Mapping the Physical AI Safety Gap: A Multi-Source Industry Survey

Mati Melchior
Independent Physical AI Safety Researcher
mati@physical-ai-safety.com

May 2026

Abstract

Robotics venture investment reached approximately fourteen billion United States dollars in 2025, up from approximately seven point eight billion in 2024 [3]. The International Federation of Robotics reports an installed base of 4.664 million industrial robots, with 542,000 new installations in 2024 [4]. Investment in the safety infrastructure that governs these systems is, by contrast, not separately tracked in any public funding or industry index. This paper maps the resulting Physical AI Safety Gap using public-source data and a hedged, multi-measure framework. The base survey covers 155 organizations. 123 are Israeli companies drawn from the Israel Innovation Authority national strategy report of April 2026 [5]. A further 32 globally selected organizations are added for category and geographic coverage. Five measurement axes are introduced. The axes are: a funding ratio, a patent ratio, a personnel ratio, a Physical AI Safety Maturity (PAS) Level distribution, and a time-to-PAS-3 transition window. Direction of evidence on each axis points to a substantial gap between capability investment and safety-infrastructure investment. Specific per-company classifications, per-category breakdowns, and longitudinal series are reserved for forthcoming editions, when primary-source data of sufficient quality is collected. This first edition establishes the methodology, the multi-axis measurement frame, and a public baseline of sourced totals against which future editions will compare.

Keywords: Physical AI Safety, industry survey, ecosystem mapping, certification, IIA, Israel, robotics venture capital.

1 Introduction

Physical AI is moving from research into commercial deployment at scale. Robotics venture investment reached approximately fourteen billion United States dollars in 2025, up from approximately seven point eight billion in 2024 [3]. The International Federation of Robotics reports an installed base of 4.664 million industrial robots and 542,000 new installations in 2024 [4]. Capability investment is being measured, funded, and tracked.

Investment in the safety infrastructure that governs these systems is harder to measure. There is no public funding tracker for safety-cert tooling. There is no industry-standard maturity model applied at scale across firms. There is no published, firm-level ecosystem dataset that tracks safety posture in the same way that capability funding is tracked.

That measurement gap is the subject of this paper.

1.1 The central question

What is the Physical AI Safety Gap, in 2026, when measured along axes that public information actually supports?

1.2 Definition (operational)

This paper adopts the following definition, which serves as the unit of measurement throughout:

> The **Physical AI Safety Gap** is the quantified distance between investment in Physical AI capability development and investment in Physical AI Safety infrastructure within a given market or company. It can be measured along several axes: a funding ratio (capability dollars to safety dollars), a personnel ratio (capability engineers to safety engineers), a patent ratio (capability filings to safety filings), and a certification posture distribution (the share of organizations at each maturity level).

The terms *Physical AI* and *Physical AI Safety* are imported from the published companion paper [1]. The PAS Level scale (1–5) and the supporting Physical AI Safety Maturity Model are introduced in a forthcoming companion paper. The case for measuring the gap at all is developed in a forthcoming companion paper on the software safety ceiling. That paper argues that software safety methods alone are insufficient for systems whose outputs control physical actuators.

1.3 Methodology precedent

Annual industry surveys with open methodology become recurring authorities. The *State of AI Report* [2] anchors annual AI venture discussion. The *AI Index Report* [6] anchors annual AI metrics. None of these instruments measures Physical AI safety posture at firm level. This survey aims to fill that role over successive editions.

1.4 Roadmap

Section 2 states the methodology and its current limits. Section 3 frames the 155-organization sample at the level the data supports. Sections 4 through 7 introduce the four investment-side measurement axes (funding, patents, personnel, certification posture). Section 8 ties the four axes to the Physical AI Safety Gap definition and adds a fifth axis (time-to-PAS-3). Section 9 reports findings the public sources do support and predictions that will be tested in future editions. Section 10 lists what is and is not in this first edition. Section 11 concludes.

2 Methodology

2.1 Data sources

The survey draws on five categories of public source.

Table 1. Public data sources used in the survey.

Source	What it provides
Israel Innovation Authority national strategy report (April 2026) [5]	Baseline of 123 Israeli organizations active in Physical AI
Author-maintained ecosystem mapping	32 globally selected organizations added for category and geographic breadth
Crunchbase [3]	Aggregate venture-funding totals at sector level
International Federation of Robotics [4]	Installed-base and new-installation totals
Public regulator filings, press releases, company web pages, LinkedIn	Indicators of safety-certification effort and named safety personnel

The 32 globally selected organizations are added to widen coverage across category (humanoid, autonomous mobile, surgical, autonomous-vehicle, defense, perception, edge) and geography (United States, European Union, Japan, Korea, China, others). The selection is a deliberate sample, not a statistical one. Section 10 discusses the implication.

