A SCALABLE APPROACH TO ADAS AND AUTONOMOUS DRIVING

ABIresearch
THE TECH INTELLIGENCE EXPERTS™

James Hodgson, Research Director

CONTENTS

INTRODUCTION

The automotive industry, in conjunction with major suppliers in the field of AI, high performance compute, mapping and location intelligence is investing heavily in the development and deployment of assisted and autonomous driving. A broad spectrum of applications will either support drivers to drive more safely, take on certain tasks on the driver's behalf, or ultimately entirely replace drivers through the automation of the entire driving process.

Each autonomous application combines a certain set of features with a degree of driver disengagement – a combination of longitudinal and lateral automation under a given degree of driver responsibility. This responsibility could require the driver to constantly supervise, or it could allow for the driver to disengage manually, visually, or cognitively, depending on the architecture of the system.

At one extreme are active safety systems, designed to support drivers only in exceptional circumstances, either through warnings or collision-mitigating maneuvers, and keep the driver in full control, while at the other end of the spectrum are driverless vehicles, which will completely replace drivers in all driving tasks, and enable disengagement to the point that no human operator is required.

Figure 1: ***Autonomous Applications – Feature / Engagement Combinations in the SAE J3016 Standard***

(Source: ABI Research)

MORE AUTONOMOUS FEATURES

SAE L2+

Features: High range of driving maneuvers

Supervision: Full and constant

SAE L2

Features: Limited lane keeping and speed control

Supervision: Full and constant

SAE L3

Features: Limited lane keeping and speed control

Supervision: Eyes off, brain on

SAE L4/5

Features: High range of driving maneuvers

Supervision: Eyes off, brain off, or no driver

TODAY **2023+** **2026+**

MORE UNSUPERVISED OPERATION

From the consumer perspective, these feature / supervision combinations seem radically different, in terms of their value, cost and the overall impact on their personal mobility experience. However, from an architecture perspective, these applications share a common set of enabling technologies, with additional components added to enable more features and greater redundancy in the more comprehensive autonomous vehicle implementations.

Therefore, the automotive industry should adopt a scalable approach to their active safety, semi-autonomous and fully driverless applications. Maximizing re-use of components between different feature/disengagement combinations will yield the following benefits for the market:

- **Cost Reduction:** Instead of building a new platform for each autonomous application, having a common set of enabling technologies powering active safety, supervised autonomous driving and unsupervised autonomous driving allows for the higher volumes of the short-term opportunities (ADAS and Level 2+) to make longer-term applications (highly automated driving and driverless) more feasible.

- **Scale-Up and Experience:** Similarly, deploying common components into ADAS, supervised autonomous driving and unsupervised autonomous driving allows for the experience gained in the shorter-term, higher-volume opportunities to form the basis of a future successful rollout of unsupervised autonomous driving.

ADAS AND ACTIVE SAFETY

The largest opportunity for automotive in the short term, advanced driver assistance systems (ADAS), support drivers by identifying obstacles and hazardous situations, and supporting the driver to avoid a collision. In most cases, this involves the use of active sensors to detect and classify road users, determining if a collision is likely to occur, and taking proactive action such as automatic braking or automatic steering.

THE CORE ROLE OF SAFETY RATINGS AGENCIES

A Level Playing Field

In general, automakers do not tend to differentiate by safety. This reflects prevailing consumer attitudes, which hold that every vehicle should be as safe as the technological state-of-the-art allows. As a result, the safety specification of a typical vehicle is driven by regulation and quasi-regulation. For example, the European GSR 2 regulation requires all new models from July 2022 and all new vehicles from 2024 to be equipped with a number of ADAS features, including autonomous emergency braking (AEB), intelligent speed assistance (ISA) and driver drowsiness and attention warning (DDAW).

Outside of formal regulation, the critical work of safety ratings agencies has a "quasi-regulatory" effect on automotive safety, and has proven highly effective in driving adoption of active-safety over the past 10 years. These agencies perform two vital functions:

- **Standardized Testing and Scoring:** By subjecting every OEM's ADAS to a common set of testing conditions, consumers are guaranteed a certain degree of performance, regardless of automaker or branding.

- **Safety Communication:** By reflecting the fitment of certain ADAS with 5-star safety ratings or top safety picks, safety ratings agencies efficiently communicate the value of active safety to consumers, driving adoption.

