Regulatory Convergence in Physical AI: A Comparative Survey of the European Union, United States, Japan, and Israel (2025–2027)

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

Version 1.0 · May 2026

Abstract

Regulatory regimes for Physical AI are emerging across major markets in the period 2025–2027. This paper surveys four jurisdictions — the European Union, the United States, Japan, and Israel — and identifies a convergence pattern. Despite different legal traditions, four common requirements appear in each jurisdiction. These are third-party-certifiable hardware safety mechanisms, AI-specific risk assessment, traceability of decision-making, and incident reporting. We term this pattern the *Regulatory Convergence Window*. The window opens with two EU regulations. The first is Regulation (EU) 2023/1230 (the Machinery Regulation, effective January 20, 2027). The second is the EU AI Act (Regulation 2024/1689). The United States layer is sectoral, anchored on the NIST AI Risk Management Framework, alongside FDA, NHTSA, and OSHA guidance. Japan operates through METI-coordinated industrial policy. Israel's Innovation Authority published an April 2026 AI Strategy identifying the certification gap as a strategic priority. Companies operating across these markets face a 24-month compliance window. Within that window, PAS Level 3 or higher (per the PAS-MM framework introduced in a forthcoming companion paper) becomes the de facto baseline for safety-critical Physical AI applications. We provide a compliance roadmap structured by company stage and by jurisdictional priority. Cross-jurisdictional sequencing analysis indicates that EU 2023/1230 compliance creates a foundation meeting approximately 70 to 80 percent of United States and Japanese requirements. The convergence is not coordinated policy. It reflects shared underlying concerns. The convergence is robust and likely to deepen.

Keywords: physical AI safety, regulatory convergence, EU machinery regulation, EU AI Act, NIST AI risk management framework, IEC 61508, ISO 13849, certification

---

Disclaimer. This paper presents policy analysis. It does not constitute legal advice. Companies should consult qualified counsel in each jurisdiction for compliance decisions affecting their products.

---

1 Introduction

Physical AI is moving from research to deployment. In 2025 and 2026, AI-driven autonomous systems were deployed in industrial cells, surgical theatres, public roads, agricultural fields, and care facilities at increasing scale. Each deployment raises a certification question that did not arise for software-only AI systems: when an AI-driven system causes physical harm, accountability turns on who certified the system as safe to deploy.

Across four major Physical AI markets — the European Union, the United States, Japan, and Israel — answers are emerging. The answers differ in legal form. They share an underlying structure. This paper examines that structure.

The European Union approach is prescriptive and unified. Regulation (EU) 2023/1230, the Machinery Regulation, takes effect on January 20, 2027. It replaces Directive 2006/42/EC with no transition period. Annex III §1.1.6 explicitly regulates AI/ML systems as machinery for the first time in EU law. The EU AI Act (Regulation 2024/1689), adopted in June 2024, layers AI-specific obligations on top of the machinery framework. Together they compose the first horizontally applied Physical AI regulatory regime.

The United States approach is sectoral and pluralistic. The NIST AI Risk Management Framework provides a voluntary backbone. Sector regulators — the FDA for medical devices, NHTSA for autonomous vehicles, the FAA for unmanned aerial systems, OSHA for workplace robotics — apply existing frameworks to AI-driven systems within their domains. State-level activity in California, New York, and others adds a further layer.

Japan's approach is industry-coordinated. METI sets strategic direction. The Japan Robot Association, through JISC, leads standards development. The PMDA approves medical robots. Industrial robots have operated under JIS B 8433 for decades. The pattern is consistent: industry leads, regulation follows, convergence with EU and United States standards emerges over time.

Israel's approach is innovation-focused. The Israel Innovation Authority (IIA) published its AI Strategy in April 2026, identifying the safety certification gap as a strategic priority. Domestic regulatory capacity is nascent. Israeli Physical AI companies, facing a small domestic market, typically prepare for EU and United States compliance from the start.

1.1 The Convergence Claim

Four jurisdictions, four legal traditions, four modes of regulatory production. When the requirements are catalogued side-by-side, four common requirements appear in each:

1. **Third-party certifiability** of safety-critical Physical AI components.
2. **AI-specific risk assessment** with documented methodology.
3. **Decision traceability** — auditable records of AI-driven decisions affecting physical outcomes.
4. **Incident reporting** to regulators within defined timeframes.

We term this pattern the **Regulatory Convergence Window**, defined as follows:

> **Definition 1 — Regulatory Convergence Window.** The Regulatory Convergence Window is the 2025–2027 period during which regulatory regimes for Physical AI in the European Union, United States, Japan, and Israel converge on common requirements: third-party-certifiable hardware safety mechanisms, AI-specific risk assessment, traceability of decision-making, and incident reporting. Companies operating across these markets face a 24-month compliance window during which they must achieve PAS Level 3 or higher (per the PAS-MM framework from P4) for safety-critical Physical AI applications.

PAS Level 3 corresponds to documented hazard analysis, validated safety functions, and evidence of third-party certification capability for the relevant product category and jurisdiction.

The convergence is not coordinated policy. It reflects shared underlying concerns: physical harm from AI-driven systems, public trust, market access pressure from large buyers (industrial integrators, hospital systems, fleet operators), and insurance-industry pricing of unmitigated risk. Where coordinated policy might produce a single international standard, concern-driven convergence produces parallel requirements with similar substance and divergent procedure.

1.2 Paper Structure

Section 2 examines the EU regulatory landscape: Regulation (EU) 2023/1230, the AI Act, the harmonized standards landscape, and notified-body capacity. Section 3 covers the United States: federal AI policy, sectoral regulators, state activity, and the regulatory pluralism challenge. Section 4 addresses Japan: METI, industry-led standards, and sectoral rules. Section 5 examines Israel: IIA's April 2026 strategy, regulatory sandbox proposals, and sectoral regulators. Section 6 develops the convergence analysis with the four-jurisdiction matrix and the 2025–2027 window thesis. Section 7 provides a compliance roadmap structured by company stage and by jurisdictional priority. Section 8 concludes.

This paper is the seventh in an eight-paper series on Physical AI Safety. Companion papers in the series develop a definitional framework [18], the PAS-MM Maturity Model [19], software safety bounds [20], an empirical industry survey [21], common-cause failure analysis [22], and a certification framework [23]. The present paper provides the regulatory analysis layer of that series. Where this paper imports concepts from companion papers, the importing paper is cited at the point of import.

