DVN ADAS/AD Applications Newsletter – December 2025

Wayve's Embodied AI in Leonberg: What It Means for Europe's L2+/L3 Ambitions

By J. Dickmann, DVN Senior Advisor and Independent ADAS/AD Engineering Consultant

Dan McCloskey, VP Wayve for Hardware and Juergen Dickmann Senior Advisor DVN.

When a European start-up claims it can rethink automated driving with end-to-end AI and a lean sensor stack, healthy scepticism is appropriate. After spending a cold December morning in Leonberg in one of Wayve's Mustang Mach-E development vehicles, that scepticism has shifted to cautious optimism. The demo did not look or feel like a polished showcase; it felt like a serious engineering programme that has already solved a surprising number of hard problems in everyday European traffic.

The invitation came from Dan McCloskey, now VP of Hardware at Wayve and previously a familiar face from the early Waymo days. My mindset before the tour was shaped by two decades of work on ADAS and automated driving, from the Mercedes-Benz "Bertha" drive in 2013 through to radar-centric perception architectures and recent L2+ deployments in China and Europe, today. Those experiences taught how deceptively simple human driving tasks become difficult for machines, particularly in dense city traffic, and how far a carefully engineered combination of vision and radar can be pushed when the sensor fusion and validation effort is taken seriously. The last decade also taught the limits of AI hype: more data is not always better data, and black-box "end-to-end" slogans tend to raise as many questions as they answer. So I was actually really looking forward to this E2E demo.

Wayve's setup in Leonberg is deliberately modest. The test and development hub sits in an industrial area on the edge of Stuttgart, close to the A8 and a representative mix of German urban, rural, and highway traffic. The vehicles themselves are standard production cars from Ford Mustang Mach-E fitted with Wayve's compute platform and sensor suite. The engineering goal is clear: keep the vehicles as close as possible to series production hardware and integrate the AI driver with minimal intrusion, so OEM partners can see a realistic path to deployment rather than a science-project demonstrator.

Discreet entrance

Sensor mounting for fast development. Radar and cameras

The compute architecture in the Mach-E I rode in reflects that philosophy. In the trunk, under the load floor, a compact rack of NVIDIA-based GPU modules delivers the inference performance needed for real-time end-to-end driving, while still leaving useful luggage space—a sharp contrast to early AV prototypes that sacrificed the entire boot to servers and cooling. The current Gen-2 stack is already powerful enough to run high-capacity vision models on multiple camera streams and log data for offline training; the announced Gen-3 platform will move to NVIDIA's latest automotive SoCs, with headroom aimed at future Level 3 and Level 4 capabilities. From an integration standpoint, the fact that this hardware can be dropped into very different vehicles without major re-engineering is as important as the raw tera-operations figure.

