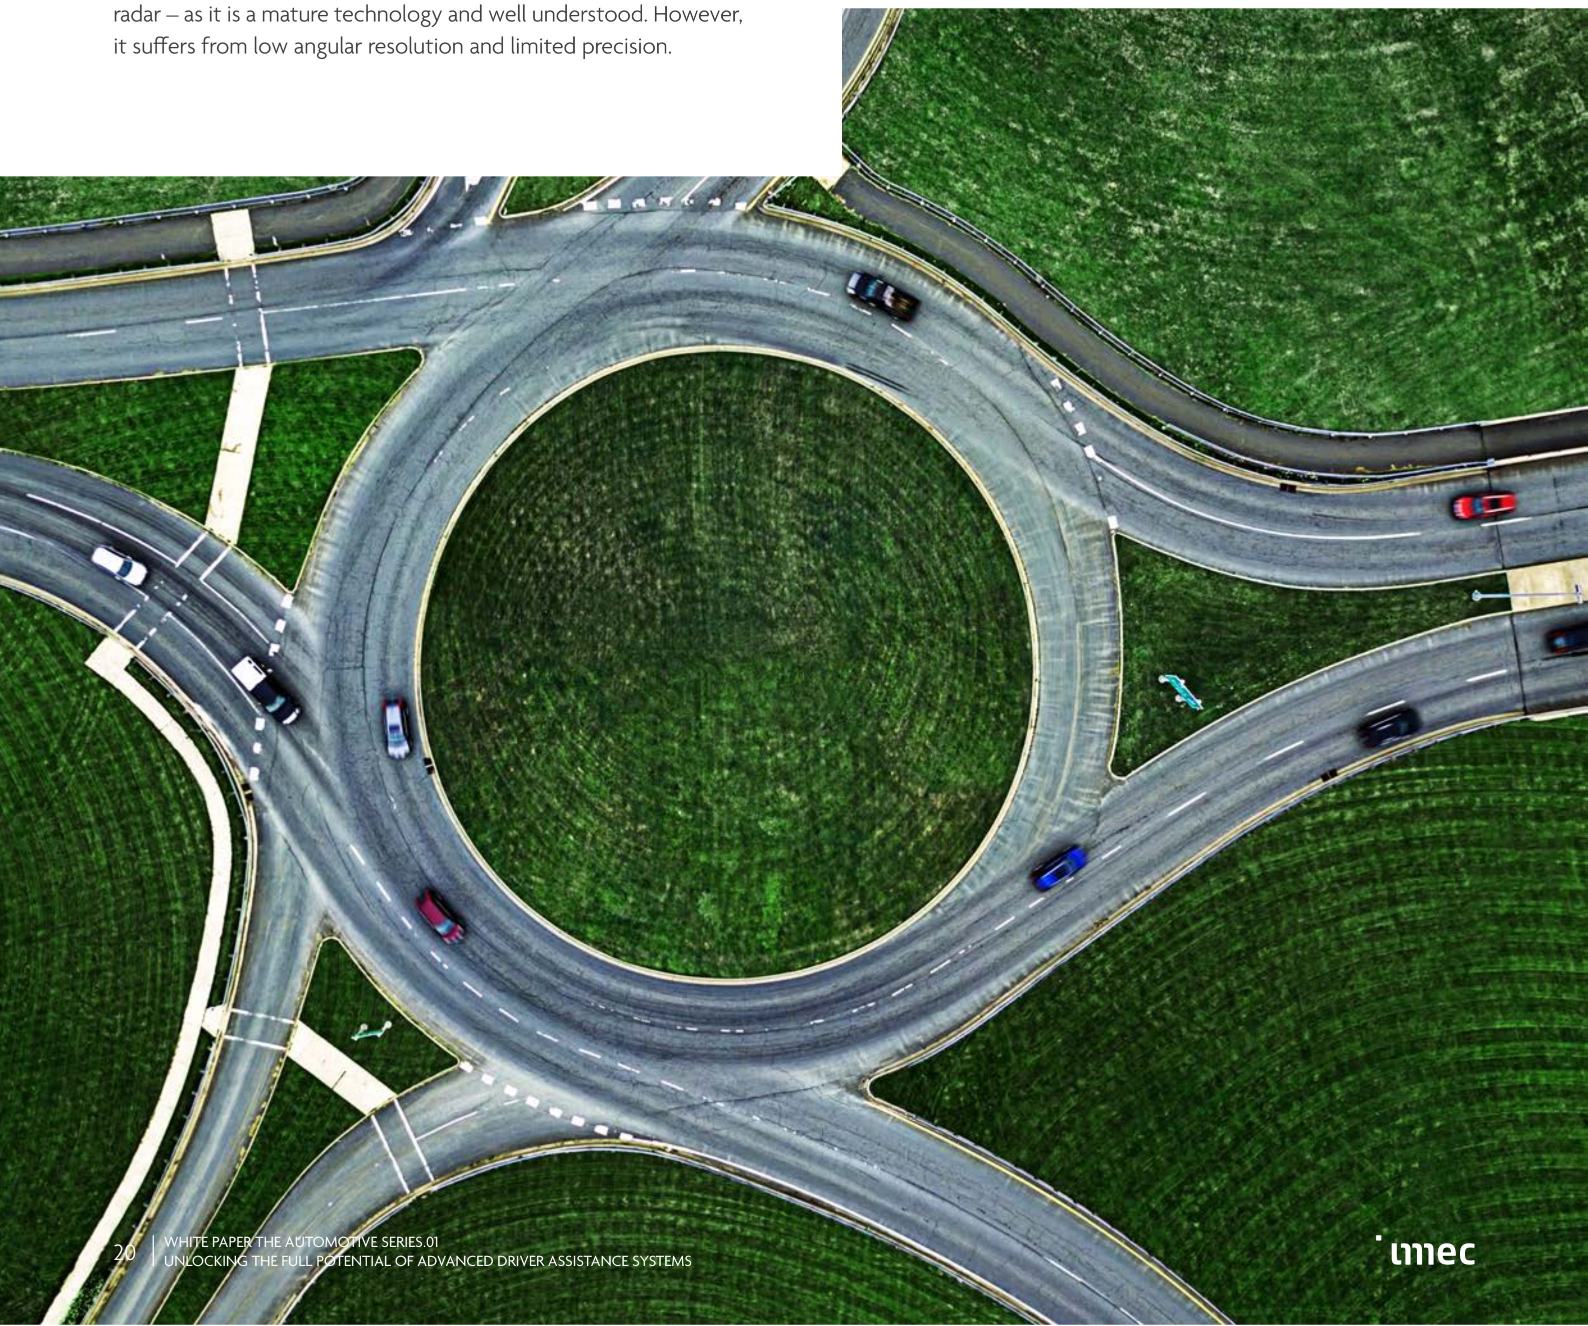

Launching a call for integrated solutions

Current sensor fusion implementations greatly suffer from a lack of integration. Each tier-1 supplier, for instance, offers its own traffic sign recognition package featuring a combination of dedicated software and hardware.

This forces automotive OEMs to buy a camera box from one supplier and a radar box from another supplier; boxes that are typically not talking to one another. This makes sensor fusion a difficult feat to achieve.

In such a siloed market, OEMs could benefit from launching a call for integrated solutions (such as an integrated radar-video sensor) that build on general recognition structures, in the same way our brain works. That could be an important first step to start investigating the potential of concepts such as cooperative sensor fusion.

