# White Paper on Embodied AI: A Pathway for Future Embodied AI Standards Development

Bridging research, deployment, and globally relevant standards for safe, scalable embodied AI.

Jointly developed by Singapore Standards Council and ASTM International in collaboration with Singapore's Agency for Science, Technology and Research (A*STAR) and the National Robotics Programme (NRP).

Published June 2026 by ASTM International and Enterprise Singapore.

enterprisesg.gov.sg

go.astm.org

# Contents


01 About This White Paper . . . . . . . . . . 3

02 Global Developments in EAI . . . . . . . . . . 5

03 Identifying Roadblocks Hindering Mass-Scale Adoption . . . . . . . . . . 9

Industry Consultations in Singapore — the Standards Gap in EAI . . . . . . . . . 9

04 Deep Dive — The Five Standardisation Gaps . . . . . . . . . . 10

Gap 1: Data Governance & Privacy Compliance . . . . . . . . . . 10
Gap 2: Learning from AI Safety & Reliability in Enterprise Systems . . . . . . . . 12
Gap 3: Workplace Safety & Risk Management . . . . . . . . . . 13
Gap 4: Interoperability & Middleware . . . . . . . . . . 15
Gap 5: Quality & Performance Certification . . . . . . . . . . 18

05 EAI Standards & Research Initiatives . . . . . . . . . . 21

Singapore Initiatives & Use Cases . . . . . . . . . . 21
Other Global Efforts to Address Standardisation Gaps . . . . . . . . . . 24

06 A Call for Collaborators . . . . . . . . . . 25

07 References . . . . . . . . . . 26

# 01 About This White Paper

The convergence of ageing populations and persistent labour shortages is driving an urgent search for automation that can operate safely and reliably in dynamic physical environments. Conventional, software-based AI systems — optimised for digital tasks — are increasingly misaligned with these demands. Embodied artificial intelligence (EAI) has emerged as a promising response, integrating AI with physical systems such as robots, autonomous vehicles, and intelligent machines so that they can perceive, decide, and act in the real world.

EAI offers transformative potential across sectors. It can automate complex real-world tasks, raise productivity in manufacturing, logistics, and healthcare, and enable new services — from robotic arms that adapt to changing production lines, to autonomous systems that optimise supply chains, to exoskeletons that personalise patient rehabilitation. By taking over dangerous, dirty, and dull work, EAI can also improve workforce safety and free human workers for higher-value roles. As a cross-cutting capability, it is likely to underpin new business models such as Robotics-as-a-Service and platform-based ecosystems that drive innovation and diversification across industries.

Rapid advances in machine learning, simulation, sensing, and robotic platforms have catalysed a surge of EAI research and early deployments. Unlike traditional industrial robotics, which typically relies on pre-programmed behaviours and structured environments, EAI systems learn through continuous sensorimotor interaction. Their perception, cognition, and action capabilities are closely coupled, allowing them to adapt to changing and unpredictable conditions in real time. However, this same adaptivity creates new safety and governance challenges. EAI systems must operate within physical constraints, manage uncertainty, and make real-time decisions in environments shared with people, infrastructure, and legacy machinery.

This white paper, jointly developed by the Singapore Standards Council (SSC), and ASTM International, and in collaboration with the Agency for Science, Technology and Research (A*STAR) and the National Robotics Programme (NRP) , examines how standards can enable safe, scalable, and economically meaningful deployment of EAI. The paper:

- Identifies five critical standardisation gaps that risk becoming strategic bottlenecks for EAI adoption: data governance and privacy; AI safety and reliability in enterprise systems; workplace safety and risk management; interoperability and middleware; and quality, performance, and procurement-grade certification.
- Maps global research and policy efforts that are beginning to address these gaps, highlighting where emerging practice can inform standards and assurance frameworks.
- Uses Singapore's deployment-oriented ecosystem as a "living lab" to surface practical lessons from real-world EAI deployments and testbeds, as well as

# 01 About This White Paper

to illustrate how deployment evidence can be channelled into research-to-standards pathways.
- Outlines opportunities for collaboration to develop deployment-ready test methods, guidance documents, and standards that can be harmonised internationally.

The paper is intended for policymakers and regulators, national standards bodies, robot manufacturers and integrators, infrastructure and facility owners, and large end users responsible for deploying EAI systems in workplaces and public spaces. For these stakeholders, it provides both a strategic overview of emerging EAI ecosystems and a practical agenda for bridging the gap between research prototypes, pilot deployments, and globally relevant standards.

The shift from disembodied AI to physically deployed, embodied systems is reshaping the basis of technological and economic competition. Intelligence is moving from static screens into robots, autonomous systems, industrial machines, and other physical platforms, creating new demands on safety, assurance, interoperability, and governance frameworks. Long-term market projections suggest that general-purpose humanoids and broader EAI systems could become a multi-trillion-dollar market by mid-century, as autonomous agents become increasingly embedded in daily life and industrial operations.

In response, many countries and regions are developing national strategies, research programmes, and regulatory regimes to support EAI ecosystems. The next section surveys these global developments to identify reference models, complementary strengths, and collaboration opportunities that can inform a coordinated international approach to EAI standards.


Certis PeTER (Patrol & Traffic Enforcement Robot) developed for Certis CISCO Security in Singapore.

1 With contributions from other Singapore agencies such as the Infocomm Media Development Authority (IMDA) and the Workplace Safety and Health Council (WSHC).

# 02 Global Developments in EAI

The development of EAI is an international endeavour, with countries around the world contributing complementary strengths in research, industrial deployment, standards development, and governance. Together, these efforts offer important reference points for the field's broader development, where joint exploration and shared learning can create mutual value. The global transition toward EAI is not simply a continuation of the current AI wave. It marks a deeper structural shift: from intelligence that operates primarily in digital environments to intelligence that can perceive, decide, and act in the physical world.

As AI moves from static screens into kinetic agents such as robots, autonomous systems, industrial machines, and other physical platforms, the basis of economic competition is also changing. Long-term projections suggest that the total addressable market for general-purpose humanoids and EAI could eventually reach US$5 trillion by 2050 (Morgan Stanley, 2025), as autonomous agents become as ubiquitous as personal vehicles. This has catalysed countries to develop national strategies to accelerate innovation and deployment of these technologies.