---

2 European Union Regulatory Landscape

The EU operates the most advanced Physical AI regulatory regime in 2026. Two horizontally applied regulations form the regime's spine: Regulation (EU) 2023/1230 (the Machinery Regulation) and Regulation (EU) 2024/1689 (the AI Act). A harmonized-standards layer translates the regulations into testable engineering requirements. A

notified-body certification layer provides conformity assessment. Each layer is examined below.

2.1 Regulation (EU) 2023/1230 — The Machinery Regulation

Regulation (EU) 2023/1230 takes effect on January 20, 2027. It replaces Directive 2006/42/EC. There is no transition period: on the effective date, the Directive ceases to apply and the Regulation applies directly across all 27 Member States.

The shift from Directive to Regulation is structural. A Directive requires national transposition; Member States retain interpretive latitude. A Regulation applies uniformly. Member States cannot vary its scope or technical requirements. For Physical AI, this matters: a system certified under the Regulation is certified for all 27 Member States simultaneously, and a system non-compliant in one jurisdiction is non-compliant in all.

Four provisions define the Regulation's posture toward Physical AI.

Annex III §1.1.6. This provision regulates AI/ML systems as machinery for the first time in EU law. Where the predecessor Directive had no concept of machine learning at the time of its 2006 drafting, Annex III §1.1.6 explicitly addresses systems whose behaviour is shaped by training data and learned models. The provision requires that AI/ML behaviour be characterised, bounded, and validated as part of the conformity assessment.

Six categories of high-risk machinery. The Regulation identifies six categories of machinery for which third-party certification is required regardless of which harmonized standards the manufacturer applies. Self-certification, permitted under the Directive for many product categories, is no longer sufficient for high-risk machinery under the Regulation. Categories include machinery with embedded safety functions implemented by AI/ML, removable mechanical transmission devices in certain configurations, and several others. The practical effect for Physical AI is straightforward: a robot whose safety functions depend on AI-driven decision-making falls within the third-party certification requirement.

Cybersecurity throughout product lifecycle. Cybersecurity controls apply throughout the product lifecycle. Documentation under the Machinery Regulation must be retained for up to 10 years. Vulnerability handling and security updates under the EU Cyber Resilience Act (Regulation (EU) 2024/2847) must be supported for a minimum of 5 years, with reporting obligations applying from 11 September 2026 and full applicability from 11 December 2027. For Physical AI, this extends conformity obligations well past the point of sale. A robot deployed in 2027 may carry security obligations into the early 2030s.

Digital declarations of conformity. The Regulation permits declarations of conformity in digital form, replacing paper-only documentation in many cases. This change, while operational rather than substantive, supports automated verification and audit by integrators and procurement bodies.

2.2 The EU AI Act — Regulation (EU) 2024/1689

The EU AI Act was adopted in June 2024 and entered phased implementation across 2025–2027. Where the Machinery Regulation governs the physical hazard layer, the AI Act governs the AI-specific obligations layered on top. For Physical AI, the two regulations operate in tandem.

High-risk AI systems. AI systems embedded in safety-critical machinery fall within the AI Act's high-risk classification. This triggers obligations beyond those in the Machinery Regulation: documentation of training data quality, transparency to users about AI involvement in safety functions, human-oversight provisions, and post-market monitoring of model behaviour.

Article 57 — Regulatory sandboxes. Article 57 requires each Member State to establish at least one AI regulatory sandbox by August 2, 2026. Sandboxes provide controlled environments in which compliance approaches can be tested before full market deployment. For Physical AI startups, sandboxes offer a path to engage with regulators early, reducing certification risk later. Member States are at varying stages of sandbox preparation as the Article 57 deadline approaches. Each Member State is required to establish at least one operational AI regulatory sandbox by that date.

Transparency, traceability, and human oversight. The AI Act requires that high-risk AI systems be designed for transparency to operators, that AI-driven decisions be traceable through audit logs, and that meaningful human oversight remain operationally feasible. For Physical AI, traceability is the operationally demanding requirement: systems must record, in tamper-resistant form, the decisions that produced safety-relevant physical outputs.

Fines. Non-compliance penalties under Article 99 follow a three-tier structure: up to €35 million or 7 percent of global annual turnover for prohibited practices under Article 5; up to €15 million or 3 percent for high-risk AI non-compliance under Articles 6 to 49 — the tier most relevant to Physical AI; and up to €7.5 million or 1 percent for supplying misleading information to authorities. The mid-tier (3 percent / €15 million) is the operative ceiling for safety-critical Physical AI manufacturers, since most Physical AI applications fall within the high-risk classification rather than the prohibited-practice category.

2.3 Harmonized Standards Landscape

EU regulations of this kind work through a layered model. The regulation states essential health and safety requirements. Harmonized standards specify how those requirements are met in technical terms. Conformity to a harmonized standard creates a presumption of conformity to the underlying regulation. For Physical AI, the harmonized-standards layer is partially populated.

Table 1. Status of standards relevant to Physical AI conformity assessment under Regulation (EU) 2023/1230 as of May 2026.

Standard	Scope	Status (May 2026)
IEC 61508	Functional safety of electrical/electronic/programmable electronic systems	Existing harmonized
ISO 13849-1:2023	Safety-related parts of control systems	Existing harmonized
ISO 13482:2014	Personal care robots	Existing harmonized
ISO 10218 series	Industrial robots (2025 update)	In process — harmonization expected late 2026
ISO/IEC TR 5469	Functional safety + AI	Emerging — technical report stage
ISO/IEC 22989	AI terminology	Existing

The gap is visible in the table. Existing functional-safety standards (IEC 61508, ISO 13849) were not drafted with AI/ML systems in mind. ISO/IEC TR 5469, which addresses functional safety in the presence of AI, remains at the technical-report stage. Until ISO/IEC TR 5469 graduates to a full standard and is harmonized under the Regulation, conformity assessment for AI-driven safety functions relies on adapting older standards to new system architectures. Notified bodies and manufacturers have begun to develop common interpretive practice. Convergent practice, not codified standard, is the operational reality for AI-specific conformity in 2026 and into 2027.

2.4 Notified Bodies and Certification Capacity

Conformity assessment under Regulation (EU) 2023/1230 is performed by notified bodies — designated organizations authorised to issue certificates within their accredited scopes. Major notified bodies active in machinery and Physical AI include TÜV SÜD [26], TÜV Rheinland, TÜV NORD, SGS [28], Pilz [27], BSI, and AFNOR.