Roadmaps and Application Trends

The specific roadmaps for introducing new testing protocols and reflecting performance in these tests in safety scores varies across the different regional NCAPs / IIHS. However, in general, the following trends in safety testing are shaping the ADAS / active safety market:

- **Detecting More Road Users:** While initial ADAS systems were highly effective at detecting other vehicles, ADAS testing roadmaps have evolved to better test the ability of ADAS systems to detect other road users, such as bicycles and pedestrians.

- **More Realistic Testing Scenarios:** Initial testing protocols tested vehicles in daylight conditions with good visibility. Increasingly, testing at night will increase the effectiveness of ADAS in real-world conditions.

- **Internal Cabin Monitoring:** As well as using active sensors to identify hazards outside of the vehicle, roadmaps from NCAPs and the IIHS are increasingly rewarding systems that use active sensors to detect and mitigate hazards unfolding within the vehicle cabin. Typical applications include fatigue and driver distraction detection, and the detection of children accidentally left behind in locked vehicles.

- **Over-the-Horizon and Non-Line-of-Sight (NLOS):** While most ADAS rely on active sensors for line-of-sight (LOS) detection of hazards, future safety ratings agency roadmaps increasingly reward the use of non-line-of-sight (NLOS) technologies, such as cellular connectivity and broadcast V2X to identify hazards over the horizon and around corners.

TECHNOLOGY IMPLICATIONS

While safety ratings agencies set out testing protocols to reward ADAS that can successfully detect and mitigate certain accident scenarios, they do so with a technology-neutral approach, allowing for automakers to choose from a range of competitive alternatives when sourcing their ADAS solutions. However, the automotive industry has developed technology solutions that deliver the required performance in a scalable and cost-effective way.

Camera-Radar Sensor Fusions and Stereo Vision

In order to robustly detect, identify and range a wide array of potential obstacles, automakers tend to adopt one of two approaches to perception. The most widely adopted approach is camera / radar sensor fusion, with the contrasting strengths of the two sensor modalities delivering a more robust perception. CMOS camera sensors are low cost, deliver rich semantic insight to aid in object classification, and through the use of optical flow analysis and structure from motion, can also deliver depth input through non-probabilistic methods. However, camera sensors can struggle in extreme lighting and weather conditions. In contrast, radar sensors have relatively poor resolution, but continue to perform in the same circumstances that compromise camera performance. Radar sensors also deliver useful inputs such as range and relative velocity.

Figure 2:   Camera-Radar Sensor Fusion

(Source: ABI Research)
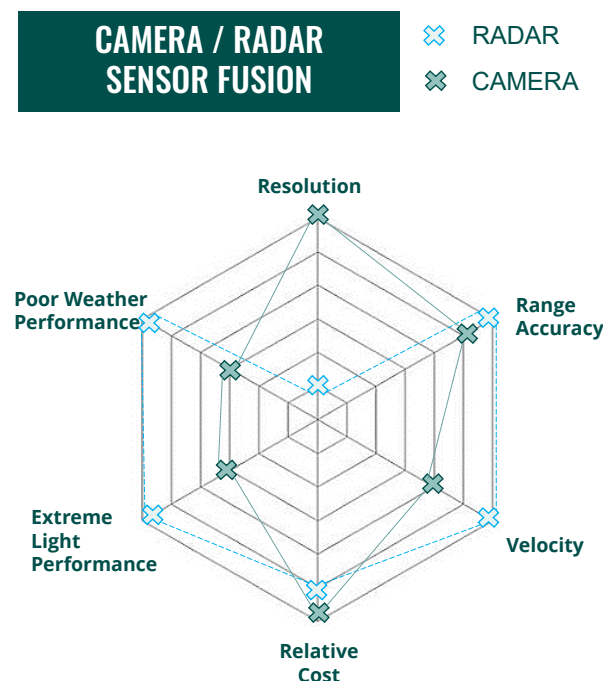
An alternative approach that has proven effective in delivering high-performance ADAS is stereovision, which leverages stereopsis to build a 3D model of the environment in the same way that humans use binocular vision to determine depth and range. This approach enables a highly compute-efficient approach to delivering the semantic, range and velocity insights required to enable robust obstacle detection and collision avoidance.