## CHINA

China has placed EAI at the centre of its industrial and technology priorities, with national plans identifying it alongside biomanufacturing, quantum technologies, and advanced communications as a focus area for the coming decade. This creates a rich landscape of activity that other countries can engage with through research exchanges, standards harmonisation, and supply-chain integration.

At the government level, China has established dedicated structures to support standardisation, including the Ministry of Industry and Information Technology's Humanoid Robot and Embodied Intelligence Standardisation Technical Committee, formed in late 2025, and has released national standards covering multiple stages of the humanoid robot lifecycle. Public investment is being deployed at multiple levels, with regional governments — including those in Beijing, Shanghai, Shenzhen, Guangdong, Zhejiang, and Hubei — supporting targeted funding programmes and industrial initiatives. Many of these efforts follow a pilot-and-scale approach, providing useful reference cases for how localised experimentation can be translated into broader rollouts.

Academic and research activity has expanded rapidly, with institutions such as Tsinghua University, Tongji University, the Harbin Institute of Technology, and the Chinese Academy of Sciences playing central roles. Dedicated data and training platforms — including initiatives at the Beijing EAI Robotics Innovation Centre and the National Local Joint Humanoid Robot Innovation Centre in Shanghai — illustrate how shared data infrastructure can support iterative model development.

On the industry side, China's installed base of industrial robots provides extensive operational data that is contributing to rapid advances in models and systems. Firms such as AgiBot and Unitree have increased the visibility of humanoid robotics, while large technology and industrial companies — including Huawei (through its

# 02 Global Developments in EAI

EAI Industry Innovation Centre) and Xiaomi — are pursuing integrated approaches that combine hardware, software platforms, and ecosystem partnerships. The scale of deployment and resulting data flywheel offer important reference points for global standards conversations on validation, lifecycle assurance, and continuous learning.

## EUROPEAN UNION

The contribution of the European Union (EU) to the global EAI landscape centres on regulatory architecture and safety-focused engineering practices. Several overlapping frameworks now govern AI-enabled robotics in Europe — including the AI Act (fully applicable from August 2026), the Machinery Regulation (EU) 2023/1230 (applicable from January 2027), and the revised Product Liability Directive — creating a comprehensive reference model for jurisdictions developing their own governance approaches.

The European Commission has allocated EUR 174 million for robotics-related projects under the Horizon Europe programme (2023–2025), with Germany's Robotics Institute initiative investing an additional EUR 350 million through 2026. Germany remains Europe's largest robot market, and European research institutions such as ETH Zurich, TU Munich, EPFL, and Imperial College produce world-class research in robot learning, soft robotics, and human-robot interaction. The European robotics industry is characterised by strong SME participation and a long tradition of focus on safety and collaborative operation, reflected in the maturity of the cobot category.

The E.U.'s experience offers valuable insights on harmonising safety, liability, and conformity assessment across multiple jurisdictions. The principles underpinning the E.U. regulatory triad — risk classification, conformity assessment, and lifecycle accountability — align closely with several of the gaps identified in this paper, including lifecycle governance, post-deployment learning, and assurance frameworks for adaptive systems. Joint dialogue and engagement through international standards bodies offer practical avenues for cross-pollination of approaches.

## JAPAN

Japan brings decades of industrial robotics leadership to the EAI conversation, combined with a strong policy focus on demographic challenges of global concern — particularly elder care, healthcare, and service robotics.

At the government and research level, Japan invests approximately USD 930 million each year in robotics R&D. The Moonshot Research and Development Programme, launched in 2020 with USD 440 million in funding through 2050, specifically targets AI and autonomous learning robots, alongside elder-care and human-machine collaboration applications. Universities including Waseda, Osaka, and the Tokyo Institute of Technology remain internationally recognised research institutions in humanoid, soft, and haptic robotics — areas where Japan's academic tradition has

# 02 Global Developments in EAI

maintained decades of continuity and where joint research collaboration could be particularly productive.

Industrially, Japan is the world's leading manufacturer of industrial robots, with FANUC, Yaskawa, and Kawasaki bringing decades of manufacturing expertise and strong global customer relationships. A notable strategic shift is underway as these traditional manufacturers integrate modern AI capabilities. FANUC has announced a partnership with NVIDIA for physical AI and released a ROS 2 driver, while Yaskawa's MOTOMAN NEXT-NHC 10DE learns packing motions by imitation from human demonstrations.

### SOUTH KOREA

South Korea has built one of the world's most deployment-dense robotics ecosystems through consistent, multi-decade policy commitment, beginning with the Intelligent Robots Development and Distribution Promotion Act of 2008. The current 4th Basic Plan on Intelligent Robots (2024–2028) commits approximately USD 2.2 billion to manufacturing and service robotics under the K-Robot Economy banner. Hyundai's acquisition of Boston Dynamics in 2021 catalysed a broader national R&D wave, with Samsung, LG Electronics, and KAIST University deepening their robotics programmes. Samsung and Hyundai are also testing humanoid platforms within their own manufacturing environments, creating an integrated loop between industrial demand and R&D investment.

### UNITED KINGDOM

The United Kingdom maintains strong academic capabilities through institutions such as Imperial College London, University College London, and the University of Edinburgh, alongside a regulatory stance that balances alignment with EU frameworks and flexibility in selected EAI domains.

The U.K.'s academic strengths in robot learning, autonomous systems, and AI safety provide natural opportunities for research collaboration with Singaporean institutions, particularly in areas such as formal verification, AI assurance, and human-robot interaction. The U.K. has also developed practical regulatory experience in domains such as autonomous vehicles and drones, where its approach to staged trials and operator certification offers useful comparative material for other ecosystems' AV and autonomous-systems pathways..

### UNITED STATES

The EAI ecosystem in the U.S. is shaped by a combination of federal research funding, industry-led deployment, and a strong tradition of standards and measurement leadership. Together, these elements offer useful points of reference for other ecosystems.