Table 2. Major EU notified bodies with capability relevant to Physical AI conformity assessment under Regulation (EU) 2023/1230.

Notified Body	Country of Origin	Strengths Relevant to Physical AI
TÜV SÜD	Germany	Functional safety, automotive, industrial
TÜV Rheinland	Germany	Industrial machinery, robotics
TÜV NORD	Germany	Functional safety, certification programmes
SGS	Switzerland	Broad-scope industrial certification
Pilz	Germany	Safety control systems, robotics expertise

Notified Body	Country of Origin	Strengths Relevant to Physical AI
BSI	United Kingdom	Standards development and certification
AFNOR	France	French market access, certification

Capacity is constrained. As of May 2026, average backlog for new Physical AI certifications ranges from 6 to 18 months across the major bodies. Two factors drive the constraint. First, the volume of pre-2027 certification activity from companies seeking to be on the market under the new Regulation when it takes effect. Second, the time required for notified-body assessors to develop expertise in AI/ML conformity.

The capacity constraint is not academic. Companies that begin formal certification engagement in 2026 may not complete assessment before January 2027. Companies that delay engagement until late 2026 face material risk of post-deadline gaps in market authorisation.

2.5 Implication for Physical AI

Regulation (EU) 2023/1230 is the binding constraint for many Physical AI companies. Four factors compose the EU's dominant role: broad scope, with machinery as defined including most safety-critical Physical AI products; mandatory third-party certification for high-risk categories; layered AI Act obligations; and constrained certification capacity. Together these produce a single dominant compliance imperative for the EU market. PAS Level 3 — meaning documented hazard analysis, validated safety functions, and third-party certification capability — represents the minimum maturity for EU market entry in human-adjacent Physical AI applications by January 2027.

---

3 United States Regulatory Landscape

The United States Physical AI regulatory landscape contrasts with the EU's. Where the EU regulates Physical AI horizontally through the Machinery Regulation and AI Act, the United States regulates Physical AI sectorally through agency-specific frameworks. There is no single Physical AI statute. Coherence emerges from a voluntary federal backbone, sector-specific obligations, and an evolving state-level layer.

3.1 The Federal AI Policy Backbone

Two federal instruments anchor the United States approach.

NIST AI Risk Management Framework (AI RMF). NIST published AI RMF 1.0 in January 2023, with subsequent updates and profiles released through 2024–2026. The framework is voluntary. It is not a regulation. Adoption, however, is increasingly assumed in federal procurement, in agency-specific guidance, and in emerging state

legislation. The framework structures AI risk management around four functions: govern, map, measure, and manage. For Physical AI specifically, the framework's emphasis on socio-technical risk assessment, traceability, and human oversight aligns closely with the AI-specific requirements emerging in EU regulation.

White House AI Executive Order. The October 2023 Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, with its 2024–2026 implementation cascade, directs federal agencies to develop AI-specific guidance within their statutory authorities. The Order does not create new horizontal Physical AI authority; it directs existing agencies to use existing authority more pointedly. The downstream effect is sector-specific.

For Physical AI companies, the federal backbone provides a vocabulary (AI RMF) and a direction of travel (the Executive Order's cascade). It does not provide a single point of certification.

3.2 Sectoral Regulators

Sectoral agencies apply existing frameworks to AI-driven systems within their domains. Five agencies are most relevant.

Table 3. United States sectoral regulators most relevant to Physical AI.

Agency	Scope	Physical AI Pathway
OSHA	Workplace safety, including industrial robotics	Currently applies general industry standards (ANSI/RIA R15.06); AI-specific guidance emerging
FDA	Medical devices, including surgical and care robotics	Software as a Medical Device (SaMD) framework; Class II/III device pathway
NHTSA	Autonomous vehicles	Federal Automated Vehicles Policy guidance; state-level patchwork beneath
FAA	Unmanned aerial systems	Part 107 (small UAS); Part 108 (BVLOS, in development)
NRC	Nuclear (limited Physical AI applications today)	Existing regulatory framework

OSHA. Workplace robot safety in the United States currently operates under general industry standards. ANSI/RIA R15.06 [17] is the de facto baseline for industrial robot safety. OSHA enforcement uses the General Duty Clause where standards do not cover specific hazards. AI-specific OSHA guidance is in development but not yet finalised.

FDA. Medical and surgical robots fall under the Software as a Medical Device framework. AI-driven systems with material clinical decision impact are typically classified Class II or Class III, requiring 510(k) clearance or premarket approval. The FDA has issued an action plan for AI/ML-based medical devices addressing the change-control challenge: how to authorise systems whose behaviour evolves through learning. Predetermined Change Control Plans are emerging as the operational mechanism. Under a Predetermined Change Control Plan (per FDA AI/ML-Based Software as a Medical Device Action Plan and 2023 PCCP guidance), a manufacturer specifies in advance the categories of model updates that will be made post-clearance, the validation procedure for each, and the trigger conditions that would force a new clearance application. The mechanism allows for continuous learning systems within a regulated envelope. For Physical AI in clinical settings, this is the operational path forward.

NHTSA. Autonomous vehicle regulation in the United States is bifurcated between federal vehicle-design authority (NHTSA) and state operational authority. NHTSA issues Federal Automated Vehicles Policy guidance and operates voluntary safety self-assessment programmes. State activity — California DMV permitting, Arizona's testing-friendly stance, Texas's statutory framework — varies markedly. The compliance picture for AV deployment is therefore state-by-state, layered onto NHTSA federal expectations.

FAA. Unmanned aerial systems operate under Part 107 for small UAS in visual line-of-sight operations. Part 108, addressing beyond-visual-line-of-sight (BVLOS) operations, is in development. AI-driven flight control, detect-and-avoid systems, and autonomous mission planning fall within FAA airspace and certification authority.

NRC. Nuclear applications of Physical AI remain limited as of 2026. Where they exist, they fall within NRC's existing licensing and oversight framework.

3.3 State-Level Activity

State activity layers onto federal sectoral frameworks. Two states are bellwethers.

California. California regulates autonomous vehicles through DMV permits. The state has passed several AI-specific laws addressing transparency, deepfakes, and algorithmic decision-making. SB-1047, the most prominent recent attempt at horizontal AI safety regulation, was vetoed in September 2024. Alternative bills addressing similar concerns have emerged in 2025–2026.

New York. New York has advanced AI Bill of Rights-style legislation, focused initially on employment and consumer-facing decisions. Physical AI implications are indirect but accumulating.