Figure 3: Stereoscopic ADAS Performance, Subaru Eyesight

(Source: NCAP)

Efficient, Low Footprint Compute

As automakers have limited avenues to differentiate through active safety, cost efficiency is key. To that end, automakers typically opt for standalone / dedicated compute that delivers the necessary applications processing, image signal processing, security and automotive-grade design. In particular, the ability to limit the functional safety burden that active safety requires to a dedicated SoC has helped automakers to limit the costs associated with active safety.

SUPERVISED AUTONOMOUS DRIVING

Level 2+

SAE J3016 defines Level 2 autonomous driving as "... steering and brake / acceleration support to the driver..." which they must "constantly supervise". For most Level 2 vehicles, this combination of longitudinal and lateral assistance comes in the form of two distinct ADAS systems operating simultaneously, such as an adaptive cruise control providing longitudinal assistance by keeping a safe distance from the vehicle in front and a separate lane keeping assistance system providing lateral assistance by keeping the vehicle in the lane of travel.

Increasingly, OEMs are bringing to market Level 2 systems that provide coordinated longitudinal and lateral assistance, leveraging compute platforms and software originally conceived for Level 3 and Level 4 systems to deliver a far more compelling consumer experience.

In practice, a wide range of features can be defined as "steering and brake / acceleration support to the driver". For example:

- Automatic lane change maneuvers

- Automatically taking highway exits or off-ramps

- Automatic overtake maneuvers to maintain target speeds on highways, with or without any human input

- Handsfree urbanized driving, navigating the complex and multi-agent environment of a city center

If these functions are executed under the constant supervision of a human driver, these deployments can be designated as Level 2 under the SAE J3016 definitions, despite offering a much richer experience than typical ACC and LKA combinations.

Semi-autonomous applications that combine multiple automated features with ongoing driver supervision have come to be known in the automotive industry as "Level 2+" or "Level 2.5" applications – a definition that is not found in SAE J3016, but which conveys the emphasis on driver responsibility, while still delivering far more features than the conventional Level 2 systems that have been deployed to date.

Advantages for Automakers and Suppliers

For automakers and their suppliers, there are a number of advantages to the Level 2+ trend:

- **Satisfying today's regulatory environment:** Developing a regulatory framework for unsupervised autonomous driving has been a lengthy process across the world. While some regions are now starting to set out the processes for how unsupervised vehicles will be type approved and insured, the overall outlook is still complex, with automakers set to navigate a patchwork of different requirements across different markets for the foreseeable future. In contrast, automakers can deploy Level 2+ systems globally within the context of the existing regulatory environment.

- **Minimize Brand Risk:** Even where formal regulation to assign liability in unsupervised autonomous driving is in effect, automakers are still wary of having their brand associated with any accident that might occur while the vehicle is driving without supervision. Even

if an investigation ultimately concludes that another road user was at fault, the impact on brand integrity that high-profile accidents can cause in the meantime produces a degree of risk aversion in many OEMs. Keeping human drivers at the center of the driving process, while supporting them with Level 2+ functions, helps risk-averse OEMs dip their toe in the autonomous future.

- **Minimize Bill of Materials (BOM):** Much of the BOM cost associated with higher levels of automation (i.e. unsupervised) comes from the introduction of redundant sensors and processing to replace the role of the absent driver. Examples include LiDAR, imaging / HD radar and duplicate AV SoCs. Keeping drivers in the loop means that advanced technologies associated with redundancy are unnecessary (as the human driver provides this redundancy). This makes Level 2+ deployments much more cost effective in the short to medium term. Over time, some additional sensor technologies, particularly imaging / HD radar, are expected to be incorporated into Level 2+ systems in order to further improve their safety. Unlike active safety systems, which tend to be shaped by safety ratings agency testing protocols, the success of Level 2+ systems will depend on their real-world performance, creating an opportunity for imaging / HD radar in the future.

Overall, a Level 2+ strategy takes advantage of the relatively lower costs, lower risk and broader regulatory accommodation of supervised automation to kick-start the autonomous vehicle revolution. OEMs, their suppliers and partners have invested billions of dollars in autonomous vehicle technology development, with almost no volume to show for it in the realm of unsupervised and driverless vehicles. Level 2+ will deliver short- to mid-term shipment volumes for enabling technologies, which will in turn reduce the cost of Level 3 and Level 4 deployments that leverage the same enabling technologies.