a) Before lifting the boot cover, b)After lifting the boot cover

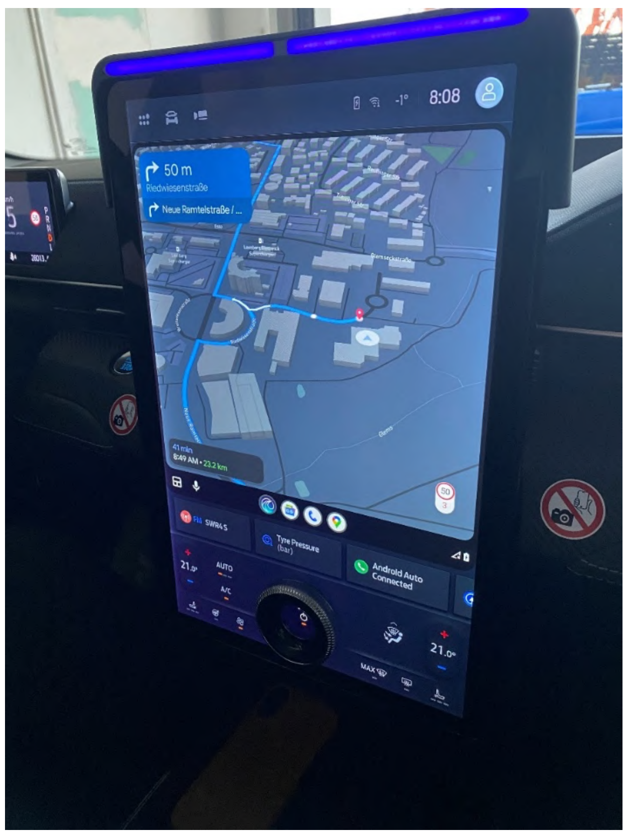

Standard screen HW used by Wayve HMI

On the sensing side, Wayve stays true to its embodied-AI thesis. The Mach-E carried a ring of automotive-grade cameras providing 360-degree coverage, including several forward-facing units with overlapping fields of view to support long-range perception and close-range redundancy. There was no LiDAR on the roof and no exotic depth sensors; radar is being phased in as the second major modality, mainly via front radar today and with the option to expand to side and rear coverage in future generations with higher autonomy. That is a lean stack by any measure, particularly for urban driving in European winter conditions, and it sets a high bar for the software.

The Route: Real-World Complexity

The route made good use of the Leonberg geography. We started before sunrise, at temperatures just below freezing, and drove into morning light without any rain or fog. When we started, we felt patches of residual ice making even walking slightly treacherous. The drive combined segments of the A8 motorway, rural roads between villages, wider city arterials, and historical town centres built for horse-drawn carriages rather than two-way car traffic. Traffic density was typical for a Monday morning around Stuttgart: a mix of commuters, cyclists, delivery vans stopping where convenient, vehicles briefly blocking or narrowing lanes for bakery or butcher shop visits. Parents with children on the school run, and pedestrians crossing at—and sometimes away from—marked crossings. Speed limits ranged from 30–50 km/h in town to 70 km/h on rural sections and 120 km/h on the motorway. In short, it was exactly the sort of messy, low-speed environment where automated systems are often exposed. The speed limits are likely to inform the current system´s capacity to function robustly without further Radar.

From the passenger seat, the first impression was the driving style. The car behaved in a way that felt recognisably "German": defensive but not timid, smooth on steering and longitudinal control, and respectful of informal right-of-way etiquette at tight pinch points. Cornering was consistently fluid, without the oscillations and late corrections that often betray over-tuned controllers or poor trajectory planning. On several occasions I mentally prepared to intervene—narrow streets flanked by parked cars, oncoming vans encroaching on the centre line, or blind corners around temporary construction—only to watch the system modulate speed early and place the vehicle confidently within its lane without drama.

T-crossing scenario worked out well

One example stands out. In a constricted old-town street with a 90-degree left turn, parked vehicles narrowed the effective width to what looked, from my perspective, barely more than a single lane. A delivery van approached from the opposite direction. The Wayve system slowed well in advance, held to the right to maximise clearance, and executed the turn without clipping kerbs or forcing the van to brake hard. That sequence required not just static path planning but an accurate sense of the Mach-E's own geometry and how much lateral margin remained as both vehicles moved. It is the kind of "simple" scenario that often reveals weaknesses in AV stacks trained primarily on wide US-style roads.

Narrow 90° left turn

Very narrow bypass in narrow road with oncoming traffic (se left corner)

Very in advance, taking car on VRUs

Pedestrian Interaction and Sensor Capability

Pedestrian interaction was equally telling. At a zebra crossing on a multi-lane urban road, with a vehicle ahead of us and a family stepping onto the crosswalk from the right, the system anticipated the need to yield before the crossing became fully visible, matching the deceleration of the lead vehicle and avoiding any abrupt last-second braking. Later, in a residential zone with children on bicycles and adults walking dogs, the car adjusted its speed to a level that felt appropriate for visibility and potential surprises rather than simply matching the posted limit. Having invested years in radar-vision fusion for VRU detection, it was interesting to see how far a predominantly camera-based stack can go when trained on rich corner-case data.