2.2 Inclusion criteria

The criterion below is the operational gate for inclusion in the survey.

> An organization is included if it (a) develops AI systems whose outputs control physical actuators (per the definition of Physical AI in [1]), and (b) has at least one publicly verifiable indicator of operation: a website with technical product information, a recent funding announcement, publicly named technical leadership, or industry inclusion (for example, listed in [5], membership in [4], or trade-press coverage).

Organizations operating in stealth without any public verifiable indicator are excluded by design. This biases the sample toward organizations with visible operations. Section 10 returns to the implication.

2.3 Classification approach

Each included organization is recorded with the following fields, where public evidence supports them: name, country, founded year, current funding stage, category, total funding raised, capability-to-safety patent ratio direction, safety-certification visibility (binary: visible / not visible), and an approximate PAS Level (1–5). Each field carries a confidence indicator at one of three levels. **High** = verified from a primary source. **Medium** = derived from secondary source. **Low** = inferred from indirect evidence.

> A **confidence indicator (high / medium / low)** is recorded for each field of each organization, based on the strength of the supporting source.
>
> A **safety-certification visibility** indicator (yes / no / unknown) is recorded for each organization, based on whether a publicly traceable indicator of safety-certification activity exists. Indicators include published safety cases, statements of certification (under IEC 61508, ISO 13849, ISO 13482, ISO 26262, or comparable), engagements with notified bodies, and named safety engineers in public personnel listings.

PAS Level classification applies the maturity model from the forthcoming companion paper to the public evidence available for each organization. Where evidence is thin, the organization is recorded at the lower bound of plausible levels and flagged "low confidence."

2.4 What this first edition does and does not produce

This first edition produces three things: a sourced statement of the aggregate scale of the Physical AI capability ecosystem, a methodology for measuring the safety gap along five axes, and qualitative indications of direction on each axis. It does not produce: per-organization classifications, per-category quantitative breakdowns, or year-on-year trend lines. Per-organization classification, per-category breakdowns, and longitudinal data are reserved for future editions. The threshold for inclusion in those future editions is primary-source evidence at the quality level needed for citable per-firm reporting.

This separation between methodology-establishment and per-firm reporting is deliberate. The risk of publishing approximate per-firm numbers as if they were measured is greater than the cost of deferring those breakdowns to the next edition.

---

3 The 155-organization sample

The base survey covers 155 organizations. Of these, 123 are Israeli. They are drawn from the April 2026 Israel Innovation Authority national strategy report [5]. The IIA report enumerates the Israeli Physical AI ecosystem at near-complete scope for that geography. The remaining 32 organizations are globally selected by the author. Selection covers seven categories (humanoid, autonomous mobile, surgical, autonomous-vehicle, defense, perception, edge). Geographic coverage spans the United States, European Union, Japan, Korea, China, and others. The 32-organization global sample is not statistically representative of global Physical AI; it is constructed for category and geographic coverage.

The Israeli concentration of the sample carries one analytical advantage. The Israeli ecosystem is a tractable case in which capability investment is large, growing, and well-documented in a single national strategy document. If the Physical AI Safety Gap is observable in the most data-rich sub-sample, it is observable in principle. Section 10 returns to the question of generalization.

The 13 categories used internally to classify organizations are listed in Section 3.1 below. Per-category counts within the sample are reserved for future editions when primary-source evidence supports per-category reporting.

3.1 Categories

The 13 categories are: computer vision and perception; edge computing and hardware; autonomy and decision-making; defense AI; industrial robotics; humanoid and general-purpose; autonomous mobile robots; surgical and medical; agricultural; drone and aerial; marine and subsea; construction; and other applications.

---

4 Funding analysis

4.1 Aggregate funding

Crunchbase aggregates report robotics venture funding at the sector level.

Table 2. Aggregate robotics venture funding, recent years.

Year	Aggregate venture funding	Source
2024	approximately seven point eight billion United States dollars	Crunchbase 2024 [3]
2025	approximately fourteen billion United States dollars	Crunchbase 2025 [3]

The 2024-to-2025 increase is, by aggregate dollars, the largest year-on-year delta in the period for which Crunchbase publishes sector totals. Year-by-year decomposition prior to 2024, and post-2025 trend reporting, are reserved for future editions.