At the research and innovation level, the National Science Foundation (NSF) supports early-stage robotics and EAI work through programmes such as the Directorate for Technology, Innovation and Partnerships (TIP); Engineering

# 02 Global Developments in EAI

Research Centers (ERCs); and the Convergence Accelerator. The Defense Advanced Research Projects Agency (DARPA) continues to advance embodied autonomy through challenge-driven programmes that have historically catalysed benchmarking and system-level integration approaches now used across the industry. U.S. universities — including Carnegie Mellon, MIT, Stanford, UC Berkeley, and the University of Washington — are consistently leading contributors to foundation-model robotics, robot learning, and vision-language-action (VLA) research. The Open X-Embodiment consortium, a collaboration of 34 laboratories that has compiled approximately one million robot trajectories from 22 robot types, represents one of the most significant shared research datasets in the field.

On the standards and measurement side, the National Institute of Standards and Technology (NIST) has developed widely adopted test methods for robots, particularly for response robots and autonomous systems, and its AI Risk Management Framework (AI RMF) provides voluntary guidance increasingly relevant to embodied systems. Standards development organisations — including ASTM International and IEEE — are actively defining safety, interoperability, and performance standards. ASTM's robotics committees, such as the committee on robotics, automation, and autonomous systems (F45), are developing test methods that address real-world deployment challenges, including public-area robotics and emerging humanoid systems.

Industry deployment is led by companies such as Boston Dynamics, Agility Robotics, Figure AI, Physical Intelligence, and Apptronik, with NVIDIA's Isaac platform serving as a widely used infrastructure layer for training robot foundation models. Deployments at companies including BMW, GXO, and Amazon are generating operational data and integration lessons that help shape what responsible, production-grade EAI deployment requires.

Despite these diverse regional strengths, a common set of friction points has emerged. **These roadblocks are not unique to any one nation; they represent systemic gaps that hinder mass-scale adoption worldwide.**

# 03 Identifying Roadblocks Hindering Mass-Scale Adoption

## INDUSTRY CONSULTATIONS IN SINGAPORE — THE STANDARDS GAP IN EAI

Despite the strong global push to advance EAI, the harder challenge may not be innovation itself, but enabling these systems to be adopted safely, consistently, and at scale. Many of the conditions needed for deployment remain underdeveloped. In particular, there is still limited coordination around the standards required to guide how EAI systems are tested, integrated, procured, governed, and deployed in real-world environments.

This is critical because EAI systems will not operate in clean or controlled settings. They will be deployed in dynamic workplaces alongside workers, existing machinery, legacy systems, and sector-specific operating procedures. Without clearer standards on safety, reliability, data governance, interoperability, certification, and risk management, enterprises may find it difficult to assess whether a solution is ready for deployment. Technology providers may also face fragmented requirements across customers and sectors, slowing commercialisation and making it harder for solutions to scale beyond pilots.

This was a key motivation behind Enterprise Singapore's July 2025 EAI standardisation workshop, organised in partnership with local stakeholders including NRP, the Building and Construction Authority (BCA), and Jurong Town Corporation (JTC). The workshop, where ASTM International was invited as a keynote speaker, brought together more than 140 industry representatives to assess real-world demand for EAI, and to surface the standards gaps, deployment challenges, and problem statements that could stand in the way of wider adoption. The key standardisation gaps identified through the workshop are as follows:

1. **Data Governance & Privacy Compliance**
2. **Learning from AI Safety and Reliability in Enterprise Systems**
3. **Workplace Safety & Risk Management**
4. **Interoperability & Middleware Standards**
5. **Quality Certification & Procurement Standards**

The subsequent section delves into these gaps in detail, exploring why they matter and where robust standards can provide essential guidance.

# 04 Deep Dive — The Five Standardisation Gaps

1

## GAP 1: DATA GOVERNANCE & PRIVACY COMPLIANCE

EAI systems are mobile, sensor-rich, and physically present in the real world, constantly collecting visual, spatial, acoustic, and operational data as they interact with people, infrastructure, and other machines. This data is generated on the ground — in workplaces and public spaces — and shared across multiple organisations, creating a unique and complex set of data governance and privacy challenges that existing data protection frameworks, while providing important foundations, were not designed to handle. Industry stakeholders have flagged uncertainty around data management as a growing barrier to wider EAI adoption.

These challenges are amplified by the nature of EAI. In EAI systems, perception, decision-making, and action are closely linked. Poor data handling may therefore have direct safety implications, not just compliance implications. For example, incomplete or poorly governed data could affect how a robot understands its environment, responds to people, or learns from past operations. EAI data is also highly context-dependent: the same sensor stream may have different levels of sensitivity depending on where the system is operating, who is nearby, and what task is being performed.

EAI systems may also use operational data to improve their models over time. This raises important questions about how data should be stored, reused, secured, and audited. As deployments scale, organisations with access to larger and better datasets may be able to improve their systems faster, creating a "data flywheel" advantage. Compounding this, EAI deployments span multiple stakeholders — each with differing data handling practices and levels of accountability — creating fragmented ownership and heightened governance risk across shared data streams.

The practical consequences of these gaps are significant. Organisations may grow hesitant to share operational data needed for model improvement and safety validation, procurement cycles could be prolonged by bespoke data clauses, and inconsistent privacy controls across vendors risk undermining buyer confidence. These issues could make it increasingly difficult to scale pilots into production environments, particularly in human-occupied spaces, and leave legal accountability unclear when incidents or near-misses occur.

There is therefore a need for shared data governance approaches across the EAI ecosystem. Data governance should not be viewed only as a compliance issue — it is a foundation for safe, trusted, and commercially viable deployment. At this stage of EAI maturity, stakeholders are not calling for rigid or overly prescriptive rules. Instead, there is a need for standards that provide clarity, consistency, and common expectations across the industry. These standards should help to:

- Establish shared terminology and role definitions;
- Define minimum governance artefacts and evidence expectations;
- Enable interoperability of governance practices across platforms and vendors;
- Support consistent assurance without constraining innovation.