States including Texas (which has advanced an autonomous-vehicle framework supportive of commercial testing), Massachusetts (which has pursued legislation on automated decision-making oversight), Illinois (which maintains long-standing biometric-privacy law that intersects with Physical AI sensing), and Colorado (whose AI Act establishes obligations on developers and deployers of high-risk AI systems) have advanced varied frameworks. The cumulative effect is a state-level regulatory mosaic that companies must navigate alongside federal sectoral rules.

3.4 The Regulatory Pluralism Challenge

The United States approach produces a distinct compliance posture. A Physical AI company entering the United States market does not face a single conformity assessment; it faces a sector-specific pathway shaped by the dominant sectoral regulator for its application, layered with state-specific obligations in the deployment states.

This pluralism creates two effects. The first is complexity: cross-sector products (an AI system used in both medical and industrial settings, for example) face multiple parallel pathways. The second is strategic flexibility: companies can sequence market entry by selecting jurisdictions and sectors where the regulatory pathway best matches their development stage.

The closest the United States has to a unified Physical AI framework is the NIST AI RMF. As an explicitly voluntary framework, AI RMF cannot bind. As a widely cited reference, AI RMF increasingly shapes what sectoral regulators expect. By 2027, AI RMF compatibility is likely to be a de facto baseline for federal procurement and a common reference point in sectoral agency guidance.

3.5 Implication for Physical AI

United States compliance is fragmented but converging. Companies entering the United States market need a sectoral strategy: identification of the dominant sectoral regulator, alignment with that regulator's existing framework, AI RMF compatibility as a cross-sectoral foundation, and state-specific compliance overlays in deployment jurisdictions. Cross-sector products require parallel pathways. Federal procurement is an additional convergence vector. Federal agencies, governed by Office of Management and Budget memoranda, increasingly cite the AI RMF in solicitations for AI-enabled systems. For Physical AI companies serving federal customers — Department of Defense, Department of Veterans Affairs, Department of Homeland Security — AI RMF compatibility is moving from optional to expected. The fragmentation is structural; convergence within sectors and across the AI RMF baseline is the operational direction of travel.

---

4 Japan Regulatory Landscape

Japan is among the world's largest Physical AI markets by deployed installed base. Its regulatory approach differs structurally from both the EU and the United States. Japan's pattern is industry-led, government-coordinated, and convergence-oriented in standards adoption.

4.1 METI and the Strategic Framework

The Ministry of Economy, Trade and Industry (METI) coordinates Japan's Physical AI policy. The FY2026 government budget includes ¥387 billion in Physical AI commitments across robotics, autonomous systems, and AI-enabled industrial transformation. This investment level signals strategic priority and provides the financial substrate for regulatory and standards activity.

Three documents anchor the framework:

Table 4. Anchor documents in Japan's Physical AI policy framework.

Document	Purpose
AI Strategy (multiple iterations 2019–2026)	National AI policy direction
Robotics Industry 5.0 framework	Sectoral robotics strategy
Society 5.0 vision	Cyber-physical integration policy

The framework is forward-looking and economic in posture. Where the EU emphasises market-access conditioning and the United States emphasises sectoral risk management, Japan's framework emphasises industrial competitiveness with safety as an integrated rather than primary lever.

4.2 Industry-Led Standards

Standards development in Japan flows through the Japanese Industrial Standards Committee (JISC), which participates in ISO and IEC processes, and through industry groups led by the Japan Robot Association (JARA). JARA's role is more substantive than that of equivalent bodies in many jurisdictions: industry self-regulation often produces de facto requirements before government regulation formalises them.

This pattern has consequences for international convergence. Japanese delegates to ISO and IEC technical committees frequently bring industry-developed practice into international standards processes, and Japanese industry adoption of finalised ISO/IEC standards tends to be rapid. The result is that Japanese practice and international standards co-evolve closely.

4.3 Sectoral Rules

Specific Physical AI rules in Japan operate at the sectoral level.

- **Industrial robots:** The Japanese Industrial Standard JIS B 8433 governs industrial robot safety. The standard has multiple parts addressing system design, operation, and integration. JIS B 8433 has been revised across decades and remains the operational baseline for industrial deployments in Japan. The standard tracks ISO 10218 closely, supporting interoperability with EU and international markets.
- **Autonomous vehicles:** The National Police Agency (NPA) issues operational rules for AV testing and limited deployment. Major manufacturer collaborations — particularly Toyota and Honda — operate under MLIT (Ministry of Land, Infrastructure, Transport and Tourism) guidance combined with NPA enforcement. MLIT's posture is permissive within designated operational design domains. Toyota's Woven City — officially launched September 25, 2025 in Susono City, Shizuoka Prefecture, with approximately 300 initial residents — and Honda's autonomous-mobility programmes are notable Japanese-domestic AV test sites.

- **Surgical robots:** The Pharmaceuticals and Medical Devices Agency (PMDA) approves medical and surgical robotics. The PMDA pathway parallels the FDA's SaMD framework in scope, with Japan-specific procedural requirements. Approval timelines are typically longer than FDA timelines but produce certificates accepted across the broader Asian market.
- **Care robots:** Care robotics is a flagship sector. METI coordinates the Robot Service Initiative, which provides a framework for safety, deployment guidelines, and procurement support across hospitals and care facilities. Japan's demographic profile makes care robotics a national-strategic priority, and government procurement programmes have created a domestic deployment scale unmatched in other jurisdictions.

4.4 The Japanese Pattern

Japan typically leads in Physical AI deployment and follows in regulatory specificity. Industry self-regulation through JARA and parallel JIS standards often precedes formal regulation. After observing implementation in Japan and abroad, Japanese regulators tend to converge on EU and United States frameworks where compatibility supports export markets.

The convergence is partial but persistent. JIS standards in industrial robotics align closely with ISO 10218. PMDA approval pathways increasingly parallel FDA SaMD. AV operational rules track international harmonisation efforts.

4.5 Implication for Physical AI

Japanese market access is achievable through industry engagement, JARA participation, and adoption of JIS-aligned international standards. Regulatory convergence with the EU and United States is partial but ongoing, and is generally favourable to companies that have already prepared for EU 2023/1230 or United States sectoral pathways. Japan is rarely the first market to require certification, but companies treating it as an afterthought miss the operational-deployment scale that the Japanese installed base supports.

---

5 Israel Regulatory Landscape

Israel is a small market with a high concentration of Physical AI startups. The regulatory posture reflects this asymmetry: domestic regulation is light, international compliance is the operational requirement, and recent IIA activity signals an emerging shift toward domestic certification capacity.