Advantages for Consumers

For consumers, the Level 2+ phenomenon represents a significant upgrade on legacy Level 2 systems, which frequently were hampered by highly constrained operational design domains (ODDs), and stopped functioning outside of certain speed ranges, or when road curvatures exceeded certain angles. By leveraging software originally intended to be deployed in highly autonomous vehicles, OEMs can deliver a far more compelling form of supervised automation than consumers have accessed before.

TECHNOLOGY IMPLICATIONS – 360 DEGREE PERCEPTION AND HIGH-PERFORMANCE COMPUTE

360-Degree Perception with Camera / Radar Sensor Fusion

Active safety ADAS offers support in exceptional circumstances in a single dimension. For example, AEB offers longitudinal support, while BSD and LKA offer lateral support. In contrast, Level 2+ systems are "pilot" systems intended to deliver a wide range of driving assistance, both longitudinal and lateral, on an ongoing basis.

In order to deliver this combined longitudinal and lateral assistance, Level 2+ systems require 360-degree perception covering the whole of the road environment. Typically, camera sensors are supported by radar sensors for the same reasons as in active safety implementations, with the contrasting strengths of the two sensors complementing each other. However, as the driver is there to constantly supervise the operation of the system, there is no need for additional redundant sensors such as LiDAR or imaging / HD radar.

Table 1:   Autonomous Vehicle Sensor Suites, SAE Levels 2 to 5

Manufacturer	Model	System	SAE Level	Camera	Long-Range Radar	Mid-Range Radar	LiDAR	US
Mercedes-Benz (2016)	E Class	DRIVE PILOT	2	2 (Stereoscopic)	1	4	0	4
Audi (2018)	A8	AI Traffic Pilot	3	5 (1 long-range, 4 short-range)	1	4	1 (Valeo SCALA)	4
Nissan (2016)	Rogue, Serena	ProPILOT	2	1	0	0	0	0
Nissan (2019)	Skyline, Q50	ProPILOT 2.0	2+	7	5	0	0	12
Tesla (2019)	Model 3	Autopilot 3.0	2+	8	1	0	0	12
Mobileye (2021)	Geely Models (Geely, Polestar, Smart)	SuperVision	2+	11	1	0	0	0
Mercedes-Benz (2022)	S Class	DRIVE PILOT	3	2 (Stereoscopic)	1	4	1	4
GM (2021)	Cruise	Cruise	5	16	11 (5 long-range, 4 mid-range/wide-angle)	10 (All short-range)	5	0
Mobileye / Luminar (2022 - 2025)		Full-stack	4/5	11	6		3 (SWIR ToF)	0
Zoox		Full-stack	5	14	10		8	0
Mobileye (2025-)		Full-stack	4/5	11	6 HD Radar		1 FMCW	0
NIO (2023)	ET7/ET5	NAD	L2+ NOA	11	1	0	1	12
NIO (2023)	EC6/ES8	NAD	L2+ NOA	11	1	0	1	12
Xpeng (2023)	G6/P7i/G9	XNGP	L2+ NOA	11	0	2	1	12
LiAuto (2023)	L9/L8/L7 Max	AD Max	L2+ NOA	11	1	0	1	12
Huawei AITO (2023)	M7/M5	Huawei ADS2.0	L2+ NOA	11	1	0	1	12
SAIC IM Motors (2023)	LS6	IM AD2.0	L2+ NOA	11	0	1	1	12

High Performance Compute

Active safety systems tend to rely on lean, energy-efficient compute, to deliver the functionality necessary to secure top scoring safety ratings or to comply with local ADAS regulations within the smallest possible cost and compute-power envelope. In contrast, Level 2+ systems not only require more compute performance to power a wide range of driving functions, but also tend to feature additional processing capacity, beyond what is necessary for functionality at point of sale in order to enable the delivery of new Level 2+ applications post sale.

This is particularly relevant due to the prevalence for subscription business models for Level 2+ systems, in which consumers pay a recurring monthly or annual fee to maintain access to the Level 2+ capabilities. In this business model, minimizing churn and maintaining value for the consumer is essential, and the regular delivery of new features is the best strategy to keeping consumers engaged and paying for the Level 2+ subscription. This strategy in turn requires that

OEMs seed the market with "headroom" in the form excess compute capacity in the ADAS / AV processing domain, creating a new opportunity for post-sales monetization.