# 04 Deep Dive — The Five Standardisation Gaps

Table 1. Key governance gaps observed in EAI deployments — why they matter for embodied systems, and where standards could add value.

| GAP AREA | WHY THIS MATTERS FOR EAI | WHAT IS BREAKING TODAY | WHERE STANDARDS COULD HELP |
|---|---|---|---|
| **Data ownership and rights in multi-party deployments** | EAI data is generated through physical interaction in shared spaces involving multiple stakeholders. | Data sharing for learning and safety is stalled due to unclear rights and responsibilities. | Common role taxonomy and baseline responsibility models. |
| **Data classification for robotics modalities** | EAI produces heterogeneous, continuous sensor data tied to physical context. | Inconsistent controls and weak auditability across deployments. | Minimal classification schema and metadata conventions. |
| **Context-aware privacy handling** | Sensitivity of data changes dynamically based on location, task, and human proximity. | Static privacy rules are over- or under-applied. | Shared context labelling and policy-switching concepts. |
| **Privacy-by-design in onboard sensing** | Physical sensors cannot be “turned off” without affecting capability. | Buyers cannot compare or verify vendor privacy claims. | Baseline control catalogues and verification artefacts. |
| **Data minimisation vs safety and forensics** | Safety learning and incident analysis require data retention. | Organisations either over-retain or under-retain data. | Lifecycle-based retention concepts and trigger-based guidance. |
| **Access control across stakeholders** | Remote access to robots and logs is common in EAI. | Ad-hoc access models increase risk and slow response. | Standard access tiers and delegation patterns. |
| **Cross-border and hybrid architectures** | EAI systems often combine edge autonomy with cloud oversight. | Data residency concerns delay deployments. | Reference governance architectures and residency tagging. |
| **Data lineage for training reuse** | Operational data feeds back into model updates. | Compliance and accountability are hard to demonstrate. | Minimum lineage and provenance metadata expectations. |
| **De-identification of multi-modal data** | Robotics data is harder to anonymise than static datasets. | Organisations avoid data pooling altogether. | Modality-specific de-identification guidance. |
| **Auditability of privacy-relevant actions** | EAI systems act in real time with limited oversight. | Weak assurance during audits or incidents. | Standard audit log fields and reporting structures. |
| **Incident data handling** | Physical incidents trigger both safety and privacy concerns. | No shared playbook for data use during incidents. | Incident data governance workflows and boundaries. |
| **Procurement-ready governance evidence** | Buyers lack a consistent way to assess governance maturity. | Procurement cycles lengthen unnecessarily. | Common governance evidence pack structure. |

# 04 Deep Dive — The Five Standardisation Gaps

2

## GAP 2: LEARNING FROM AI SAFETY & RELIABILITY IN ENTERPRISE SYSTEMS

The dynamic, environment-driven nature of EAI systems introduces a distinct set of compounding challenges for AI safety and reliability. EAI systems may exhibit emergent behaviours that were never explicitly programmed or anticipated, making it difficult to verify all possible operating conditions through conventional certification processes. Unlike traditional machines, EAI systems may continue to evolve through updates, retraining, or adaptation to new environments. This raises questions about how safety certification can be meaningfully maintained over time. Risk assessment is further complicated by the complex interactions between perception algorithms, planning models, and environmental inputs, which can give rise to hazards that are difficult to anticipate at the design stage. The probabilistic nature of machine learning models also undermines conventional verification techniques that assume deterministic, explainable software behaviour. Finally, the integration of AI with physical machinery introduces cyber-physical safety risks, where software failures, data corruption, or cybersecurity attacks can directly influence physical behaviour — requiring safety assurance approaches that address functional safety and cybersecurity in tandem.

Many existing industrial safety standards were developed for deterministic machines whose behaviour can be fully specified, tested, and verified before deployment. For these reasons, many existing robotics and machinery safety standards are insufficient to fully address the risks introduced by EAI. Current benchmarks, such as the revised ISO 10218:2025 standard, still primarily address static, pre-programmed robots. These standards do not yet account for vision-language-action (VLA) based systems that continually update and learn post-deployment. Furthermore, traditional product liability chains break when a robot autonomously modifies its own parameters. The lack of an international liability framework for post-deployment learning remains a critical barrier to enterprise adoption and insurance, justifying a shift toward “Lifecycle Governance” and robust “Maturity Models.”

Addressing these challenges will require new ways of ensuring that EAI systems remain safe over time, including tools to monitor system behaviour in real time, methods to validate AI models, and sector-specific standardised environments for testing autonomous systems. Standards have a meaningful role to play in making this possible. Purpose-built EAI safety frameworks can provide structured ways to assess the risks from autonomous decision-making and adaptive behaviours — covering how reliably a robot perceives its environment, how it transitions safely between tasks, and how its motion planning is validated. Given that EAI systems continue to evolve after deployment, standards can also set out best practices for ongoing monitoring, documentation, post-update validation, and audit trails. Finally, standards can introduce standardised testing and benchmarking environments, such as simulations and datasets, to evaluate system robustness in diverse conditions, enabling consistent performance comparisons and evidence-based safety claims.

# 04 Deep Dive — The Five Standardisation Gaps

Table 2. Safety and reliability gaps for adaptive embodied systems, and where standards development could provide value.

| GAP AREA | WHY THIS MATTERS FOR EAI | WHAT IS BREAKING TODAY | WHERE STANDARDS CAN HELP |
|---|---|---|---|
| **Lack of safety assurance frameworks for EAI** | Adaptive, learning-based behaviours introduce uncertainties that existing workplace safety standards — designed for predictable machines — cannot assess, exposing workers to autonomous decision-making risks. | Existing standards struggle to certify systems whose behaviour is emergent and not fully predictable. | Evaluate AI perception reliability, define safe-state transitions for anomalies, ensure redundant sensing/control, and validate AI motion planning. |
| **Lack of standardised testing methods in realistic environments** | EAI performance is highly sensitive to environmental context, making static testing insufficient. | Inability to consistently compare system performance across vendors or deployments due to lack of common protocols. | Define metrics for perception, navigation, and manipulation, with stress tests for uncertainty, HRI validation protocols, and benchmark datasets for realistic conditions. |
| **Lifecycle governance for AI systems** | Continuous learning and adaptation mean EAI systems are not static, requiring ongoing validation beyond initial deployment. | Safety assurance typically ends at deployment, leaving post-update behaviour ungoverned. | Frameworks for monitoring performance and maintaining safety across the operational lifecycle — documentation, post-update validation, degradation monitoring, and audit trails. |
| **Integration of functional safety and cybersecurity** | Cyber vulnerabilities can influence physical behaviour and translate into physical safety hazards. | Separate treatment of functional safety and cybersecurity fails to address combined cyber-physical risks. | Robust cybersecurity for safety-critical systems through secure updates, anomaly detection, and integrated cyber-physical risk assessment. |

3

## GAP 3: WORKPLACE SAFETY & RISK MANAGEMENT

EAI is increasingly being adopted in workplaces to address labour shortages, improve productivity, and enhance worker safety. In Singapore, EAI is being leveraged in healthcare, where robots like DigiRob — developed by NCS and A*STAR — help automate tasks in nursing homes, such as taking vital signs and delivering medication.