On the motorway, the behaviour was less spectacular but reassuring. Lane keeping at 120 km/h was stable, with no noticeable hunting, and merges onto and off the A8 were handled cleanly. The system left slightly more headway than an assertive human driver might choose, but within a comfort band that European customers are likely to accept. Given the current sensor configuration, which does not rely on rear radars, this conservatism is a sensible design choice. It also matches the regulatory reality: these are still supervised Level 2+ functions, with the human driver legally and practically responsible for monitoring and intervention.

How This Compares to Classical AV Architecture

How does this compare to the architectures European OEMs have been pursuing since the early 2010s? Classical AV 1.0 stacks grew out of systems engineering: modular perception pipelines, explicit rule-based planning, HD maps to anchor everything, and long validation cycles each time a sensor or vehicle platform changed. That approach produced impressive one-off demos—including the Bertha drive I participated in—but struggled with scalability. Every new geography, every new vehicle derivative, effectively reset parts of the engineering and safety case, which explains why sophisticated automated functions remain confined to a narrow subset of premium models and markets.

Wayve's embodied-AI approach attacks that rigidity. By training a single, large model to map sensor inputs directly to driving decisions, updated continuously with data from a mixed fleet of test and partner vehicles, they aim to reduce the re-engineering burden for each new platform or region. For an OEM looking to roll out L2+ functions across a wide portfolio, licensing such a foundation model could offer a faster path than building and validating a full stack in-house. At the same time, it shifts the technical challenge towards data curation, offline evaluation, and the explainability needed for European type approval. My own work on safety cases and standards like ISO 26262 and UNECE regulations makes me acutely aware that this is where end-to-end systems will face their hardest questions.

Leonberg as a Test Market

The Leonberg hub itself is part of the answer. Testing in dense, regulation-heavy German traffic, under the scrutiny of local authorities and with OEM partners nearby, forces the system to mature in an environment that leaves little room for shortcuts. The company's investment in German testing infrastructure and partnerships with established test-validation specialists signals a strategic commitment that goes beyond opportunistic marketing.

Discussions during and after the drive touched on validation, safety argumentation, and the path from supervised L2+ to conditionally automated L3 and beyond. Time was too short to explore these topics in the depth they deserve; a follow-up interview is planned to dig into questions such as how to structure safety cases for opaque models, what role simulation plays in covering rare events, and whether Europe should aim for incremental L3 steps or focus directly on limited-ODD L4 services.

Open Questions

There are, of course, open points. We did not have heavy rain, snow, or low-sun glare on this particular morning. Those conditions will be critical for any camera-centric system. The current radar configuration still looks sparse from the perspective of someone who has championed high-resolution radar as a key enabler for robust automation. And the broader policy debate—how much better than an average human driver must an AI system be before society accepts unsupervised operation—remains unresolved. These are not criticisms unique to Wayve; they apply to the entire industry.

What the Leonberg drive did demonstrate is that a European start-up (Wayve pioneered E2EAI in 2017), launching years after the US and Chinese pioneers, can field an L2+ system that handles complex real-world urban scenarios with a maturity comparable to better-funded competitors.

For engineers and decision-makers in Europe's automotive industry, Wayve's progress is a reminder that embodied AI is no longer a speculative concept. It is a working technology that now needs to prove its safety, scalability, and economic value under the toughest regulatory regimes we have.

Our DVN impression is, that business wise, Wayve also looks promising. Public internet information says, Wayve has signed definitive agreements with Nissan in December 2025 to integrate the "Wayve AI Driver" into the next-generation ProPILOT system—first series production in Japan from FY2027, with global rollout including North America planned. Already in September 2025, $500 million flowed in from Nvidia, fuelling the UK firm's expansion into Germany, Japan, and the US while betting on scalable AV2.0 technology powered by generative AI. Wayve really app.

For DVN, Wayve actually seems like Europe's last chance to play a part in the AV spectacle.

DVN note: We will follow up with a structured interview in London to probe validation, safety case methodology, scaling strategy, and future sensor suite. For comments or contributions, contact the DVN editorial team.

After this morning in Leonberg, that challenge looks demanding but far from impossible.

J. Dickmann is an independent ADAS/AD engineering consultant with 20+ years' experience in Radar and Radar perception, sensor fusion, perception architectures, sensor set-up definition and 20+ years of ADAS and AV driving.