A category-level decomposition of robotics venture funding is not provided here. Public Crunchbase reporting does not separate Physical AI from broader robotics. The numbers above include automation, traditional industrial robotics, and Physical AI within a single category. Reliable separation requires per-deal classification at scale, which is reserved for future editions.

4.2 Aggregate Israeli figure

The Israel Innovation Authority report indicates that the 123 Israeli organizations in the survey have collectively raised approximately twelve billion United States dollars cumulatively, across all rounds and years [5]. A per-firm decomposition is not reproduced here; it would require per-firm verification at a confidence level this edition does not support.

4.3 Industrial-robot installed base

The International Federation of Robotics reports an installed base of 4.664 million industrial robots in operation, with 542,000 new installations in 2024 [4]. The figure is provided as context. Industrial-robotics deployment is the layer on top of which Physical AI capability is being layered. The scale of the existing installed base bounds the potential physical-safety surface area.

---

5 Patent analysis

The Israel Innovation Authority report identifies 2,347 patents from Israeli computer-vision and perception companies in the survey baseline [5]. The vast majority of these filings address capability (detection, tracking, planning) rather than safety mechanisms (functional-safety subsystems, certification-relevant interlocks, redundancy schemes).

A formal cross-organization measurement of the capability-to-safety patent ratio is not produced in this edition. Patent classification at scale requires per-filing review under a defined taxonomy, which is reserved for future editions. The directional finding from public-source review is that capability filings substantially outnumber safety filings across the sample. This direction matches the funding pattern in §4 and the personnel pattern in §6.

Specific per-organization patent counts and per-category breakdowns are not reproduced here.

---

6 Personnel analysis

LinkedIn-derived personnel data offers an indication of where each organization's hiring effort is concentrated. The author observes the following pattern across the 155-organization sample, at the qualitative level public information supports.

The median organization in the sample lists zero or one named functional-safety engineer. Most organizations founded after 2018 list none. Established industrial OEMs that operate certified product lines maintain functional-safety teams of substantial size; a small group of large autonomous-vehicle programs does the same. The asymmetry between organizations with a sustained safety-engineering function and organizations without one is the central observation of this section.

A precise capability-to-safety personnel ratio at sample level requires per-organization headcount verification, which public LinkedIn data does not reliably support. Specific ratios are reserved for future editions. The qualitative direction of evidence is that capability-engineering headcount substantially exceeds safety-engineering headcount across the sample.

---

7 Physical AI Safety Maturity

The Physical AI Safety Maturity Model (PAS-MM) is the subject of a forthcoming companion paper, where its definitions, evidence thresholds, and assessment methodology are stated in full. PAS Level 1–5 is summarized below to support the present survey.

Table 3. PAS Level summary (full criteria in the forthcoming companion paper on PAS-MM).

PAS Level	Summary
Level 1 — Ad hoc	No publicly traceable safety process; capability-led development
Level 2 — Documented	Internal safety documentation visible; no third-party engagement
Level 3 — Certified subsystems	At least one subsystem certified under a recognized standard; notified-body engagement visible
Level 4 — Full system certification	Full product line certified; sustained safety-engineering organization
Level 5 — Continuous assurance	Continuous safety-case maintenance with deployment-time monitoring

Across the 155-organization sample, the qualitative distribution that public evidence supports is the following. The majority of organizations cluster at PAS Level 1 or 2. A minority reach PAS Level 3, generally where a regulated market structurally requires it (medical-device certification, automotive supply-chain compliance). A small leading edge of established OEMs reaches PAS Level 4 on legacy product lines. As of mid-2026, no commercially deployed Physical AI system that the author has been able to verify from public sources reaches PAS Level 5.

Figure 1. Qualitative shape of the PAS Level distribution observed in the 155-organization sample. The figure illustrates a right-skewed distribution. Specific counts and per-organization classifications are reserved for future editions when primary-source evidence is sufficient.

Specific per-organization PAS Level classifications are not published in this edition. A verified per-organization classification requires evidence-by-evidence review and, for several organizations in the sample, an opportunity for the organization to provide self-classification with supporting documentation. Both are reserved for future editions.