In other countries such as China, EAI is heavily integrated into manufacturing, with a Xiaomi humanoid robot, powered by a VLA model, performing precision assembly tasks such as self-tapping nut installation on the company's Beijing EV production line. Similarly, in the U.S., EAI is enhancing industrial and construction operations: Boston Dynamics' Spot, a versatile quadruped robot, supports site inspections, data collection, and safety monitoring across construction, energy, and infrastructure sectors.

EAI presents an inherent tension in the context of workplace safety. While it is deployed to protect workers by taking over hazardous and physically demanding

# 04 Deep Dive — The Five Standardisation Gaps

tasks, the introduction of autonomous and physically capable machines into human-occupied spaces gives rise to a new category of safety risks. The autonomy and adaptability that make EAI valuable also render its behaviour harder to predict and control, meaning that a poorly governed or inadequately tested system can itself become a source of harm. The risks extend beyond conventional industrial hazards to include uncertainty in system behaviour, perception failures, autonomous decision-making errors, and complex human–robot interactions. Existing industrial safety frameworks developed for conventional automation systems cannot fully address these challenges.

As EAI systems become more autonomous, workplaces must contend with two distinct but related dimensions of human-machine safety.

**Physical incidents arising from shared spaces**
Robots behaving unexpectedly — through sudden starts or stops, speed changes, or movements outside programmed paths — can result in workers being struck by, or caught between, moving machines. Such incidents may stem from sensor failures, software errors, poor spatial layout, or inadequate demarcation of robot zones. Near misses and unplanned contact should be treated as early warning signs and prompt a review of existing controls.

**Cognitive safety**
As workers transition from direct machine operation to supervisory oversight, they are increasingly required to monitor multiple systems simultaneously, interpret complex feedback, and intervene when anomalies arise. Poorly designed interfaces and unpredictable robot behaviour can elevate cognitive workload and erode situational awareness, leading to slower decision-making and delayed responses to hazards. Workers may also develop automation bias — over-trusting systems and failing to intervene when necessary — or conversely lose confidence and override automated processes in ways that introduce new risks.

Standards cannot eliminate all risks associated with EAI in the workplace, but they play a critical role in defining acceptable safety practices, establishing common evaluation methods, and setting guardrails that keep systems operating within manageable risk boundaries. One priority area is human-robot interaction. Existing collaborative robot guidelines were not designed to account for the higher levels of autonomy and adaptive behaviour characteristics of EAI systems. Standards could address these gaps by covering intuitive human-machine interface design, standardised signals for communicating robot status and intent, safety-rated sensors for human detection, guidelines for safe distances and zones, and tiered risk assessment frameworks calibrated to varying levels of hazard.

Singapore's Workplace Safety and Health Council (WSHC, 2026) has recommended best practices for safer human-robot interaction, covering both physical safeguards and organisational measures:

# 04 Deep Dive — The Five Standardisation Gaps

Table 3. WHSC recommended controls for safer human–robot interaction.

| ENGINEERING CONTROLS | ADMINISTRATIVE CONTROLS |
|---|---|
| Integrate proximity and collision sensors (e.g., cameras, LiDAR, ultrasonic, infra-red) to detect movement and prevent contact injuries. Provide physical barriers or safe zones where tasks demand separation between humans and robots. Install and regularly check emergency-stop (E-stop) and manual override mechanisms for quick response during abnormal operations. Implement automatic speed or power reduction when humans enter shared spaces. Design EAI systems to be user-friendly, reliable, and inherently safe for the task environment, using visual indicators, auditory alerts, and standardised operational signals to improve situational awareness. | Conduct thorough risk assessments (RA) and job hazard analyses focusing on human–robot interaction across installation, operation, and maintenance phases. Provide EAI-specific WSH training for workers and supervisors, including safe start-up, shutdown, and manual intervention procedures. Establish clear communication protocols and visual cues (e.g., safety distances, cautionary signs). Brief workers on the scope and limitations of each E-stop — whether it acts on a zone, subsystem, or the full system — and display this clearly at each E-stop point. Review incidents, near misses, feedback, and safety performance data regularly, and encourage a safety culture where any team member can stop unsafe operations and report hazards without fear of reprisal. |

## 4

## GAP 4: INTEROPERABILITY & MIDDLEWARE

Interoperability and middleware form the connective tissue of EAI systems. While interoperability has long been discussed in robotics, its importance is significantly amplified in EAI, where intelligent behaviour emerges from coordinated perception, decision-making, and action across embodied agents, software services, and physical infrastructure. In this context, middleware is not merely a technical integration layer; it becomes a system-level enabler of safety, scalability, and trust. Without robust interoperability foundations, EAI deployments become brittle — intelligent behaviours cannot be reliably transferred across sites, fleets cannot be expanded without bespoke engineering, and safety assurance becomes fragmented.

As EAI deployments expand into dense, shared environments such as hospitals, transport hubs, and public spaces, robots, software platforms, and physical infrastructure must communicate and coordinate seamlessly. The proliferation of proprietary robotic solutions has created an “Interoperability Paradox” where operational deadlocks arise in shared environments, and industry stakeholders consistently identify this as one of the most persistent barriers to scaling EAI, even where individual robot capabilities are technically mature. Interoperability challenges are more acute in EAI than in conventional robotics for several reasons:

– **EAI systems operate as ecosystems, not standalone machines.** Intelligent behaviour depends on real-time coordination across multiple robots, fleet managers, perception services, digital twins, and site infrastructure. Multimodal perception models require the real-time exchange of state and intent to maintain semantic consistency across heterogeneous hardware, while next-generation orchestration protocols such as Zenoh have emerged to address the discovery inefficiencies of traditional middleware in high-density urban environments (Open Source Robotics Foundation, 2026).

