

10 Questions to the Humanoid Humanoid Industry

Ten open questions, ten hypotheses, and the evidence behind each one

Lars-Fredrik Forberg

Presented on behalf of Humanoid.guide

CEO, Kinetic Blocks AS

Future Pulse 2026

The Dutch Robotics Conference

Evoluon, Eindhoven

15 September 2026


01

QUESTION 01

How many humanoid robots are actually actually doing paid industrial work today?
today?

Four independent sources now let this be built from the bottom up, and the answer is hundreds in genuine production work rather than the tens of thousands the shipment figures imply

Humanoids deployed in the first half of 2026, and what is left after each known deduction



Key insights

- Accenture put 19,000 deployed in the first half of half of 2026, with 97 percent sourced from China, China, and said most of it is still proof of concept, concept, internal testing or education.
- Up to 70 percent of Chinese output in that half went to training centres, about 13,000 machines across 53 facilities with 34 more under construction.
- IDC found more than 85 percent of the 2025 base in performances, education, data collection and guided tours, reaching the same conclusion from a different direction.

Named industrial deployments we can count: count: Agility at GXO, Aptronik at Mercedes and Mercedes and Jabil, Figure at BMW, Humanoid at Humanoid at Schaeffler from late 2026, Wandercraft and Renault with 350 units planned planned before end 2027

30 to 50 percent of a human worker's worker's task efficiency, per per Unitree's founder

9 of 9 named Chinese deployments deployments on a fixed scene, single task, usually overnight



02

QUESTION 02

If the models are this good, what is it that still that still does not work?

Manipulation benchmarks are close to solved in the lab, while the best of ten policies reaches 12.8 percent on real hardware and European safety testing puts almost every collision result in the yellow to red band

Real-world task success rate by policy and robot embodiment, RoboDojo leaderboard, percent

	ARX X5	Piper	Piper X	All 18 tasks
π 0.5	13.3	21.7	3.3	12.8
InternVLA-A1	3.3	18.3	0	7.2
GalaxeaVLA	0	13.3	0	4.4
Xiaomi-Robotics-0	8.3	3.3	0	3.9
X-VLA	1.7	8.3	0	3.3
GR00T-N1.7	1.7	3.3	0	1.7
π 0	0	5.0	0	1.7
StarVLA- α	0	5.0	0	1.7
Spirit v1.5	0	1.7	0	0.6
DM0	0	0	0	0
Human teleoperation	100	100	100	100

Key insights

- LIBERO is treated as solved at 95 to 99 percent, and LIBERO-Plus drops the same models below 30 percent by moving the camera inside the same simulator.
- Fraunhofer IPA has run Unitree G1, Dobot Atom, Agibot G2 and Neura machines through 66 tests, and the G1 at 50 kilograms is not permitted for power and force limited operation under EU norms.
- Every one of the nine named Chinese deployments runs on a fixed scene with a known object set and no untrained public present, so the world is arranged for the machine.

Nine of the ten policies complete no task at all on Piper X, the third embodiment

RoboDojo, the sim-and-real benchmark behind these success rates success rates



03

QUESTION 03

What does it cost to teach a robot one task?
task?

Three separate stages in Stuttgart put the cost of teaching a single task at four hours, at five to ten hours, and at six months, which is the distance between a demonstration and an industrial line

Time to train a single task, as stated from three separate stages on one day

Fraunhofer IPA

One cable-handling task



4 hours

240 teleoperation trials

Boston Dynamics

One autonomous policy



5 to 10 hours

Behaviour cloning through VR

Renault and Wandercraft

One industrial operation



6 months

Teleoperation, world model, simulation

Key insights

- Fraunhofer needs 240 trials, about four hours of of teleoperation, for one cable-handling task, with task, with data efficiency as the explicit research research goal.
- Renault and Wandercraft need six months per industrial operation against a target of days, and kitting is still unsolved because dexterous hands broke on metal parts within days.
- Rhoda post-trains on one to ten hours of trajectory trajectory data per task, which is the low end of the of the same range and the number every buyer buyer should be quoting against.

3 hrs

to retask a humanoid today, against the five minutes the VDMA says industry needs

350

Wandercraft units planned into Renault plants before the end of 2027

Every speaker described a shortage of the right data, recorded on operating industrial lines, rather than a shortage of data in general



Renault and Wandercraft, the Calvin robot at work on the line



04

QUESTION 04

Should you buy your training data, build the
build the channel, or take it for nothing?
nothing?

