


Frontier Challenges Bringing AI to 1B Robots

Alex Kendall, Founder/CEO, Wayve

CVPR 2026, Denver USA



Robots commercially deployed today

Application	Number	Source
Agility Digit humanoid	40	https://newmarketpitch.com/blogs/news/humanoid-robotics-digit-deployment-tracker
Waymo robotaxis	3K	https://techcrunch.com/2026/03/27/waymo-skyrocketing-ridership-in-one-chart/
Tesla FSD	1M	https://electrek.co/2026/01/28/tesla-discloses-fsd-subscriber-count-first-time-1-million/
Industrial robots	5M	https://ifr.org/ifr-press-releases/news/global-robot-demand-in-factories-doubles-over-10-years
Roomba robots	50M	https://www.sec.gov/Archives/edgar/data/1159167/000115916725000011/irbt-20241228.htm
Mobileye ADAS	200M	https://www.businesswire.com/news/home/20250130605548/en/Mobileye-Release-Fourth-Quarter-and-Full-Year-2024-Results-and-Provides-Business-Overview/

Does autonomy work?



Can it generalize?



03 mph



2X SPEED



02 mph

2X SPEED





04 mph

WAYVE

Welcome Aboard!
Our 57 Seater!

57 S

IN EMER


TURN HA
PULL DOO

2X SPEED




13 mph


2X SPEED





09 mph






WAYVE

21 mph



14 mph





10 mph

4X SPEED





21 mph

渋谷区神国町


2X SPEED





00 mph







Zero-shot driving in 500 cities

New cities over time



Zero-shot training overview

CITIES	TRAINING DATA	% OF TOTAL CITIES
219	ZERO DATA	43%
260	<10KM DATA	51%
341	<100KM DATA	67%

NUMBER OF VISITS TO EACH ROAD SEGMENT



Is generalization all you need?

- Generalization = safety
- Generalization = deployment size
- Generalization enables licensing business model
- Generalization expands impact beyond driving

Our goal is to build generalizable, cost-efficient autonomy across heterogenous platforms

Contrarian ideas



Wayve has been pioneering end-to-end AI since 2017



2018

World first demo of reinforcement learning teaching a car to drive



2018

World first demo of end-to-end learned driving on public urban roads



2021

World first demo of AI trained in London driving in 5 new, unseen UK cities



2022

World first demo of AI generalizing across different vehicle platforms



2023

First LLM and GenAI breakthroughs for autonomy: LINGO and GAIA




2025

First to show zero-shot driving in 500+ cities across Europe, North America and Asia




Industry is shifting towards AV2.0



End-to-end learning from perception to control


AV1.0 Classical



HIGHLY ENGINEERED STACKS
RULE-BASED SAFETY SYSTEMS

*Safety through rules,
maps and redundancy*


AV1.5 Hybrid



E2E MODEL CHECKED BY
RULE-BASED SAFETY SYSTEM

*Safety through
rule-based checker*

AV2.0 E2E AI



END-TO-END DRIVING
LEARNED RISK INTELLIGENCE

*Safety scales with the
intelligence of the AI*




Reinforcement Learning at Wayve today

- Millions of hours of driving experience
- On- and off-policy data flywheel with large diverse fleet
- Multi-billion parameter critics and world models
- Large multi-node RL jobs as the default
- Real distribution shifts (long tail, weather, geography, sensors)



Idea Thursday – Ride-along Monday

The Four Day Loop



From “I have an idea about distributional critics” to “I just rode a car driven by it” – same sprint.




The Journey



Multi-Agent Reinforcement Learning in Simulation for Urban Driving

Video rendered at 10x real-time



Efficient tokenized input

Distributed training for
billions of agent steps

Simple rewards:
don't crash, make progress





700x real time learning




Emergent Behaviour



Scenario: through-right-vs-giveway-straight
Reward: 0.025 World: soak-test-world
Speed: 0.00m/s Lane part: 60
Accel: 0.00m/s² Should move: no
Pitch: 0.00deg Desired speed: 8.94m/s
Ends in 14.2sec Vehicle density: 0.0
of agents: 2





Robust Autonomy Emerges from Self-Play [Cusumano-Towner et al. 2025]



Distilling to a non-privileged driving policy

INTERACTIVE SIM - WORLD: 090c14-7c91ff-0025

MODEL ID: lionfish-fierce-peach / session_2023_03_17_10_59_28_aml-v100-1nn-paul-dagger_v0_step....


IN CONTROL: S2 MODEL [TRAIN SIM]

MODEL TYPE: birdseye_costmap_mcn

TIMESTAMP: 39.40sec
VELOCITY: 5.43m/s
DISTANCE: 208.38m
SPEED LIMIT: 8.94m/s
EPISODE ENDS: 117.10sec
REASON: vehicle




BEV COSTMAP:



End-to-End Urban Driving by Imitating a Reinforcement Learning Coach [Zhang et al. 2021]

Rendered procedural simulation is
insufficiently realistic and diverse
for real-world embodied AI





Generative AI for Autonomy

June 2023– first cut



4x playback speed



Diverse Futures



Generated by GAIA-1



Action Conditioning




Generated by GAIA-1



GAIA-2 – March 2025



Front Left



Front




Front Right



Back Left





Back Right



Multi-view | Higher resolution | Geographic diversity | Controllability | Scale







Real Data

GAIA-2

GAIA-2

GAIA-2

GAIA-2

GAIA-2

Back Left

Front Left

Front Forward

Front Right

Back Right

Generating out-of-distribution data



GAIA-3 – Dec 2025



Realism | Controllability | Geographic diversity | Embodiment transfer



Increased realism



Ego-vehicle control





Drift Left

Drift Right

Too Fast

Too Slow



Drift Left

Drift Right

Too Fast

Too Slow



Drift Left

Drift Right

Too Fast

Too Slow



Drift Left

Drift Right

Too Fast

Too Slow



Drift Left


Drift Right

Too Fast

Too Slow



Embodiment transfer



Embodiment transfer




Could we learn to drive in a
world model?