# 04 Deep Dive — The Five Standardisation Gaps

- **Context-awareness is distributed,** requiring the real-time exchange of state, intent, and environment information across systems.
- **Safety and autonomy are emergent properties,** shaped by interactions between robots, humans, and infrastructure rather than by any single component. Safety requirements must be stratified by sector; while isolated factory cells prioritise efficiency under ISO 10218, human-rich environments like hospitals demand stricter interaction safety protocols as defined in ISO 13482 (ISO, 2025).
- **Learning and adaptation rely on shared interfaces, as operational data,** events, and feedback loops flow through middleware layers. Frameworks like FailSafe-VLM demonstrate that cross-robot action generalisation depends on middleware's ability to provide a "reasoning" layer that adapts to unpredictable environmental changes.

Several structural weaknesses are compounding interoperability challenges at the ecosystem level. The absence of coordinated interoperability baselines means that even where relevant standards exist, they are inconsistently adopted, causing system architectures to diverge and integration work to be duplicated — a problem historically illustrated by coordination failures between robots from different vendors in healthcare settings prior to frameworks like Open-RMF. Procurement practices compound this further, as robots are frequently purchased as individual assets without explicit interoperability requirements, resulting in incompatible fleets and costly custom integrations. The middleware ecosystem remains fragmented, with developers facing inconsistent upstream interfaces and performance bottlenecks arising from high-frequency, large-model inference (Huang et al, 2024). Successful pilots frequently stall when moving to production due to bespoke APIs and the absence of conformance testing — a challenge that initiatives such as the Mass Robotics Interoperability Standard are beginning to address in environments like the Punggol Digital District. Underlying all of this is limited visibility on how interoperability frameworks and middleware stacks will be maintained and governed over time, which dampens industry confidence and long-term investment.

Stakeholders are looking for practical, deployment-oriented interoperability standards rather than prescriptive ones that dictate system architecture. The priority is clarity around interface expectations across robots, middleware, and infrastructure, with enough flexibility to accommodate local deployment realities, support conformance and assurance, and enable repeatable integration across sites and vendors. To address fragmentation, several strategic principles can guide standards and middleware development:

- **Vendor-agnostic orchestration layers:** Systems can adopt Open-RMF to ensure diverse fleets can coordinate path planning and infrastructure access without being locked into proprietary silos (Open Source Robotics Alliance, 2024).
- **Transition towards deterministic control:** Hybrid architectures that bridge high-level reasoning with industrial-grade programmable logic controllers allow robots

# 04 Deep Dive — The Five Standardisation Gaps

to retain the adaptability of EAI while meeting the rigorous precision required for manufacturing (Hyundai Motor Group, 2026).

- **Multi-modal data generation and interactive deployment:** Standards must enable "intelligent apprentice" robots to learn from limited human demonstrations and adapt to specific operator preferences (CNRS@CREATE, 2026).
- **Workforce harmonisation:** Alongside technical interoperability, learning outcomes for middleware platforms must be standardised across education and industry to ensure a consistent talent pipeline.
- **Decoupling of cognitive reasoning and physical execution:** Through standardised protocols — Agent2Agent (A2A) for horizontal peer-to-peer task negotiation (Google, 2025) and the Model Context Protocol (MCP) for orchestrating connections between agents and their external tools (Anthropic, 2024).

Table 4. Key interoperability and middleware gaps, and where standards development can provide immediate value.

| GAP AREA | WHY THIS MATTERS FOR EAI | WHAT IS BREAKING TODAY | WHERE STANDARDS CAN HELP |
|---|---|---|---|
| **Absence of a shared interoperability baseline** | EAI systems depend on coordinated multi-agent and infrastructure interactions. | Architectures diverge across deployments; integration is re-engineered each time. | Define common interoperability layers and scope boundaries. |
| **Inconsistent middleware interfaces** | Intelligent behaviour relies on reliable exchange of state, intent, and events. | APIs and semantics differ across vendors and sites. | Shared interface definitions and event semantics (Orchestrating Embodied Systems, 2026). |
| **Weak procurement signalling** | Interoperability is rarely enforced at purchase time. | Fleets become incompatible by design. | Referenceable interoperability requirements for procurement. |
| **Lack of conformance pathways** | EAI safety and reliability depend on predictable behaviour. | "Interoperable" claims cannot be verified. | Conformance concepts and testable criteria. |
| **Fragmented robot–infrastructure integration** | EAI frequently operates within buildings and shared spaces. | Bespoke lift, door, and BMS integrations. | Standardised robot–infrastructure interaction models — e.g. interfacing with lifts and automated doors (Singapore Standards, 2025). |
| **Limited portability of autonomy logic** | Autonomy stacks depend on middleware consistency. | Behaviour does not transfer cleanly across sites. | Stable abstractions between autonomy and hardware. |
| **Unclear cybersecurity alignment** | Middleware is a shared attack surface. | Security controls vary widely. | Alignment with established information-security controls. |
| **No clear stewardship model** | EAI ecosystems evolve continuously. | Uncertainty over long-term maintenance. | Defined roles for upkeep, evolution, and alignment. |

# 04 Deep Dive — The Five Standardisation Gaps

5

## GAP 5: QUALITY & PERFORMANCE CERTIFICATION

As EAI systems move from controlled pilots into mission-critical deployments, the industry faces a significant certification gap. Existing verification and validation (V&V) methodologies (ISO, 2025) were designed for deterministic, pre-programmed systems and are less suited to EAI's adaptive, AI-driven behaviour in unstructured environments. There is no globally recognised test regime for validating that behaviours trained in simulation will transfer safely to the real world. Statistical evidence shows that an AI policy achieving 95% accuracy in a controlled lab environment often drops to 60% when deployed in the field due to this "sim-to-real gap."