5.1 The IIA AI Strategy (April 2026)

The Israel Innovation Authority (IIA) published its AI Strategy in April 2026. The report covered 123 Israeli Physical AI companies and identified the safety certification gap as a strategic concern. A key recommendation, in the original Hebrew:

אוטונומיות" מערכות לבטיחות הסמכה מסלולי "פיתוח

[Development of certification tracks for autonomous system safety.]

The recommendation positions safety certification as a national-strategic capability rather than purely a market-access cost. The framing is consistent with the IIA's broader posture toward Physical AI as a strategic export sector for Israel: domestic certification capacity reduces dependence on foreign notified bodies and supports Israeli companies entering EU and United States markets.

The April 2026 report identifies several specific gaps: limited domestic notified-body capacity, absence of dedicated Physical AI safety certification programmes within the Standards Institution of Israel (SII), and the cost burden on early-stage companies of pursuing foreign certification. The recommendations propose addressing these gaps over the next several years.

The IIA report frames the certification gap in industrial-policy terms. Israeli Physical AI companies operate in sectors — industrial robotics, autonomous vehicles, agricultural automation, drone systems, medical robotics — where international export is the dominant revenue path. Foreign certification costs, often denominated in tens to hundreds of thousands of euros, are a material burden for early-stage companies. Domestic certification capacity, if developed, would lower those costs and accelerate time-to-market. The IIA's strategic framing treats certification as economic infrastructure, comparable to export credit or trade-mission support.

The SII has historically focused on adoption of international standards rather than development of domestic Physical AI safety certification. The April 2026 report signals that this posture is under review. Proposed expansion would build SII capability for AI-specific conformity assessment, with the medium-term aim of operating as a notified-body-equivalent for Israeli companies seeking EU market access.

5.2 Regulatory Sandbox Proposal

The IIA report and parallel discussions within the Israeli regulatory community include a proposal for a Physical AI regulatory sandbox. The proposal references two international precedents: Singapore's CETRAN (Centre of Excellence for Testing and Research of Autonomous Vehicles) [24] for AV-specific sandboxing, and the United Arab Emirates' Regulations Lab [25] for cross-sector regulatory experimentation.

A sandbox would offer Israeli Physical AI startups a controlled environment to test compliance approaches, develop documentation, and engage with regulators before full market deployment. The proposal is aligned with the EU AI Act's Article 57 sandbox model, providing potential interoperability advantages for Israeli companies entering the EU market.

As of May 2026, the sandbox proposal is in active discussion but not yet operational. The IIA's role would likely be coordinating; sectoral regulators (transport, health) would participate within their domains.

5.3 Sectoral Regulators

Israeli sectoral oversight follows a pattern familiar from larger jurisdictions, scaled to the country's size.

Table 5. Israeli sectoral regulators with Physical AI relevance.

Body	Domain	Posture
Standards Institution of Israel (SII)	Technical standards adoption	Adopts ISO/IEC standards; nascent dedicated Physical AI capability
Ministry of Health	Medical devices including surgical robots	Approval pathway parallels FDA / PMDA
Civil Aviation Authority (CAA)	Drones and aerial systems	Israeli drone-pilot licensing; airspace coordination
Ministry of Transport	Autonomous vehicle pilots	Permit-based AV testing

5.4 The Israeli Pattern

The Israeli pattern is shaped by market size. The domestic market is too small to support Physical AI startups at scale on its own. Most Israeli Physical AI companies prepare for EU compliance, United States compliance, or both, from the start. Domestic regulation tends to enable innovation by deferring to international standards and avoiding parallel domestic-only requirements.

The April 2026 IIA report represents a policy shift. The recognition that domestic certification capacity is itself a strategic asset signals that Israel's posture is moving from "defer to international" toward "enable international through domestic capacity."

5.5 Implication for Physical AI

Israeli companies face the EU 2023/1230 constraint regardless of domestic regulation. The IIA's April 2026 strategy signals that domestic certification infrastructure is a near-term policy priority but not yet an operational one. Israeli Physical AI companies should plan for international certification on conventional timelines while monitoring domestic developments that may, over the next 12 to 24 months, provide complementary domestic options.

---

6 Convergence Analysis

The four-jurisdiction survey above presents the regulatory landscape as four distinct legal regimes. The remainder of this paper develops the central thesis: that despite these distinctions, the four regimes converge on common substantive requirements during the period 2025–2027. This section presents the convergence evidence; Section 7 provides the operational implications.

6.1 The Four-Jurisdiction Convergence Pattern

Across the EU, United States, Japan, and Israel, four common requirements emerge in 2025–2027:

1. **Third-party certifiability** of safety-critical Physical AI components.
2. **AI-specific risk assessment** with documented methodology.
3. **Decision traceability** — auditable record of AI-driven decisions affecting physical outcomes.
4. **Incident reporting** to regulators within defined timeframes.

These requirements converge despite different legal frameworks, indicating shared underlying concerns rather than coordinated policy.

The presence of each requirement in each jurisdiction is summarised in Table 6.

Table 6. The four-jurisdiction convergence matrix. Each cell records the regulatory pathway through which the listed requirement applies in the listed jurisdiction. "Yes" indicates a horizontally applied legal requirement. "Sectoral" indicates application through agency-specific regulation. "Industry-led" indicates de facto requirement through standards or association rules.

Requirement	EU	United States	Japan	Israel
Third-party certifiability	Yes — Regulation (EU) 2023/1230	Sectoral (FDA for medical, NHTSA-aligned for AVs)	Industry-led, emerging through JIS / JARA	Following EU and United States; IIA proposing domestic tracks
AI-specific risk assessment	Yes — AI Act high-risk classification	Yes — NIST AI RMF (voluntary, increasingly cited)	Industry guidance via JARA; aligned with international standards	IIA emphasis (April 2026 strategy); SII forthcoming
Decision traceability	Yes — AI Act traceability obligations	Sectoral (FDA Predetermined Change Control Plans, NHTSA standing general orders)	Industry guidance	Following EU and United States

Requirement	EU	United States	Japan	Israel
Incident reporting	Yes — Regulation (EU) 2023/1230 + AI Act post-market monitoring	Sectoral (NHTSA standing general orders, FDA MDR, OSHA record-keeping)	Industry guidance	Following EU and United States

The matrix shows the convergence directly. Every jurisdiction contains every requirement. The pathway differs — the substance does not.