Figure and Skild both built their own collection channels this month while LightWheel released 100,000 hours free on Hugging Face, which squeezes the vendor market from both directions at once

Three routes to training data, and what each one costs

Buy

\$3 to \$40

Per hour for egocentric human human video on the vendor market

\$70 to \$98

Per hour for real machine data data from Chinese training centres

Zero

Capital committed before the the first hour arrives

Build

\$1bn+

Figure's twelve-month data and and compute commitment

43,200

Hours uploaded per day through the Index app

1 to 3

Dollars on collection for every every three on quality control, control, per Skild

Take

100,000

Hours in LightWheel's EgoSuite-EgoSuite-Open100K on Hugging Hugging Face

15,000+

Tasks, across 128 scene types in types in seven categories

\$0

Price. First batch live, the rest rest rolling out

Key insights

- Skild spends three dollars on quality control for for every dollar spent collecting, and Chinese training centres produce three usable hours from from every eight recorded.
- Real machine data from those centres costs 500 to 700 yuan an hour, which is 70 to 98 dollars, against 3 to 40 for egocentric human video.
- Underneath all of it sits general web video, free and free and effectively unlimited, which Rhoda has has now shown improves a real industrial task. task. What stays scarce is the demonstration recorded on the customer's own floor.

Chinese training centres produce three usable hours from every eight recorded, and Boston Dynamics keeps annotation in-house because a supplier relationship does not reproduce the loop




05

QUESTION 05

Can a robot learn to work by watching video
video that has nothing to do with robots?
robots?

Scaling pre-training on general web video, containing no actions and almost nothing about robots, took at-speed completion on a real customer task from 4 to 85 percent across seven measured checkpoints

At-speed completion rate against pre-training quality, seven checkpoints, one industrial task, percent



DINO FD on held-out web video, log scale, lower is better, so pre-training quality increases to the right right

Model size sweep: XS 4, S 65, M 75, L 85

Compute sweep on M: 58, 67, 74 at 0.08x, 0.18x, 0.37x


200+ hours of real-robot evaluation, 100 to 260 trials per policy **31 pts** spread between budgets at a quarter of the data, 8 at full data

Key insights

- How well a model predicts held-out web video, measured before it has seen the task at all, ranked the seven checkpoints in the same order the robot did.
- Dyna-2 reached the same conclusion from egocentric human video, scaling from 20 to 53 percent across four orders of magnitude with no robot data in pretraining.
- The honest counter is that human and robot limbs move differently, that usable footage still has to be labelled by hand, and that Rhoda measured a correlation across one architecture family on one task.



Rhoda on the task the seven checkpoints were measured on



06

QUESTION 06

How should you tell a robot what you want it
want it to do?


On tasks the model has never seen, one video demonstration in context reaches 66 percent where language prompting reaches 9, which makes how the task is specified a bigger lever than the architecture

Average per-step success rate, demonstration prompting against language prompting, controlled study, percent

Tasks seen in pretraining, at 1,000 hours



Tasks never seen, at 100,000 hours



Key insights

- Both arms of the study used identical data, architecture and compute, so the gap comes from from the prompt rather than the model.
- Language wins at small scale, 53 against 43 percent on seen tasks at 1,000 hours, and loses decisively once pretraining grows.
- The commercial version of this question is retasking retasking time, which is three hours today and and needs to be five minutes.

380

post-training episodes a language-prompted policy needs to match what one demonstration in context achieves, with post-training only overtaking it at 2,000

11 min

from recording one human demonstration to the robot running a ten-minute task it had never seen, measured end to end on a single evening



Skild S1 executing a task it has never seen, from one demonstration




07

QUESTION 07

What is the one question your supplier would rather you did not ask?

Miles or hours between human interventions is the reliability measure frontier operators run their fleets on, and Europe is now the only region building an independent test regime for it

Miles between faults, and what each band requires of the operator



Key insights

- Below ten miles between faults a person walks beside the robot, and between ten and twenty-five it requires remote supervision.
- Fraunhofer IPA tests collision force, cybersecurity, energy use and cleanroom particle class, and is taking the results into ISO Working Group 12.
- Rhoda found two checkpoints from the same training run 28 points apart on the robot, and that validation loss picks almost the worst of them, so nothing short of a real trial tells you what you have.

66

tests in the Fraunhofer Fraunhofer IPA framework: capability, capability, functional safety, cybersecurity, energy use and cleanroom particle class class

Unitree G1 at 50 kg is not permitted for power and force limited operation under EU norms

Results now going into ISO Working Group Group 12


08

QUESTION 08


Unitree closed 460 percent up, then fell 30 percent. What was actually priced?