EAI certification must account for the close link between perception, decision-making, and physical action. Testing frameworks should therefore assess not only whether a robot completes a task, but also how well it reasons, adapts, and responds to changing conditions (SMU, 2025). They must also account for the gap between simulation and real-world deployment. A robot that performs well in simulation may behave differently when exposed to real-world factors such as friction, uneven surfaces, sensor noise, lighting changes, or human movement (NIST, 2025).

EAI systems face three main categories of vulnerability. First, external stressors such as poor lighting, adverse weather, crowded spaces, or unexpected obstacles can weaken perception and reduce performance, especially in dense urban or public environments (NIST, 2025). Second, internal degradation such as hardware wear, sensor drift, or actuator fatigue can cause performance to decline over time, reinforcing the need for materials, components, and monitoring systems that can withstand physical stressors (Fan et al., 2026). Third, risks can arise at the interface between perception and reasoning: an EAI system may misinterpret an unfamiliar environment, make an incorrect inference, or generate a physically unsafe action plan. These risks, including AI hallucinations in physical contexts, are not yet fully addressed by existing certification frameworks (NIST, 2025).

Certification frameworks will therefore need to test EAI systems for more than task completion. They should assess robustness, adaptability, reasoning quality, and safe behaviour under changing real-world conditions. Systems should also be able to detect and respond safely to unexpected physical events, ranging from accidental interruptions to deliberate cyber-physical attacks. Without standardised ways to measure cognitive quality and performance resilience, stakeholders have no reliable basis for comparing systems or quantifying the impact of AI improvements. This leads to several systemic friction points:

- **Procurement uncertainty and the "trust gap":** Users face uncertainty in identifying certified systems for high-risk deployments. Public road trials for autonomous shuttles, for example, require safety operators until they meet updated cybersecurity and update-management standards (NRP, 2025).
- **Liability and responsibility blind spots:** It remains unclear whether certification absolves vendors from responsibility when an autonomous system causes real-world harm, necessitating new legal frameworks for performance-based liability.

# 04 Deep Dive — The Five Standardisation Gaps

- **Absence of risk classification:** There is no shared taxonomy to classify EAI risks across sectors, leading to a "diagnostic void" for managers attempting to quantify AI improvements in human-rich environments (NIST, 2025).
- **Future navigation reliability:** Standard systems often fail in GPS-denied environments. A potential solution could be "quantum verification" via portable cold-atom inertial sensors to ensure resilient positioning (National Quantum Strategy, 2024).
- **The failure of "front-loaded" certification:** One-time certifications are insufficient for continuous-learning systems whose safety profiles change daily, requiring ongoing re-validation protocols (NIST, 2025).
- **The lack of "physical common sense":** EAI systems can produce mathematically optimal plans that are physically perilous, necessitating benchmarks like BEHAVIOR-1K that evaluate grounding in physical common sense (Li et al., 2024).

At this stage of EAI development, the goal of standards is to provide a structured path toward trustworthy autonomy. This means establishing clear boundaries between what manufacturers are responsible for and what operators are obligated to ensure, and shifting safety frameworks from simple mechanical safeguards to more sophisticated behavioural guardrails. As EAI systems are increasingly designed for collaborative environments, certifying for social and psychological safety becomes as important as preventing physical collisions (ISO, 2025). The following guiding principles hold for EAI certification:

- **Continuous multi-level certification:** Verification must occur at the algorithm, module, and system levels, extending across the entire AI lifecycle rather than being a one-time "gatekeeper" event (NIST, 2025).
- **Use-case-specific envelopes:** To mitigate performance drift, certification should be bounded to specific operational envelopes, such as hospital wards or outdoor loading docks, rather than general-purpose claims.
- **Fundamental behavioural verification:** Standards must deploy formal methods, such as the process analysis toolkit (PAT), to mathematically prove that autonomous behaviours satisfy strict safety and liveness properties in concurrent, real-time systems (Liu, 2026).
- **Differentiated application standards:** Standards must account for functional variance, distinguishing between requirements for stationary robotic-arm motion planning and mobile AMR path planning.
- **External "guardian" systems:** Safe autonomy requires independent safety monitors, such as control barrier functions, that sit outside the AI's neural network to ensure the system functions within verified specifications (Xie, 2025).

# 04 Deep Dive — The Five Standardisation Gaps

Table 5. Key certification gaps and where standards development can provide immediate value.

| GAP AREA | WHY THIS MATTERS FOR EAI | WHAT IS BREAKING TODAY | WHERE STANDARDS CAN HELP |
|---|---|---|---|
| Deterministic vs probabilistic regulatory mismatch | Current industrial standards rely on binary logic, while EAI operates on probabilistic weights. | Forcing probabilistic intelligence into deterministic regulatory checkboxes. | Define deterministic “Safety Shields” (e.g., Signal Temporal Logic compliance) that override AI in safety-critical moments (Poskitt et al., 2024). |
| Risk classification frameworks | EAI operates across a spectrum of risk, from isolated factories to public spaces. | Buyers cannot distinguish between “hobby-grade” and “industrial-safe” systems. | Establish a 5-level “Maturity Model” from ad-hoc testing to formal mathematical verification (NIST, 2025). |
| Productivity & performance benchmarking | Economic viability depends on objective evidence of productivity gains. | Performance metrics are siloed and vendor-specific. | Standardise test methods for Task Completion Time (TCT) and Cognitive Workload (CWL) reduction using benchmarks like BEHAVIOR-1K (Li et al., 2024). |
| Continuous V&V protocols | External risks cause performance drift that bypasses one-time certifications. | Validated systems become unsafe post-deployment due to fatigue or drift. | Standardise runtime monitors and “black box” recorders for lifecycle-long transparency and accountability (NIST, 2025). |
| Sim-to-real validation | Real-world testing is expensive and dangerous. | Over-reliance on simulation results that fail under real-world friction and gravity. | Standardise high-fidelity “digital twin” protocols and reality-mirroring benchmarks to mitigate drift. |
| Infrastructure interoperability | Diverse fleets must coexist in shared smart buildings. | Costly bespoke integrations for lifts and doorways. | Adopt SS 713:2025 and TR 130:2025 for standardised robot-to-building data exchange (SSC, 2025). |

# 05 EAI Standards & Research Initiatives

Identifying these gaps is only the first step. To move from identifying bottlenecks to clearing them, governments and organisations around the world — including Singapore — have leveraged their unique ecosystems to pioneer technical solutions that can potentially serve as real-world blueprints for these critical gaps.