Typically, the software content of a Level 2+ system encompasses more AI and neural networks, requiring neutral network acceleration IP to be included in most Level 2+ SoCs. This typically takes the form of GPU cores, FPGAs, or more tailored ASIC or accelerator cores, with the best SoC composition dictated by the software roadmap of the AV deployer.

These Level 2+ compute platforms are more centralized than the discrete smart sensors commonly employed for active safety, but in practice, the market is likely to still employ some discrete compute clusters alongside a high performance SoC. In some cases this will take the form of preprocessing sensor data before fusion within the central SoC, but will also apply to some applications that tend to reside on dedicated compute, such as autonomous parking.

Driver Monitoring Systems

With the driver playing a critical role as ongoing observer in all Level 2+ applications, most OEMs include a robust driver monitoring system in their Level 2+ systems to ensure that the driver continues to engage in observing the outside environment. Camera-based driver monitoring is essential to determine direction of gaze and driver pose.

HD Maps and Location Intelligence

Digital maps with rich attribution can help the vehicle to better localize within the lane of travel, even in the absence of visible lane markings, and to navigate busy intersections. AV-specific attribution such as traffic sign lane relevance and roadsign semantic content can also complement active sensors in understanding local road restrictions, even if the camera sensor is compromised by extremes of lighting or occluded signs.

Overall, compared to ADAS / active safety systems, Level 2+ applications require a higher density of camera and radar sensors to accommodate 360-degree perception, higher compute power to deliver more functions and headroom for future updates, driver monitoring to ensure the balance between autonomous system and human driver is maintained, and location intelligence / mapping with AV-specific attribution.

UNSUPERVISED AUTONOMOUS DRIVING

LEVEL 3 AND LEVEL 4

In Level 3 and Level 4 systems, human drivers are able to disengage from the driving process. In Level 3 implementations, this disengagement is only partial, with drivers able to disengage both manually and visually from the driving process, but they need to be available to resume control if the Level 3 application cannot safely navigate the road situation, with a reasonable handover period to bring the driver back into the loop. In the event that the driver fails to respond to a handover request, the vehicle must be capable of executing a minimum risk maneuver that will bring the vehicle to a safe state (e.g. stationary at the side of the road).

Conversely, Level 4 systems are able to fulfill all driving tasks within the operational design domain without any input, oversight or backup from a human driver.

While these two implementations (Level 3 and Level 4) will be very different from the perspective of consumer value, from an enabling technology perspective they both have the same additional requirement on top of Level 2+ architectures – redundancy. As human drivers disengage, their supervisory role must be replaced with technology-based alternatives, both in perception and processing.

Robotaxis and Driverless Commercial Vehicles

While Level 4 vehicles are equipped with a much higher density of AV technologies, their short-to medium-term volumes are expected to be in a different magnitude than the Level 2+ and Level 3 applications deployed in the passenger vehicle market. Indeed, fully driverless vehicles will be deployed in a robotaxi context, with fleet operators employing as few vehicles as possible to fulfill the mobility demand.

Driverless vehicle deployments in support of people transit on public roads are still highly limited, and expected to remain so until legislation evolves to accommodate the introduction of driverless vehicles at scale. However, commercial vehicle use-cases are expected to benefit from driverless operation in the short-term owing to a number of factors:

- **Clear Business Case:** Whether to compensate for a lack of available drivers, or to reduce losses attributable to driver error, it is easier to build a business case in a fleet context.

- **Off Public Highway:** Many commercial vehicle use cases for fully driverless operation can be realized in off-highway contexts, representing a much lower-hanging fruit opportunity than the complex urban environment, which features multiple vulnerable and sometimes unpredictable agents. Marshalling, moving trailers and repositioning assets in yards and other controlled private environments can deliver enterprise value in a comparatively easy environment.

Chart 1:   New Vehicle Shipments by SAE Level
World Markets: 2018 to 2032

(Source: ABI Research)

TECHNOLOGY IMPLICATIONS—REDUNDANCY IN PERCEPTION, PROCESSING, AND SOFTWARE

In order to replace human drivers with respect to supervision, autonomous vehicle perception in Level 3 and Level 4 deployments requires redundancy in ranging, velocity and semantic insight. Two particular sensor modalities are set to complement camera-based machine vision – HD / imaging radar and LiDAR.