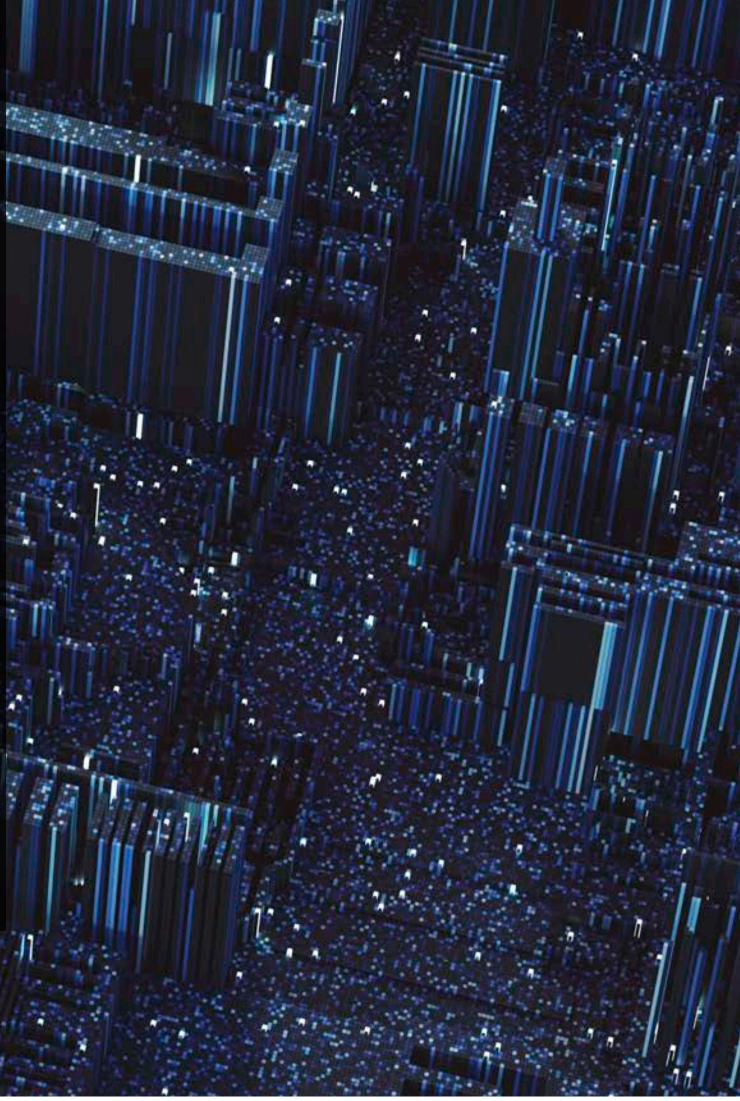

An important challenge remains: improving accuracy

Clearly, sensor fusion is a topic that is still fully developing. One of the main challenges that remains to be addressed, relates to achieving the **detection accuracy** that is required to accommodate truly autonomous vehicles.

To give a concrete example: today, the best technology option for detecting vulnerable road users (i.e. camera technology) achieves a 90% accuracy rate. In an assisted driving scenario, with a human driver who is ultimately in control, that is fine. Still, it makes truly autonomous driving in a non-mapped and challenging urban environment a distant dream.

Perceptive systems need to become much better at detecting specific contexts – so that they can anticipate situations just like humans do. A human driver, for instance, will spontaneously assess the risk of someone suddenly appearing from behind a parked vehicle – based on context and experience. An autonomous vehicle, however, does not have this context and simply resorts to rule-based decisions. In other words: to make autonomous driving a reality, perceptive systems need to have a much better 'understanding' of their surroundings. And that is where artificial intelligence (AI) comes in.

AI: mimicking human intelligence to adequately react to the intentions of other road users

The introduction of artificial intelligence (AI) is just another example of how human limitations are being offset by means of technology. In the automotive domain, this process started off with the advent of the motor car in the 1880s – substituting people's physical power with a much better alternative: the combustion engine. As a next step, thanks to AI, 'smart' cars will also optimize our cognitive (driving) skills.

As previously discussed, building smart cars that can assess – and adequately react to – the intentions of other road users will require a clever combination of increasingly powerful and versatile sensors. But that is not sufficient: they will need to be complemented with an AI engine that is capable of learning and – ultimately – reasoning.

Demystifying AI: how do smart cars learn?

AI stands or falls with its ability to learn. But how does this **(machine) learning** process actually work?

Just like humans acquiring knowledge, AI needs to be trained using lots of **training data** (e.g. thousands of pictures of road signs, pedestrians, traffic lights, etc.). In a second step, algorithms are used to derive patterns from those training data. The ultimate objective: making the AI engine recognize objects in totally new settings. When it sees a picture of a road sign it has not yet encountered, for instance, it should be able to confirm with x% probability that this is a picture of a road sign.

To accommodate this process, (lots of) **computing capacity** is needed to train larger and more complex models – using structured and unstructured data.