6.2 The Convergence Claim

The four jurisdictions converge on the same four requirements not because of coordinated policy but because of shared underlying concerns. Four causes appear most consistently in the policy record:

Table 7. Underlying concerns producing the four-jurisdiction convergence.

Cause	Mechanism
Physical harm from AI-driven systems	Direct safety concern; primary driver in EU and FDA-regulated US sectors
Public trust	Sustains political support for continued AI deployment
Market access pressure	Large buyers (industrial integrators, hospital systems, fleet operators) require compliance
Insurance industry pricing	Underwriters price unmitigated AI risk; certification reduces premiums

Each cause is general — none is jurisdiction-specific. Each has been documented in EU, United States, Japan, and Israel policy records. The convergence is the predictable downstream effect of these shared concerns operating through different legal cultures.

This causal account has implications for the durability of the convergence. Policy convergence driven by coordination (a treaty, a standards body, a single intergovernmental forum) can be reversed when coordination weakens. Concern-driven convergence is more durable: it persists as long as the underlying concerns persist. The concerns above show no signs of weakening. The convergence is robust and likely to deepen.

6.3 The 2025–2027 Window

The current period is the window during which convergent requirements are being codified into binding regulation. Figure 1 shows the timeline.

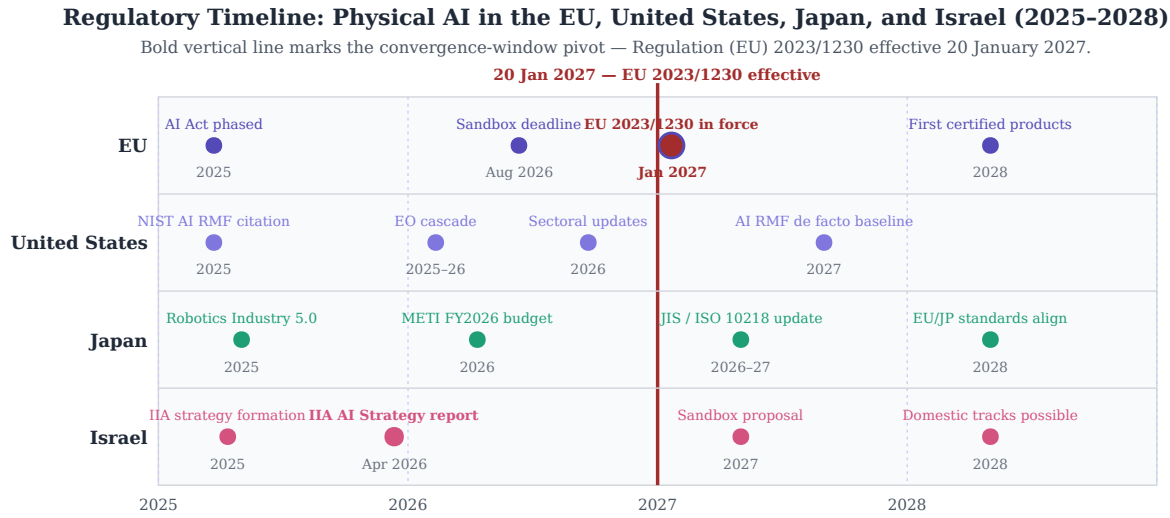

Figure 1. The 24-month compliance window runs from mid-2025 through mid-2027.

Source: synthesis of Regulation (EU) 2023/1230, Regulation (EU) 2024/1689, NIST AI RMF, METI policy documents, and IIA April 2026 AI Strategy report.

Figure 1. Figure 1 — Regulatory timeline across the European Union, United States, Japan, and Israel from 2025 through 2028, with milestones marked per jurisdiction and a vertical pivot at the effective date of EU Regulation 2023/1230 on January 20, 2027.

Figure 1. Regulatory timeline across the four jurisdictions, 2025–2028. Each lane shows jurisdiction-specific milestones. The bold vertical marker indicates the convergence-window pivot — Regulation (EU) 2023/1230 effective 20 January 2027.

For reference, the milestone data behind Figure 1 is summarised below.

Table 8. Milestone data underlying Figure 1.

Period	EU	United States	Japan	Israel
2025	AI Act phased implementation begins	NIST AI RMF widely cited; Executive Order cascade	Robotics Industry 5.0 implementation	IIA AI strategy formation
2026 H1	AI Act sandbox deadline (Aug 2, 2026)	Sectoral guidance updates; state-level activity intensifies	METI FY2026 budget activates	IIA April 2026 AI Strategy report
2026 H2	EU 2023/1230 final preparation; notified-body assessments	Continued sectoral cascade	Standards updates aligned with ISO 10218 revision	Sandbox proposal refinement

Period	EU	United States	Japan	Israel
2027	**Regulation (EU) 2023/1230 effective January 20**	Sectoral pathways increasingly converge on AI RMF baseline	Convergence with EU machinery posture begins	Domestic certification track proposals advance
2028	First wave of certified Physical AI products on EU market	State-level patchwork accumulates; federal action possible	Convergence with EU/US standards deepens	Sandbox operationalisation possible

By 2027, the four jurisdictions' substantive requirements will be substantially aligned, even where procedural pathways differ. Companies that fail to achieve PAS Level 3 or higher by then will face market-access challenges proportional to their exposure across these markets.

The window is open. It opened in 2025 when the AI Act and Regulation (EU) 2023/1230 entered phased preparation. It pivots in January 2027 with the Machinery Regulation's effective date. It closes — in the sense that companies entering the market thereafter without compliance face genuine regulatory exposure rather than operational friction — across 2027 and 2028. The 24 months between mid-2025 and mid-2027 are the operational window for companies to move from current state to compliance.

---

7 Compliance Roadmap

This section translates the convergence analysis into operational guidance. The roadmap is structured along three dimensions: compliance posture (where a company stands), company stage (development phase), and jurisdictional priority (sequencing across markets).

7.1 Compliance Posture

Compliance posture characterizes a company's progress toward jurisdictional regulatory requirements:

- **Pre-compliance:** company has not begun certification activity for the relevant jurisdiction.
- **Compliance-in-progress:** company has engaged a notified body or equivalent and has documented certification roadmap.
- **Compliance-achieved:** company has certification for relevant product category in the relevant jurisdiction.
- **Compliance-mature:** company has multi-cycle certification with renewal, including operational data sharing.