HD/Imaging Radar

Radar has been adopted in the automotive space, powering ADAS perception for many years. While the modality is valued for its orthogonality to camera sensors, and its robust performance in poor lighting and weather, conventional configurations suffer from a number of weaknesses.

- **Better Elevation / Vertical Resolution:** Poor resolution along the vertical axis makes it difficult to determine the height of obstacles, and therefore to distinguish stationary vehicles from stationary street furniture, such as overhanging gantries or road signage.

- **Better Azimuth / Horizontal Resolution:** Poor angular resolution makes it difficult to distinguish different agents on the road, such as vehicles and pedestrians. A pedestrian crossing perpendicular to the path of the vehicle will often fail to be detected or distinguished from other agents in the scene.

- **Minimizing Sidelobes / False Positives:** The presence of grating lobes alongside the main lobe of the radar sensor's radiation pattern results in multiple false positives. This is, once again, caused by the limited number of channels in the current range of automotive radar chipsets. In practice it can be difficult to distinguish between the genuine reflection of a less reflective object (such as a pedestrian) and the sidelobe of an adjacent, more reflective object (such as a parked truck).

The solution to all of the weaknesses discussed above is adding more virtual channels to the transceiver array. This is achieved by increasing the number of physical Radio Frequency (RF) channels (Transmit (Tx) and Receive (Rx)), and, therefore, massively increasing the number of virtual channels through typical Multiple Input, Multiple Output (MIMO) techniques.

These new-generation radar transceivers produce far higher data volumes than the status quo, requiring a new generation of radar processing capable of handling the new data volumes, as well as enabling the new applications made possible by the significant increase in resolution, such as mapping.

CASE STUDY: Arbe Robotics

Founded in 2015, Arbe specializes in the development of transceivers, processors and platforms for high definition radar. In order to boost the number of Tx and Rx channels in its next-generation radar technologies, Arbe has leveraged Global Foundries' 22 FDX / 22nm FDSOI CMOS process, delivering 24 output and 12 input channels in its RFIC chipset, while retaining a low cost per channel. These physical channels can then enable over 2000 virtual channels, delivering angular resolution of 1° azimuth (horizontal) and 2° elevation (vertical) and a range of 300m. In order to process this large volume of data, Arbe has developed a radar processing unit (RPU) to integrate embedded signal processing algorithms within its baseband. This enables the processing chip to accommodate 48Rx and Tx channels while minimizing power consumption.

In September 2020, Continental, a leading Tier 1 supplier of automotive radar modules, announced the ARS540, a next-generation 4D radar device designed to address the legacy shortcomings of radar sensors. The ARS540 features 12 TX and 16 RX channels, 1.75 times the number of antenna channels on the previous flagship radar module, the ARS430. The ARS540 is therefore able to deliver 192 virtual antenna channels, driving a significant improvement in azimuth and elevation resolution. Practically, this system has accurate height estimation to detect overhanging signs and tunnel rooves, sufficient angular resolution to distinguish road agents in complex urban environments (e.g. a motorcycle next to larger parked vehicle) and detection of non-overridable ground obstacles.

In order to deliver compute performance required to process an eight-fold increase in the number of virtual channels, Continental has opted to use AMDs Zynq™ UltraScale+ adaptive SoC, leveraging the 16nm technology to create a radar-based point cloud that can deliver on the 4 "dimensions" of range, velocity, azimuth and elevation. In comparison to the previous generation Zynq 7000 adaptive SoC the Zynq UltraScale+ SoC delivers 20 times the raw data processing power and 10 times the object-tracking capability.

LiDAR

LiDAR has been a mainstay of autonomous vehicle prototype perception suites since the days of the DARPA Grand Challenges, and continues to be an important modality in unsupervised autonomous driving deployments. Continuous development over the past 15 years has yielded numerous iterations of the core principle of LiDAR technology, which uses infrared light to range objects with sufficient resolution for object classification. Some of the different technology options for LiDAR solutions include:

- **Pulse / Time-of-Flight (ToF) vs. Frequency Modulated Continuous Wave (FMCW):** In Pulse/ToF configurations, separate pulses of light are emitted at a fixed frequency, with the response time of reflections measured to determine the range of different points. Algorithms are then leveraged to identify and track objects, classifying and measuring velocity, etc. A more recent alternative is FMCW LiDAR, which operates according to the same principle as conventional automotive radar. Rather than sending brief laser pulses at a fixed frequency, this next generation of LiDAR technology modulates the frequency of the continuously emitted laser light. Comparing the frequencies of the emitted wave and reflections with a reference oscillator enables range detection and velocity detection.