The first public price in this sector arrived and gave back a third of itself within a week, in the same quarter that growth decelerated from 333 percent to about 40

Unitree valuation, USD billion



Unitree revenue growth, percent



Key insights

- Figure sits above a thousand times sales on third-party revenue estimates, against low tens for listed hardware companies at steady state.
- Agility is listing at 2.5 billion dollars with 65,000 operating hours and over 300 million dollars of multi-year orders behind it.
- The number of limited partners behind venture funds has halved since 2022, concentrating capital into fewer and larger cheques while Cartwheel Robotics sits in involuntary Chapter 7.

\$2.5bn pre-money valuation at which Agility is going public by SPAC, with 65,000 operating hours and over 300 million dollars of multi-year orders behind it

1,000x Figure's multiple on third-party revenue estimates, against low tens for listed hardware

-50% limited partners behind venture funds since 2022

Sources: The Economist, 23 August 2026; Caixin; Bloomberg; Reuters, August 2026; Business Times, 24 August 2026; Agility Robotics and Churchill Capital Corp XI, June Capital Corp XI, June 2026; Robots & Startups, May 2026. The one-week figure is approximate.



09

QUESTION 09

What does this industry look like in 2032,
2032, once the robots are ordinary?

AMSTERDAM ROBOTICS

REPARATIE EN INNOVATIE

ROBOTICA
REPARATIE
& SERVICE

DE FIETSENMAKER
REPARATIE & ONDERHOUD



Three of the five layers in this value chain do not yet exist as businesses, and the integrators who deploy the fleets and own the operating data will carry more weight than the makers

Five layers, and how much of each exists as a distinct business today

Components	Consolidated. One supplier covers 60 to 70 percent of all humanoid makers	Mature
OEMs	150 to 200 brands sitting on a much smaller number of real number of real manufacturers	Consolidating
Integrators and fleet operators	Buy, deploy, lease, and own the operating data	Forming now
Second-hand and residual residual value	No resale market and no residual curve, which is why leasing costs more than buying	Does not exist
Service, parts and aftermarket	Workshops, limb swaps, batteries, wear parts, accessories	Does not exist

Agility, Apptronik and AGIBOT all sell robots as a service themselves, which is what happens when the integrator layer has not been built yet. It does not survive the first fleet of a thousand machines.

Key insights

- Behind 150 to 200 humanoid brands sits a much much smaller number of real manufacturers, and and one supplier provides harmonic drives to 60 to 60 to 70 percent of them.
- There is no second-hand market and no residual value curve, which is why leasing a humanoid today costs roughly twice what buying it costs over three years.
- Any limb on the latest Atlas can be replaced in under five minutes, and Renault's motors overheat after two to three hours of heavy lifting. Both are workshop arguments.

Sources: Werner Kraus, Fraunhofer IPA, from fourteen company visits in China, July 2026; Aya Durbin, Boston Dynamics; Renault and Wandercraft, all presented in Stuttgart, presented in Stuttgart, September 2026; published RaaS pricing benchmarks, 2026.




10

QUESTION 10

What can Europe actually own in 2030?

China controls roughly 90 percent of magnet processing and actuation is 40 to 60 percent of the bill of materials, leaving Europe the integration, certification and operations layer

Estimated bill of materials for Optimus Gen 2, USD thousand



Where the value sits along the chain



McKinsey expects the supply chain to split in two rather than one side winning, with export markets likely to accept Chinese hardware while restricting software and data flows. That is the mechanism behind both the FCC determination and the new Agibot plant in Serbia.

Key insights

- McKinsey expects the supply chain to split in two rather than one side winning, with Europe differentiating on safety-certified, high-assurance deployment.
- Export markets may accept Chinese hardware while restricting software and data flows, which is the mechanism behind both the FCC determination and the new Agibot plant in Serbia.
- Humanoid raised 152 million dollars with Bosch and Schaeffler alongside it, Wandercraft is putting 350 units into Renault, and Verity won the IERA award at ICRA in Vienna.



Three things to take back to the office

1

Ask suppliers for hours in paid operation before you ask for units shipped

Almost every published figure counts deliveries, while the number that predicts value is time in production at a customer site.

2

Price the cost of teaching one task before you price the machine

Four hours, ten hours or six months is the difference between a pilot and a plant, and it is the number that never appears in a quotation.

3

Build the position in the three layers that do not exist yet

Fleet operations, residual value and service. The component race is decided; the workshop on the corner is not.