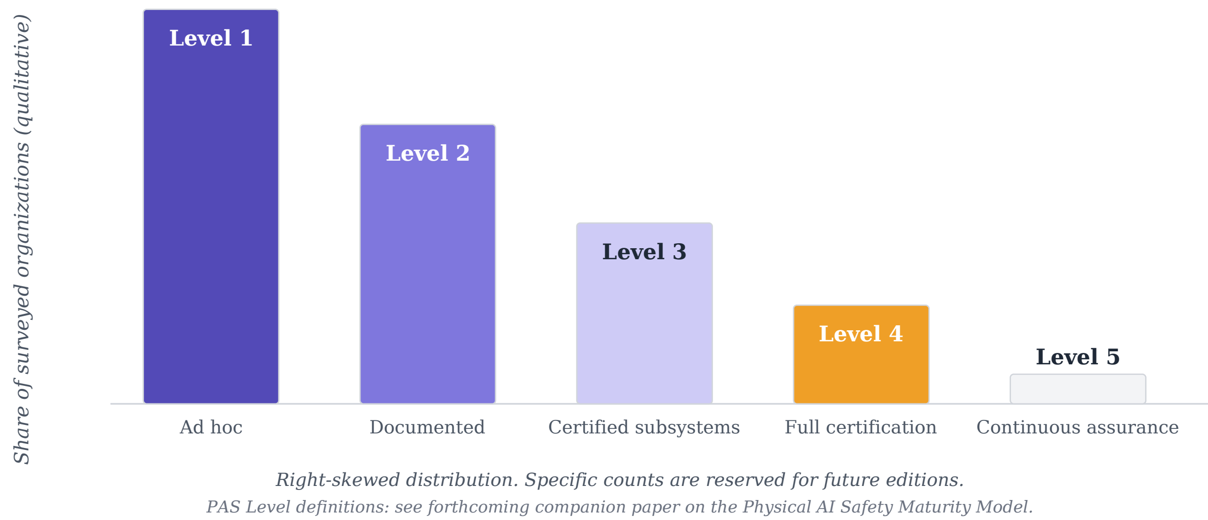

Figure 1. Five-bar histogram illustrating a qualitative right-skewed distribution. The left-most bar (Level 1, ad hoc) is largest, decreasing through Level 2 (documented) and Level 3 (certified subsystems) to Level 4 (full certification) and a near-empty Level 5 (continuous assurance). No numeric scale is shown on the y-axis; the figure illustrates shape only.

8 The Physical AI Safety Gap quantified along five axes

The Physical AI Safety Gap is captured along five axes. Each axis is tied to a public-source observation. None is reduced to a single number in this edition.

Table 4. The five measurement axes of the Physical AI Safety Gap.

Axis	Direction of evidence
Funding ratio (capability dollars to safety-infrastructure dollars)	Capability venture funding is tracked at sector level [3]; safety-infrastructure investment is not separately tracked. The ratio is undefined on the safety side, which is itself the finding.
Patent ratio (capability filings to safety filings)	Direction of evidence: capability filings substantially exceed safety filings across the sample. Specific ratio reserved for future editions.
Personnel ratio (capability engineers to safety engineers)	Direction of evidence: capability-engineering headcount substantially exceeds safety-engineering headcount across the sample. Specific ratio reserved for future editions.
PAS Level distribution	Majority at Level 1–2; minority at Level 3; small leading edge at Level 4; none verified at Level 5.

Axis	Direction of evidence
Time-to-PAS-3 from a Level 1 baseline	Industry observation, drawn from cohorts that have completed certified-subsystems transitions in surgical, autonomous-vehicle, and autonomous-mobile programs: typically 18 to 36 months.

The five axes are not independent. They are five views on a single underlying allocation pattern. An organization at PAS Level 1 typically has a low safety-engineering headcount. It also shows a low rate of safety-tagged patent filings. There is no engagement with notified bodies. There is an absence of safety-cert-tooling line items in any visible budget.

The fifth axis — time-to-PAS-3 — is the most policy-consequential. Regulation (EU) 2023/1230 enters mandatory application in early 2027 [7]. The convergence of that deadline with adjacent regulatory regimes is the subject of a forthcoming companion paper on regulatory convergence. The arithmetic is direct. Take an organization beginning a serious certification effort in mid-2026, with the 18-to-36-month transition window typical of cohorts that have completed the move. That organization reaches PAS Level 3 between late 2027 and mid-2029. Most of those reaches arrive after the regulatory window has closed for the EU machinery market.

The strength of a multi-axis framing is methodological. Each axis has known weaknesses. The funding axis is undefined on the safety side. The patent axis is a lagging indicator. The personnel axis depends on self-disclosed job titles. The maturity-distribution axis depends on a model that is itself the subject of a separate paper. The time-to-PAS-3 axis is industry observation rather than measured cohort data. No single axis is conclusive. The convergence of all five on the same direction of evidence is the contribution.