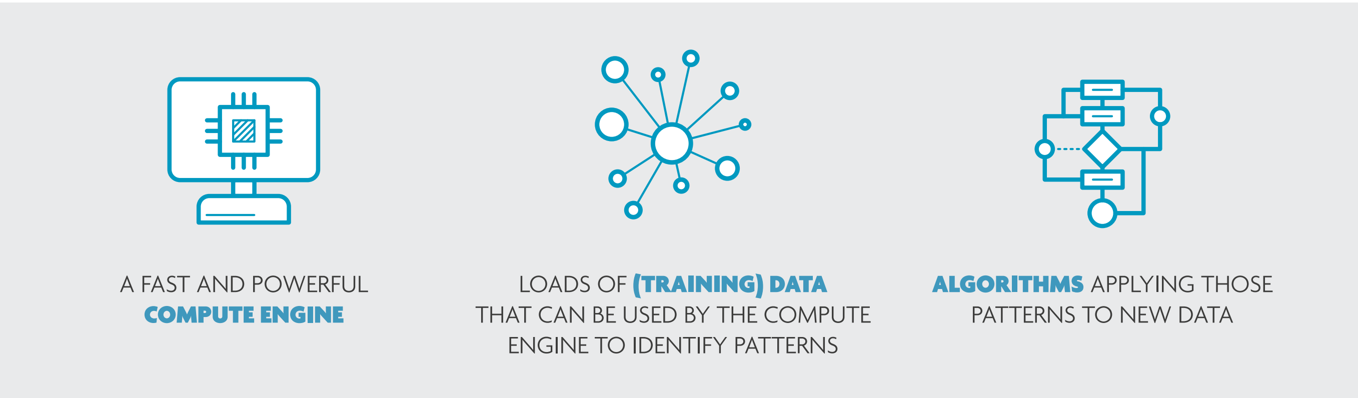

Artificial intelligence builds on the combination of computing capacity, (training) data and algorithms. Source: EDUbox AI

Machine learning comes in various flavors. One of its latest additions includes **deep learning** – a technique that enables smart systems to 'think' in multiple layers, just like humans do. Deep learning is particularly relevant when dealing with very complex situations such as driving, which involves lots of action and interaction (monitoring one's surroundings, accelerating, changing gears, braking, anticipating the intentions of other road users, etc.). In other words: deep learning is perfectly tuned to the learning needs of smart vehicles.

The concept of deep learning goes hand in hand with the use of (artificial) **neural networks:** networks that mimic the human brain. While the human brain is made up of connected networks of neurons, neural networks aim at acting like interconnected brain cells, so that they can learn and make decisions in a more humanlike manner.

Automotive Intelligence

Artificial intelligence can be defined as the ability of a technical solution to mimic human intelligence – characterized by behavior such as sensing, learning, understanding, decision-making and acting.

In an automotive context, AI is typically described as “the technology enabling a vehicle to learn and operate independently – and to ultimately make its own decisions – based on past experiences (in the form of data)”.

Tiny AI

In order to realize this vision of intelligent vehicles, multiple extreme edge devices will need to cooperate locally and share with one another partially processed information and partial decisions. They will need to involve the next hierarchical level – the edge – when acceptable (e.g. with respect to latency, privacy, etc.) or when needed (with respect to available compute power, storage or energy capabilities). In turn, edge devices will mutually share data and decisions as well, involving local or global cloud servers when acceptable or needed.

To support such vision of connectivity at the edge, intra-device and inter-device algorithms will be needed. At imec, that approach is called **tiny AI**.

From Cloud AI to (Extreme) Edge AI

Today, lots of (machine) intelligence resides in the cloud – where virtually unlimited computing power is available. That being said, the emergence of applications such as next-generation ADAS and autonomous driving will put an end to **cloud AI**’s dominant position; A concept such as cooperative sensor fusion requires AI solutions to function autonomously – in a distributed way and running at the **(extreme) edge**.

After all, the next generation(s) of smart vehicles will simply not have the luxury of being able to rely on enormous amounts of data sets to make (real-time) decisions. The computing power, nor the latency, nor the required energy will be available to do so. Moving in the direction of (extreme) edge AI, an architecture can be built that uses very small, highly energy-efficient chip sets to process limited data sets in a cleverer way.