World Models
[Ha & Schmidhuber, 2018]



DreamerV1-3
[Hafner et al. 2019-2025]

Dreaming about driving



Wayve's 2018 Blog Post: <https://wayve.ai/thinking/dreaming-about-driving/>





Wayve's 2018 Blog Post: <https://wayve.ai/thinking/dreaming-about-driving/>


What is the reward function for embodied AI?




Hand designed rules



Imitation + intervention



Spatio-physical reasoning



Language models



Preference labeling



Something new?



New industry consensus:
End-to-end learning for driving

Contrarian belief #2:

For the same reasons, end-to-end learning is best for safety and simulation too!







GAIA


Learning risk in simulation, training on crash scenarios generated by GAIA-3



Learning risk in simulation, training on crash scenarios generated by GAIA-3







Synced time: 2026-02-10 1:15:53.783 AM CET



NCAP Testing
Data Collection

Model fails ❌



Counterfactual
Resimulation

Model passes ✅



Testing in London

Failed to Slow ❌



GAIA

Counterfactual
Resimulation

Model Slows Speed ✅



Testing in Bay Area

Failed to Avoid Vehicle ❌



GAIA


Counterfactual
Resimulation

Model Changes Lanes ✅



Testing in Japan

Failed to Give Way ❌




Counterfactual
Resimulation

Model Stops for Runner ✅

SCALING THE AUTONOMY ECOSYSTEM

Contrarian belief #3:
Platform will reach greater scale
than vertically integrated autonomy

3 business models exist for autonomy



STELLANTIS

WAYVE



NISSAN
MOTOR CORPORATION



WAYVE





WAYVE




nvidia.

Uber

A sleek, silver autonomous vehicle is shown from a rear three-quarter perspective. The car is positioned in the center of the frame, with its rear facing slightly to the left. The background is a gradient of blue and purple, with concentric, glowing lines radiating from the car, suggesting a sensor field or a futuristic environment. The text "WAYVE | Uber" is overlaid in the center of the image, in a white, sans-serif font. The car has "WAYVE" and "Uber" logos on its rear and side.

WAYVE | Uber

AI deployment strategy




...

Contrarian belief #4: cost matters!

Are supervised 'reasoning' approaches too label and data hungry?

More interesting questions for Wayve:

1. Can we drive 10M miles virtually within an hour to do validation with world models?
2. How can we achieve L5 with mass-market \$1000 hardware per car?
3. How to assemble diverse data partners for efficient fleet learning?
4. Can world models run in real time on embedded hardware?
5. Learning one representation agnostic to vehicle/sensor/compute?
6. How can we generalize validation evidence between applications?
7. Physical AI MuZero moment where a single model can explore, plan and critic in real time?



Supported by:



STELLANTIS



Uber

Qualcomm



Contrarian belief #5:
Consumer vehicles, *before*
Robotaxis, *before*
Mobility, *before*
Manipulation.



Autonomous driving will be the first frontier of Embodied AI

	Personal vehicles	Robotaxis	Goods / Logistics	General Mobility	Domestic Robotics	Warehouse Robotics
Availability of large-scale data	●	●	●	●	●	●
Hardware maturity	●	●	●	●	●	●
Global supply chains	●	●	●	●	●	●
Compelling business case	●	●	●	●	●	●
Regulation	●	●	●	●	●	●



What can we learn from self driving?

Constrained environments and simplification (Phoenix AZ, HD-maps) failed to create general purpose solutions

Design for global scale from beginning, tackle hardest problems first

Need a strategy for large-scale data collection and self-supervised learning for training and validation

Is large scale motion capture data a mistake for manipulation?

Manufacturing new robots or large capex infrastructure shifts are not feasible to change industries

Are bespoke robots the right approach vs bringing intelligence to incumbent systems?



Embodied AI Vision for Success

- Slot easily into our messy world
- Be economically scalable as a business
- Be loved by public and regulators
- Meet human safety / reliability expectations
- Add value and utility to our lives



wayve

L A B S





vciy

L A B



50 PhDs 11 Best Paper Prizes 2 Longuet-Higgins Prizes ...

CA Vancouver GB London US Sunnyvale JP Tokyo DE Stuttgart Tel Aviv

wayve

Policy learning

Foundation models

World models

Simulation

Research engineering

Model-based RL

Large-scale training

Distillation

Offline RL

Reward models

Spatial Intelligence

Data, data, data

Language for embodied AI

Cross embodiment learning

Special projects

Model compression

...

50 PhDs 11 Best Paper Prizes 2 Longuet-Higgins Prizes ...

CA Vancouver GB London US Sunnyvale JP Tokyo DE Stuttgart ...

Please reach out at
labs-careers@wayve.ai
for FTE and intern
applications



Q&A