A company's compliance posture is jurisdiction-specific: the same company may be compliance-achieved in the EU and pre-compliance in Japan. Strategic compliance posture management matches investment to deployment plans.

7.2 Roadmap by Company Stage

Compliance posture and safety-engineering maturity scale with company stage. Pre-seed and seed-stage companies are typically too early for formal certification engagement; the work at this stage is hazard identification, basic safety architecture, and documentation hygiene that will support later certification. Series A companies should engage third-party assessors at advisory level. Series B is the typical stage at which formal certification engagement becomes both feasible and necessary for primary-market entry. Series C and later companies face cross-jurisdictional compliance as a routine operational function. Public and mature companies operate with multi-cycle certification and renewal across deployment jurisdictions.

A detailed staircase mapping company stage to specific maturity levels and compliance postures is the subject of a forthcoming companion paper on the Physical AI Safety maturity model. The present paper treats PAS Level 3 as the operative regulatory threshold for safety-critical applications targeting any of the four jurisdictions in the convergence window; deeper stage-by-stage mapping is left to that companion work.

7.3 Roadmap by Jurisdictional Priority

For most Physical AI companies, the priority sequencing is:

1. **Primary market** — most often the EU, given the binding January 2027 deadline and the breadth of Regulation (EU) 2023/1230's scope. A subset of companies whose primary market is United States medical (FDA), United States automotive (NHTSA), or Japan (METI sectoral) may sequence differently.
2. **United States** — sectoral; pathway depends on application (FDA, NHTSA, FAA, OSHA, NRC). AI RMF compatibility is the cross-sectoral foundation.
3. **Japan** — industry engagement through JARA; JIS-aligned standards adoption; PMDA pathway for medical robots.
4. **Israel** — domestic regulation often deferred to international compliance; IIA developments worth monitoring as 2027 and 2028 progress.

This sequencing is a starting heuristic, not a universal rule. Companies should adjust based on actual deployment plans, customer commitments, and capital position.

7.4 Cross-Jurisdictional Efficiency

Achieving compliance with Regulation (EU) 2023/1230 generally creates a foundation that meets approximately 70 to 80 percent of United States and Japanese requirements. The reasons are structural: EU 2023/1230 demands broad third-party assessment that overlaps with FDA SaMD, NHTSA documentation expectations, and JIS-aligned industrial robot standards. AI Act traceability and risk-assessment provi-

sions overlap with NIST AI RMF function areas. Figure 2 visualises both the convergence pattern and the cross-jurisdictional coverage data.

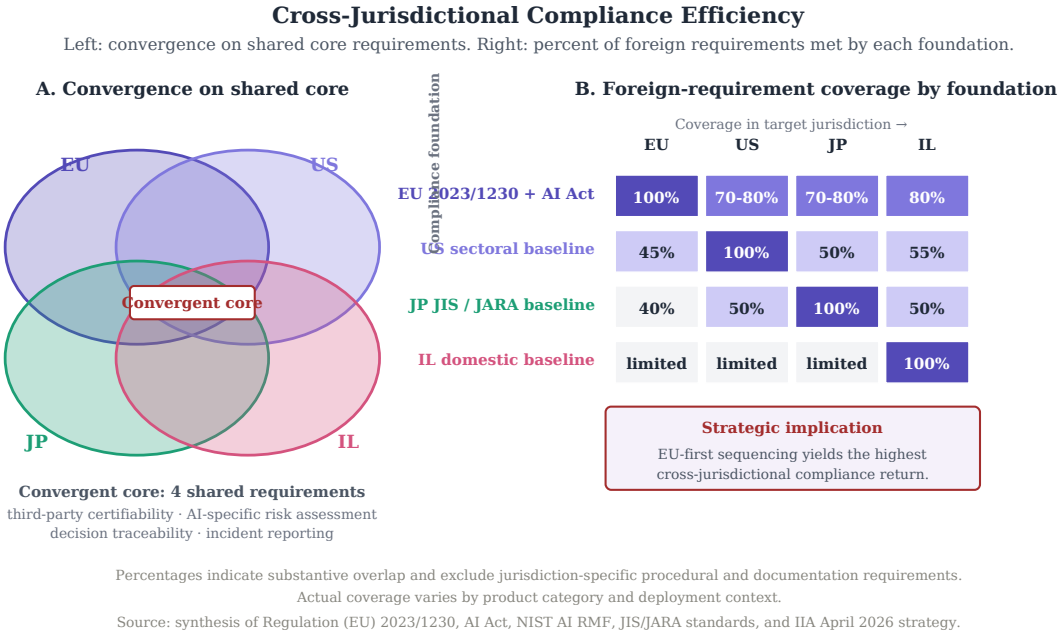

Figure 2. Figure 2 — Two-panel diagram. The left panel shows the EU, United States, Japan, and Israel regulatory regimes as four overlapping regions converging on a shared core of four common requirements. The right panel is a coverage heatmap. EU compliance covers approximately 70 to 80 percent of United States and Japanese requirements and approximately 80 percent of Israeli requirements. United States sectoral baseline covers approximately 45 percent of EU requirements. Japan JIS baseline covers approximately 40 percent of EU requirements. Israeli domestic compliance has limited overlap with foreign requirements. Strategic implication: EU-first sequencing produces the highest cross-jurisdictional return on compliance investment.

Figure 2. Cross-jurisdictional efficiency. Left panel: convergence on shared core requirements. Right panel: percent of foreign-jurisdiction requirements met by each compliance foundation. Percentages indicate substantive overlap and exclude jurisdiction-specific procedural and documentation requirements. Actual coverage varies by product category and deployment context.

For reference, the coverage data behind the right panel of Figure 2 is summarised below.

Table 9. Cross-jurisdictional coverage data underlying Figure 2.

Compliance Foundation	EU	United States (Sectoral)	Japan	Israel
Regulation (EU) 2023/1230 + AI Act	100% (by definition)	~70–80% (sector-dependent)	~70–80% (industrial robotics highest overlap)	~80% (Israeli companies often pre-aligned to EU)
United States sectoral baseline (e.g., FDA SaMD only)	~40–50%	100% (within sector)	~50%	~50–60%
Japan JIS / JARA baseline	~40%	~50%	100%	~50%
Israel domestic baseline	Limited overlap (small set of domestic-only requirements)	Limited	Limited	100%

Strategic implication: for most Physical AI companies operating across multiple jurisdictions, EU-first sequencing produces the highest cross-jurisdictional return on compliance investment. Companies whose primary deployment is United States-medical, United States-automotive, or Japan-industrial may have application-specific reasons to sequence differently, but should treat cross-jurisdictional efficiency analysis as a routine planning input.