Edge AI is characterized by local processing, using a microprocessor that has been decoupled from the sensing hardware. Extreme edge AI features local processing as well – yet makes use of a microprocessor that has been embedded on the sensing device. Source: imec.

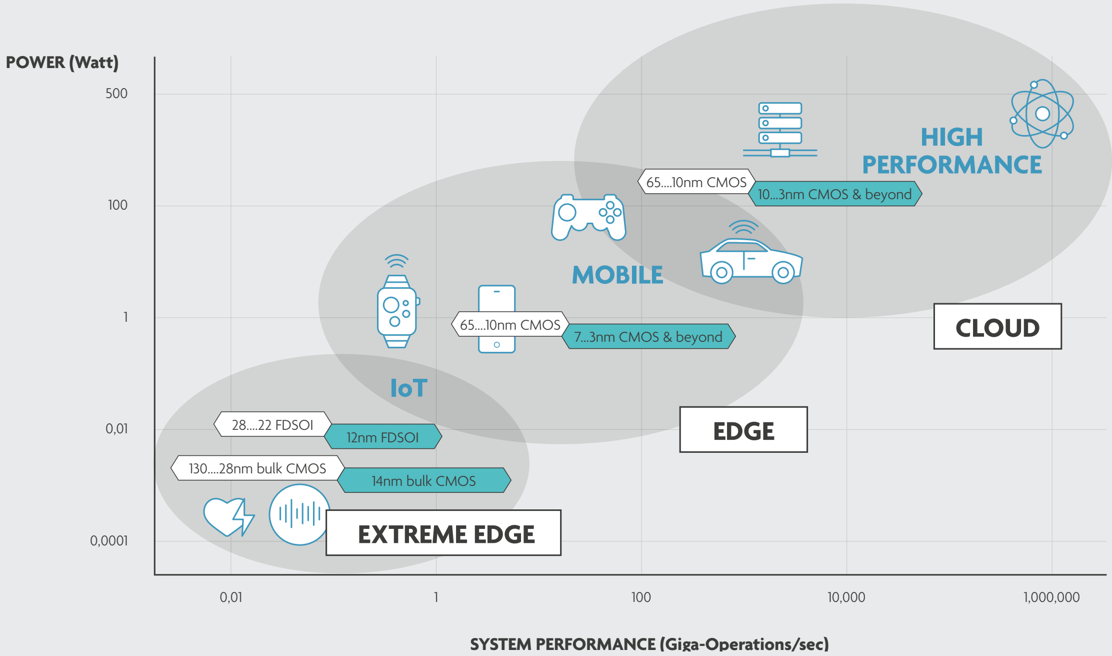

Advantages of such an approach are plenty:

- **Real-time decision-making**: (extreme) edge applications generate results that are stored locally on the device or sensor. They can process data and make independent (real-time) decisions.
- **Low-bandwidth communication**: since less data needs to be transmitted to the cloud and back, (extreme) edge solutions can do with low-bandwidth (and low-cost) communication.
- **Enhanced power-efficiency**: (extreme) edge devices are characterized by a reduced power consumption (thus improving battery life).
- **Improved data security**: by processing data locally, lots of vulnerability and privacy concerns – which are typical of IoT solutions – are solved.

From narrow AI to broad AI for better decision taking

So far, research in the AI domain has largely been focusing on narrow AI – with AI techniques being developed to solve specific problems. For instance, the first generation of smart cars that is being tested today still contains dozens of narrow AI (sub)systems: there are systems that analyze images from the on-board cameras, other systems aim at identifying road junctions and traffic signs, others are looking into navigation, etc.

To accommodate more complex technologies (such as autonomous vehicles), however, we will need to resort to broad AI approaches. Broad AI aims at developing AI solutions that can tackle problems in much the same way that humans can.

Broad AI has:

- The ability to take knowledge from one area and apply it somewhere else;
- The ability to make plans for the future based on past knowledge and experiences;
- The ability to reason;
- The ability to adapt to the environment as changes occur.

Longer-term perspective: getting smart vehicles to interact with one another

As highlighted in this section, quite some new insights have recently been introduced to make smart(er) cars a near-future reality. Yet, if there is one more challenge that merits to be called out, it is the (lack of) interaction between smart vehicles.