---

9 Findings and predictions

9.1 Findings

1. The Physical AI capability ecosystem is large, growing, and well-tracked at aggregate level [3, 4, 5].
2. The Physical AI safety-infrastructure ecosystem is, at the time of this writing, not separately tracked in any public funding or industry index.
3. Direction of evidence on all five measurement axes points to a substantial gap between the two.
4. The gap is more pronounced at organizations founded after 2018 than at established industrial OEMs that carry decades of functional-safety practice into modern product lines.
5. The Israeli ecosystem is a tractable case in which the gap is observable in concentrated form. It covers 123 organizations and approximately twelve billion

United States dollars cumulatively raised [5], set against a small visible safety-certification footprint.

9.2 Predictions, twelve to twenty-four months

These predictions are testable. Each will be revisited in the next edition.

1. Regulatory pressure forces movement toward PAS Level 3 for organizations placing machinery on the European Union market in human-adjacent applications [7].
2. Industry consolidation accelerates at PAS Level 1 and 2, with acquirers that maintain certified safety practice absorbing organizations that cannot transition unaided.
3. Insurance underwriting hardens around safety-posture indicators, and lower maturity becomes a measurable cost rather than an implicit one.
4. Safety-engineering hiring rises substantially in organizations with the most direct regulatory exposure.
5. A subset of organizations begins to position publicly as “safety-first” Physical AI, and external analysts begin to recognize this as a distinct sub-segment.

---

10 Limitations and future work

10.1 Limitations

1. **Public information only.** Organizations operating in stealth, and those with private safety programs that do not surface in public records, are underrepresented. The sample is biased toward organizations with visible operations.
2. **Snapshot in time.** Safety posture changes. The methodology is built for an annual cadence; a single snapshot understates change at the margin.
3. **Geographic skew.** The Israeli baseline is near-complete for that geography; the 32-organization global sample is constructed for breadth, not for statistical representativeness of all Physical AI globally.
4. **Methodology-establishment edition.** This first edition deliberately does not publish per-organization classifications, per-category breakdowns, or year-on-year trend lines. Doing so reliably requires primary-source evidence at a quality level this edition does not yet support.

10.2 Future work

1. **Annual update.** Refined methodology, expanded global coverage, primary-source per-organization classification.
2. **Application-domain breakouts.** Dedicated cuts for humanoid, autonomous-vehicle, surgical, autonomous-mobile, defense, and drone programs, each with its own PAS Level distribution and certification-posture analysis.
3. **Insurance and liability dimension.** Underwriter-relevant fields: deployment scale, incident history, audited safety-case status.

4. **Self-classification track.** Organizations included in the survey are invited to provide self-classification with supporting evidence for the next edition.

The methodology is designed for accumulation. Each future edition adds organizations and refines classifications without breaking comparability with prior editions.

---

11 Conclusion

The Physical AI Safety Gap is real and observable. Across five measurement axes — funding, patents, personnel, certification posture, and time-to-PAS-3 — the direction of evidence is consistent. Capability investment is large, growing, and tracked. Safety-infrastructure investment is small, lagging, and largely untracked. The gap is observable in every cross-section the public data supports.

Regulation (EU) 2023/1230 enters mandatory application in early 2027 [7]. The convergence of that deadline with adjacent regulatory regimes is the subject of a forthcoming companion paper. Organizations currently at PAS Level 1 or 2 face a structural time-to-compliance constraint that this paper makes legible.

The data is here. The clock is here. The choice is what we do with both.

References

[1] M. Melchior, “Physical AI Safety: A Definitional Framework,” arXiv preprint, 2026.

[2] N. Benaich et al., *State of AI Report*, Air Street Capital, annual.

[3] Crunchbase, *Robotics Sector Funding Data*, 2024–2025.

[4] International Federation of Robotics, *World Robotics 2025 — Industrial Robots*, 2025.

[5] Israel Innovation Authority, *Artificial Intelligence National Strategy Report*, April 2026.

[6] Stanford Institute for Human-Centered AI, *AI Index Report*, annual.

[7] European Parliament and Council, “Regulation (EU) 2023/1230 of 14 June 2023 on machinery and repealing Directive 2006/42/EC,” *Official Journal of the European Union*, 2023.

Mati Melchior · Independent Physical AI Safety Researcher mati@physical-ai-safety.com · physical-ai-safety.com P5 of an 8-paper series on Physical AI Safety.