- **NIR vs. SWIR:** Most LiDAR solutions leverage NIR emitters (800nm to 905nm), an approach which enables the use of common laser technologies, which are already adopted at scale and benefit from economies of scale. However, this wavelength can permeate the human retina, resulting in strict regulation limiting the output power of lasers in the NIR wavelength, which in turn limits the range of NIR LiDAR. SWIR wavelengths above 1400nm do not permeate through the human cornea, allowing for much higher output levels, and thus longer ranges.

- **Scanning:** Original LiDAR leveraged large rotating mirrors to scan the emitted light across the scene – an approach that results in a high degree of wear and tear. Alternatives include digital MEMS mirrors and optical phased array beam-steering techniques.

Overall, LiDAR delivers a blend of range detection, velocity measurement and relatively robust performance in poor lighting and weather conditions, while also delivering the necessary resolution for object classification, semantic segmentation and free space detection. Therefore, LiDAR is regarded as an essential tertiary sensor modality or "third opinion" in unsupervised automation.

Finally, an interesting recent dynamic, especially in competitive EV markets like Greater China, is the increasing use of LiDAR to complement sensor sets even in Level 2+ platforms as a way to differentiate their cars from popular EV makers with "camera-only" strategies.

Figure 4:   Camera, HD Radar, and LiDAR Sensor Fusion

(Source: ABI Research)
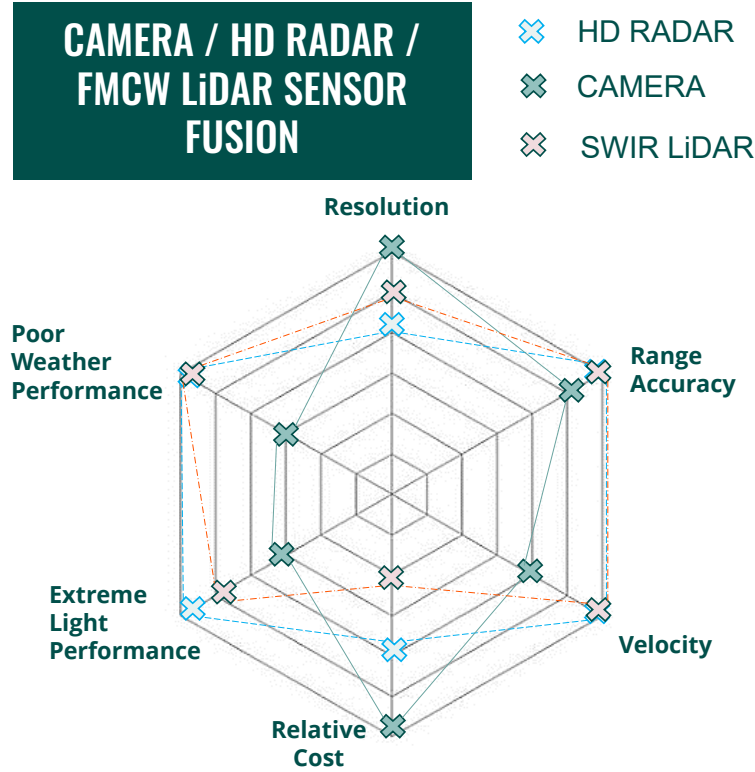
High Performance Compute

Unsupervised autonomous driving requires a higher level of compute performance over Level 2+. Firstly, HD / imaging radar and LiDAR are both high resolution sensors with compute-intensive algorithms for ranging, velocity, object classification, semantic segmentation and mapping. Therefore, the addition of multiple HD / imaging radar and LiDAR sensors requires greater compute performance.

In addition, duplication of compute resources in order to guarantee performance in the event of an error is important to ensure graceful failure, or to return the vehicle to a safe state through the execution of a minimum risk maneuver.

The general trajectory will therefore be compute that is more performant, more heterogeneous and more centralized than in supervised autonomous vehicles. Once again, a balance will be struck between the use of smart sensors with data preprocessing, and the ingestion of raw digital data within a centralized module.