YouTube movies of (semi-)autonomous cars interacting with one another are pretty impressive and get huge amounts of views. However, when taking a closer look, one will see that they are typically cars from the same manufacturer and that they can only do their magic in a pre-mapped or highly predictable environment.

The underlying issue is that, today, AI agents are not yet contextually adaptive in a flexible way. What's more, they have no knowledge of the state-of-mind of other objects crossing their path: they cannot anticipate how other objects (robots/cars) – let alone vulnerable road users – will react in a given situation. In that domain, particularly, more research is required – featuring promising research tracks such as online action detection and the use of long short-term memory (LSTM) architectures.

Hybrid AI: adding reasoning to the mix using cognitive systems

The following scenarios illustrate where AI currently falls short:

- Imagine an AI-enabled car driving through a street and 'spotting' the statue of a person riding a horse. Today, that car might perform an emergency stop in the middle of the road as it thinks it is being confronted with a real horse; simply because it is not familiar with the concept of a statue.

- A similar thing might occur in a shopping street, where each mannequin in each shopwindow might be regarded as a real, living person.

- And how will AI-enabled cars react to pedestrians not using the sidewalk?

AI's current inability to deal with the above scenarios results from limitations in its training stage: if an AI system has been trained using data from a location where no statues or mannequins can be found, it has no way of knowing what a statue or a mannequin is. And the same goes when it encounters rare or unexpected traffic scenarios (such as pedestrians not using the sidewalk).

In order to take driver assistance technology to the next level and pave the way for self-driving vehicles, we will need to equip AI with the ability to 'reason' – just like humans do. And that means replacing today's (mainly data-driven) AI approaches with hybrid AI.

The topic of hybrid AI makes for a rather recent research track. It starts from the observation that data-driven AI is very fragile: it requires enormous amounts of data, it does not provide explanations or justifications, it is susceptible to biased datasets, and it can only be used for prediction (not for models that can be understood and interpreted by humans). At the other end of the spectrum, knowledge-based AI faces significant limitations as well, since manually encoding all required knowledge is simply impossible; some of it should be learned instead.

So that is where hybrid AI comes in. It integrates learning and reasoning – using a combination of data and knowledge – to get to better cognitive systems. Cognitive systems are AI systems that perform human-like tasks in natural environments. Importantly, cognitive systems are not only able to make predictions; they can explain their predictions as well. The latter is essential for AI systems that should adhere to certain standards regarding ethics, safety, or even just legal regulations.

The importance of bringing reasoning into the equation can be illustrated by means of the following example. When a person is hiking through South-Africa and comes across a lion, he or she instinctively knows it is time to run. But when that person sees the same lion in a zoo, he or she will feel totally at ease: human reasoning tells us that – this time – there is a glass wall in-between ourselves and the lion (a wall one can look through, but it is definitely there). Tomorrow's smart cars will need to possess similar AI reasoning principles in order to adequately deal with unexpected or unknown situations.

imec expertise

Researchers at imec - Ghent University (Belgium) excel at tracking pedestrians and cyclists in the most challenging scenarios (i.e. urban centers).

Researchers from the Image Processing and Interpretation (IPI) research group (an imec research group at Ghent University, Belgium) have been leading the way in the development of cooperative sensor fusion. They have defined the underlying architecture and have built an initial proof-of-concept.

Leveraging cooperative sensor fusion, the IPI research group has a proven track record of locating pedestrians and cyclists in urban centers more accurately than anyone else. While these are incremental improvements, being able to locate vulnerable road users with cm accuracy allows (self-driving) vehicles to anticipate dangerous situations much more quickly. The team's most recent test results indicate, for instance, that their first moment of detection outperforms competitive approaches with a quarter of a second.