7.5 Roadmap Application

The combined picture: compliance posture, company stage, and jurisdictional priority compose a planning matrix. A Series B company entering EU-medical in 2027 should be in compliance-in-progress (advanced) for EU as of mid-2026, with active third-party engagement. A pre-seed company building toward EU-industrial deployment in 2028 should be in pre-compliance with documented hazard-analysis work, planning the resource ramp to engage notified bodies in 2027. The matrix is a planning tool. It does not substitute for jurisdiction-specific legal counsel or technical engagement with notified bodies.

---

8 Conclusion

This paper has surveyed the regulatory landscape for Physical AI across four jurisdictions: the European Union, the United States, Japan, and Israel. Despite different legal traditions, four common requirements appear in each: third-party certifiability, AI-specific risk assessment, decision traceability, and incident reporting. The 2025–

2027 period is the window during which these requirements are being codified into binding regulation across the four jurisdictions.

The implications are concrete. For companies, the convergence creates a 24-month compliance window during which PAS Level 3 or higher must be achieved for safety-critical Physical AI applications. EU-first sequencing produces the highest cross-jurisdictional return on compliance investment for most companies. Notified-body capacity is constrained; engagement initiated late in the window faces real risk of post-deadline gaps. For regulators, the parallel emergence of common requirements suggests opportunities for international cooperation in conformity assessment, standards development, and incident-data sharing — opportunities that, if pursued, would lower compliance costs without lowering substantive safety standards. For investors, regulatory diligence should now include a compliance-posture assessment for portfolio companies operating in regulated Physical AI markets. Broader policy and governance literature relevant to these implications is collected in references [29][30][31][32].

Companies should treat the window as a strategic priority rather than a future concern. The Regulatory Convergence Window is open. It will not stay open.

References

[1] European Union, "Regulation (EU) 2023/1230 of the European Parliament and of the Council of 14 June 2023 on Machinery (Machinery Regulation)," *Official Journal of the European Union*, OJ L 165, 2023.

[2] European Union, "Regulation (EU) 2024/1689 of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (AI Act)," *Official Journal of the European Union*, 2024.

[3] European Union, "Directive 2006/42/EC of the European Parliament and of the Council of 17 May 2006 on Machinery," *Official Journal of the European Union*, OJ L 157, 2006. (Predecessor to Regulation (EU) 2023/1230; replaced 20 January 2027.)

[4] United States Office of the President, "Executive Order 14110 of October 30, 2023: Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," *Federal Register*, vol. 88, p. 75191, 2023.

[5] National Institute of Standards and Technology, "AI Risk Management Framework (AI RMF 1.0)," *NIST AI 100-1*, 2023.

[6] U.S. National Highway Traffic Safety Administration, "Federal Automated Vehicles Policy and subsequent automated-vehicle guidance," U.S. Department of Transportation.

[7] U.S. Food and Drug Administration, "Software as a Medical Device (SaMD) Guidance and AI/ML-Based Software as a Medical Device Action Plan."

[8] Israel Innovation Authority, "AI Strategy Report," 2026. Including survey of 123 Israeli Physical AI companies and recommendations on certification tracks for autonomous system safety.

[9] Ministry of Economy, Trade and Industry (METI), Japan, “Robotics Industry 5.0 framework and related AI Strategy iterations,” 2019–2026.

[10] International Electrotechnical Commission, “IEC 61508 — Functional Safety of Electrical/Electronic/Programmable Electronic Safety-Related Systems.”

[11] International Organization for Standardization, “ISO 13849-1:2023 — Safety of Machinery — Safety-Related Parts of Control Systems — Part 1: General Principles for Design,” 2023.

[12] International Organization for Standardization, “ISO 13482:2014 — Robots and Robotic Devices — Safety Requirements for Personal Care Robots,” 2014.

[13] International Organization for Standardization, “ISO 10218 series — Robots and Robotic Devices — Safety Requirements for Industrial Robots,” 2025 update in process; harmonisation under Regulation (EU) 2023/1230 expected.

[14] International Organization for Standardization and International Electrotechnical Commission, “ISO/IEC TR 5469 — Artificial Intelligence — Functional Safety and AI Systems,” technical report stage as of May 2026.

[15] International Organization for Standardization and International Electrotechnical Commission, “ISO/IEC 22989 — Information Technology — Artificial Intelligence — Artificial Intelligence Concepts and Terminology.”

[16] Japanese Industrial Standards Committee, “JIS B 8433 — Industrial Robots — Safety Requirements.”

[17] American National Standards Institute and Robotic Industries Association, “ANSI/RIA R15.06 — Industrial Robots and Robot Systems — Safety Requirements.”

[18] M. Melchior, “Physical AI Safety: A Definitional Framework,” forthcoming companion paper, 2026.

[19] M. Melchior, “The PAS-MM Maturity Model for Physical AI Safety,” forthcoming companion paper, 2026.

[20] M. Melchior, “Software Safety Ceiling in Physical AI Systems,” forthcoming companion paper, 2026.

[21] M. Melchior, “Empirical Survey of Physical AI Industry Safety Posture,” forthcoming companion paper, 2026.

[22] M. Melchior, “Common-Cause Failures in Physical AI Systems,” forthcoming companion paper, 2026.

[23] M. Melchior, “A Certification Framework for Physical AI Safety,” forthcoming companion paper, 2026.

[24] Centre of Excellence for Testing and Research of Autonomous Vehicles (CETRAN), Singapore, “Operational framework for autonomous-vehicle testing.”

[25] United Arab Emirates Regulations Lab, “Cross-sector regulatory experimentation framework.”

[26] TÜV SÜD, “Public guidance on functional safety and AI conformity assessment.”

[27] Pilz GmbH, “Safety control systems and machinery regulation guidance.”

[28] SGS, “Conformity assessment publications relevant to Physical AI.”

[29] Brookings Institution, “Selected policy papers on AI regulation.”

[30] Stanford Institute for Human-Centered AI (HAI), “Selected policy papers on AI governance.”

[31] International Association of Privacy Professionals (IAPP), “Regulatory landscape reports.”

[32] World Economic Forum, “Reports on AI governance.”

Mati Melchior · Independent Physical AI Safety Researcher mati@physical-ai-safety.com · physical-ai-safety.com P7 of an 8-paper series on Physical AI Safety.