AMD (XA) Zynq™ UltraScale+™ MPSoC 7EV and 11EG devices deliver highly programmable capacity, performance, and I/O capabilities enabling high-speed data aggregation, preprocessing, and distribution, as well as compute acceleration for L2+ to L4 levels of automation. These devices offer over 650,000 programmable logic cells and nearly 3,000 DSP slices; a 2.5 times increase versus the previous device generation and a five times increase in system-level performance per watt over Zynq-7000 adaptive SoCs. The XA 7EV is equipped with a video codec unit for H.264/h.265 encode and decode, while the XA 11EG includes 32 12.5Gb/s transceivers. They are supported by the company's new unified software platform, Vitis™, which enables developers to take advantage of the hardware adaptability. As such, the AMD (XA) Zynq UltraScale+ MPSoC has been used in a variety of ADAS applications in addition to the Continental ARS540 4D Imaging Radar platform. Some recent announcements include Subaru (Forward Camera), Denso (LiDAR) and Aisin (Surround View, Automated Park Assist). Baidu Apollo runs on XA Zynq UltraScale+ MPSoC (XAZU5EV). Pony.ai also uses AMD technologies in its AV platform. With the addition of these high-performance devices, AMD delivers processing flexibility and scalability for L1 to L4 systems.

As Artificial Intelligence (AI) becomes increasingly important for next-generation advanced sensors, AMD will address the increased AI performance requirement with its next-generation Versal™ AI Edge series.

CONCLUSIONS

Figure 5:   Impact of Driver Disengagement on Sensor and Compute Requirements

(Source: ABI Research)

| | ADAS | PILOT | | |
| | ACTIVE SAFETY | SUPERVISED | UNSUPERVISED | |
			With human backup	*No human backup*
SAE Levels	L1 – L2	L2+	L3	L4
Driver Engagement	Eyes on, brain on, hands on	Eyes on, brain on, hands off	Eyes off, brain on, hands off	Eyes off, brain off, hands off
Sensors	Camera / radar sensor fusion, stereo	360º, camera, radar sensor fusion	360º, camera, radar sensor fusion, 1 LiDAR, HD Radar	360º, camera, radar sensor fusion, multi LiDAR, HD Radar
Processing	Lean ASIC	Heterogeneous HPC, NNA	Heterogeneous HPC, NNA, redundant processing	Heterogeneous HPC, NNA, redundant processing

+ Sensor Density
+Compute Headroom
+ Heterogenous Compute

+ Sensor Diversity
+Compute Headroom
+ Heterogenous Compute

+ Sensor Density and Diversity
+Compute Headroom and Redundancy

The biggest factor driving the adoption of greater compute power, heterogeneity, sensor volume and sensor diversity is driver disengagement. However, while higher levels of driver disengagement have a clear technology requirement, there is a substantial degree of component re-use between active safety, supervised automation and unsupervised automation.

SCALABILITY AND RE-USE

Cost Reduction

A scalable approach that maximizes re-use of components across applications will allow for higher volume opportunities in the short term, such as ADAS and supervised autonomous driving to pave the way for a more cost competitive deployment of unsupervised autonomous driving in the future. Similarly, the shipment of unsupervised automation in constrained operational design domains will scale up production of tertiary sensor technologies, enabling cost-effective broad-ODD deployment of unsupervised automation in the longer term.

Scaling Up Experience and Refinement

Beyond the obvious cost advantages of component re-use, scalable architectures benefit from the real-world experience of applications that have already been deployed into the market. Since the launch of BlueCruise, Ford has accrued over 100 million hands-free miles from multiple countries, giving critical insights that will refine future algorithms. Similarly, Mobileye leverages the experience of EyeQ4 systems powering ADAS to build the REM / roadbook map that plays a critical role in its Supervision and Chauffeur products.

Overall, the only feasible approach to delivering on feature-rich and unsupervised automation is to construct today's supervised autonomous applications on an architecture that has potential to scale through the addition of technologies that will replace the supervisory role that human drivers play today.

ABiresearch.
THE TECH INTELLIGENCE EXPERTS℠

ABOUT ABI RESEARCH

ABI Research is a global technology intelligence firm delivering actionable research and strategic guidance to technology leaders, innovators, and decision makers around the world. Our research focuses on the transformative technologies that are dramatically reshaping industries, economies, and workforces today.