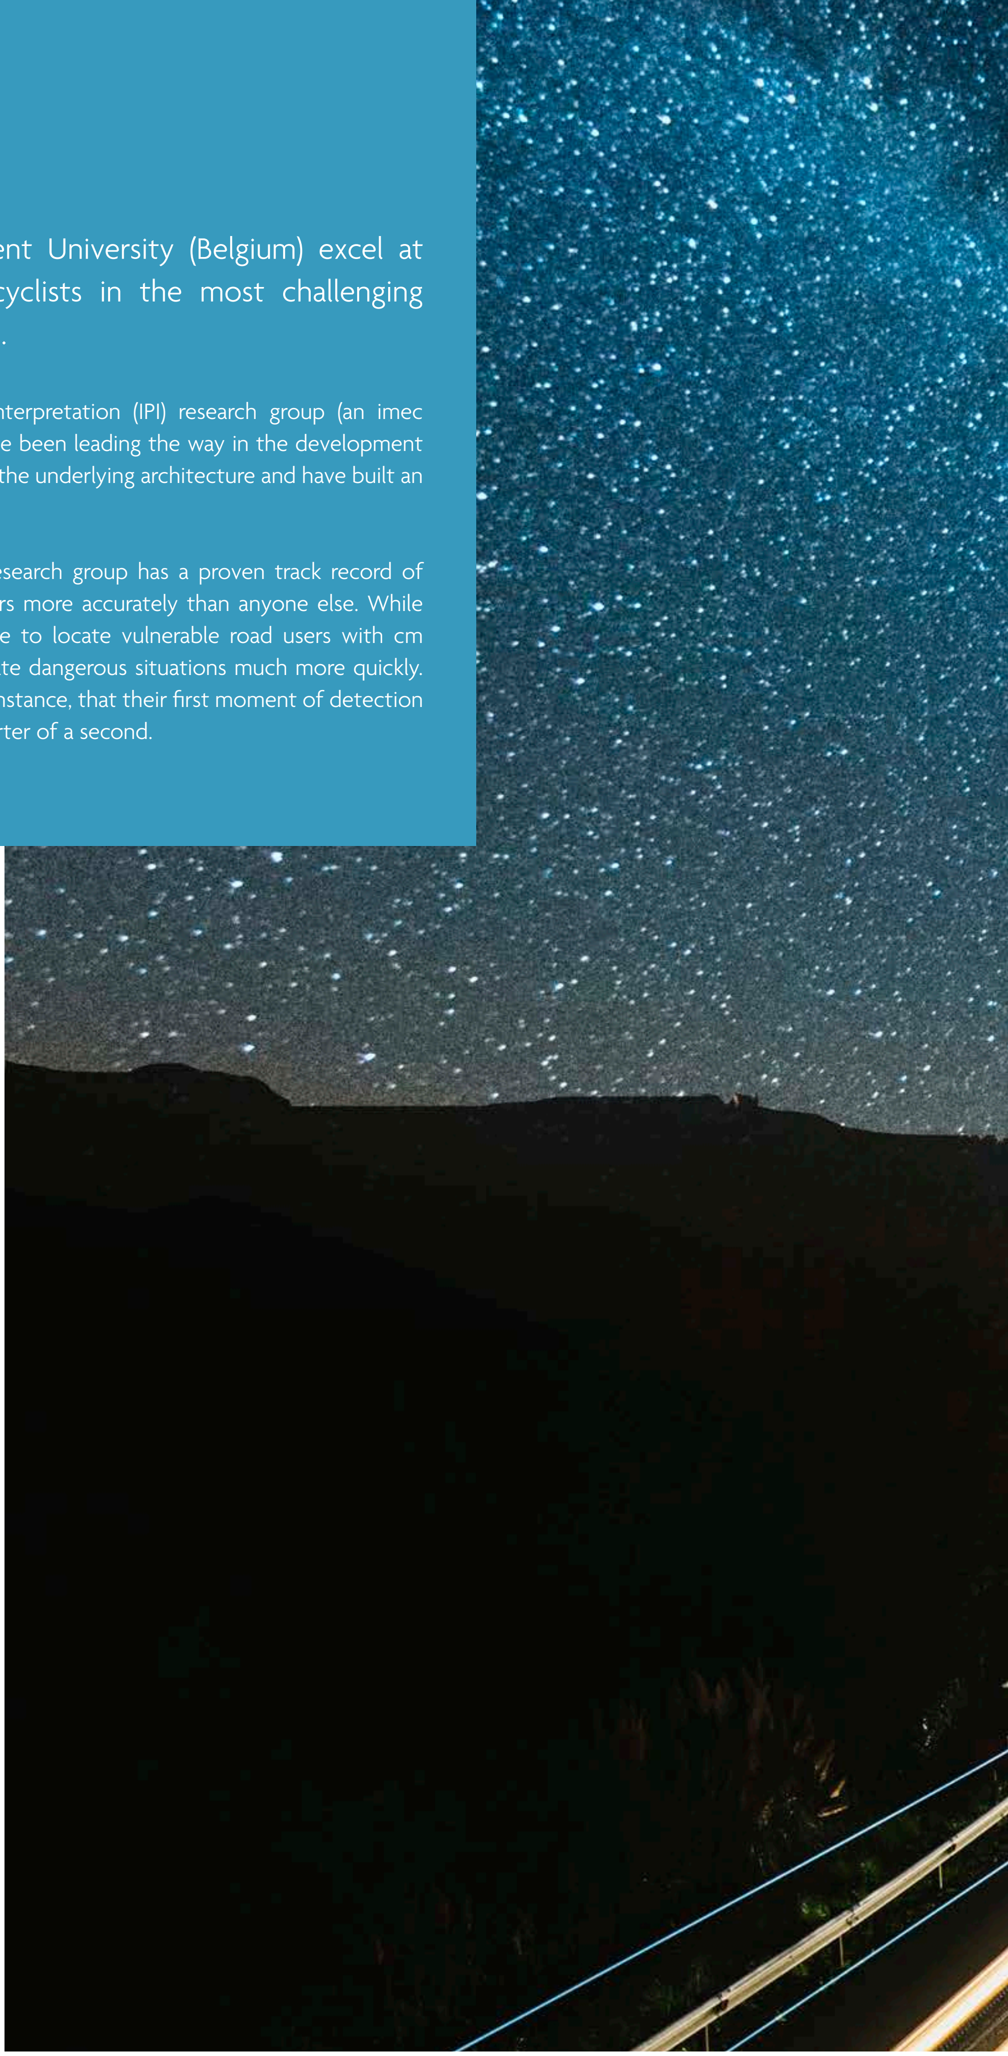

Conclusion

Current ADAS systems suffer from substantial shortcomings: their sensors lack accuracy, they work in silos instead of in concert, and they only provide input, leaving it to the driver to take action.

This white paper focused on **two domains for improvement**:

1. **anti-collision sensor technology** – Although nano- and digital technology can make camera, radar and lidar systems more efficient, probably no single technology option covers all ADAS requirements.

2. **sensor fusion and AI** – Delivering a full perceptive (3D) model and acting upon it will be the job of smart algorithms that increasingly run on energy-efficient computing devices on the (extreme) network edge.

Unlocking the full power of ADAS therefore requires you to **leverage several hardware and software capabilities**, and to look for partners within and outside the traditional automotive value chain.

As a **technology-agnostic R&D center**, imec can help you to investigate, compare and validate different approaches towards building your ADAS system. **If you're interested in working with us, feel free to contact Philip Pieters, business development director at philip.pieters@imec.be and he will soon be in touch.**

Authors

Joris Beyne, business development manager
Mieke De Ketelaere, program director
Ilja Ocket, program manager
Kathleen Philips, director
Prithi Ramakrishnan, imec North America lead for market and business development
Xavier Rottenberg, scientific director
Dimitri Van Cauwelaert, technical/business developer
Wim Van Daele, news and science editor
David Van Hamme, researcher

End notes

1 https://www.who.int/news-room/fact-sheets/detail/the-top-10-causes-of-death (May 24th, 2018)

2 https://www.euronews.com/2019/08/20/road-fatalities-which-eu-countries-are-the-most-dangerous (August 21st, 2019)

3 https://etsc.eu/more-enforcement-needed-to-reverse-decline-in-dutch-motorway-safety/ (October 15th, 2019)

4 https://www.vrt.be/vrtnws/en/2020/03/12/number-of-road-deaths-in-belgium-on-the-rise-again-after-3-year/ (March 12th, 2020)

5 https://trimis.ec.europa.eu/?q=project/vision-zero-initiative#tab-outline

6 https://www.digitaltrends.com/cars/2016-nhtsa-fatality-report/

7 Based on ADAS fitment data published by research firm SBD Automotive – www.sbdautomotive.com

8 https://fortune.com/2016/03/13/cars-parked-95-percent-of-time/