## SINGAPORE INITIATIVES & USE CASES

Singapore's EAI landscape is being shaped by a combination of 1) national robotics ecosystem-building; 2) deployment-oriented interoperability and infrastructure integration efforts; and 3) broader national AI research and governance programmes that increasingly intersect with embodied, agentic, and safety-critical applications. Singapore has announced investment to strengthen public AI research capabilities and talent, framed as supporting broader ambitions under National AI Strategy 2.0 and its Research, Innovation & Enterprise 2030 Plan. Singapore's NRP coordinates across public agencies, industry, and research to translate robotics capabilities into real-world deployments, with growing emphasis on middleware, interoperability, and site-readiness as adoption scales. Supporting initiatives such as RoboNexus and RoboCluster extend this further by accelerating enterprise development and connecting partners across sectors.

Singapore is also collaborating with international players on EAI research. The CNRS–Singapore international project "Embodied AI" was launched on 1 January 2026, involving Singapore agencies including NRF and NRP (CNRS@CREATE, 2026). Singapore has also referenced CNRS participation through its CREATE presence to strengthen EAI research collaboration. Additionally, NVIDIA's Singapore research hub, launched in May 2026, strengthens Singapore's role as a regional node for EAI and next-generation AI infrastructure. Through collaborations with local universities, research institutions, industry, and government agencies, the hub supports joint research, talent development, and testbed-based validation within Singapore's broader AI ecosystem.

Beyond research, Enterprise Singapore and the Singapore Economic Development Board (EDB) are advancing a dual mandate: catalysing industry-led EAI use cases while attracting and growing EAI startups in Singapore. By targeting real automation pain points in sectors such as maritime, aviation, and the built environment, these pilots can create demand anchors for adoption, provide reference deployments for both local and global EAI startups, and strengthen Singapore's value proposition as a launchpad for EAI solutions. In parallel, these deployments generate practical evidence to inform standards development.

Collectively, these initiatives indicate a shift from isolated robotics deployments toward ecosystem-level enablement for EAI, where interoperable middleware and infrastructure integration, assurance and governance primitives, and scalable talent and enterprise adoption pathways are increasingly treated as first-order constraints. This provides a timely backdrop for standards-oriented discussion: as EAI scales

# 05 EAI Standards & Research Initiatives

across shared environments and multi-party operational contexts, consistent baselines for interoperability, data handling, and assurance become central to reducing friction and enabling responsible adoption.

To effectively use Singapore as a "living lab" and reference model that grounds global standards in reality, the following Singapore-specific use cases — mapped to the five gaps — provide relevant validation.

Table 6. Singapore-specific use cases mapped to the five standardisation gaps.

| GAP | SINGAPORE USE CASES |
|---|---|
| **Data Governance & Privacy Compliance** | Singapore is pioneering privacy-enhancing technologies (PETs) and accountability frameworks to govern the continuous multimodal data streams generated by EAI agents in public spaces.<br><br>**A*STAR FedPro — decentralising data for privacy.** To mitigate the privacy risks of centralised data storage, A*STAR's "FedPro" platform uses Federated Learning so multiple robotic agents across institutions can train models locally on the edge. Only learned model parameters are shared with a central aggregator, ensuring raw sensitive data never leaves the local device — also addressing the "data flywheel" bottleneck by letting firms pool learned models without sharing proprietary raw data. |
| **Learning from AI Safety & Reliability in Enterprise Systems** | National AI R&D and governance initiatives materially affect EAI readiness where autonomy, safety assurance, and accountability are required.<br><br>**AI Verify Foundation & IMDA — Global AI Assurance Sandbox.** A testing ground where companies building Generative AI applications collaborate with specialist testers against IMDA's Starter Kit for Testing of LLM-Based Applications. With 13 prequalified testing firms, an analogous EAI testing sandbox could contribute to industry-applicable EAI testing standards. (AI Verify Foundation, 2026)<br><br>**Model AI Governance Framework for Agentic AI.** The world's first governance framework explicitly targeting agentic AI, mandating "risk bounding" and defining "human accountability" checkpoints before irreversible actions — a potential reference model for EAI deployments. (IMDA, 2026)<br><br>**Centre of Excellence for Testing & Research of Autonomous Vehicles NTU (CETRAN) & TR 68.** Singapore's dedicated AV testing facility validates vehicles against TR 68 (behaviour, functional safety, cybersecurity, data formats), as seen with ST Engineering's Sentosa shuttles, which retain a human safety driver until the system meets TR 68's cybersecurity and software-update requirements. (SSC, 2021) |

# 05 EAI Standards & Research Initiatives

| GAP | SINGAPORE USE CASES |
| --- | --- |
| **Workplace Safety & Risk Management** | Singapore implements hard regulatory and safety frameworks for high-risk EAI deployments in public and industrial spaces.<br><br>**Autonomous vehicle testing at Punggol.** Autonomous shuttles by WeRide and Grab keep human safety operators on board until vehicles comply with Singapore's TR 68 series — a case study in phased, standards-driven safety. (LTA, 2025)<br><br>**HTX XPLORE & H2RC — public safety benchmarking.** HTX's XPLORE facility tests how officers perform cognitively and physically alongside AI under stress, while the Home Team Humanoid Robotics Centre (H2RC) aims to set global benchmarks for humanoids in dangerous civil-defence operations. |
| **Interoperability & Middleware Standards** | Singapore has developed SS 713:2025 to govern robot-lift and doorway data exchange, and TR 120:2025 for premises owners, AMR suppliers, and integrators managing multiple fleets in shared spaces. IMDA's AMR × Digital Leaders initiative aims to help 500 enterprises adopt and integrate AMRs with building infrastructure.<br><br>**Changi General Hospital — the genesis of Open-RMF.** In 2018, robots from different manufacturers paralysed each other in shared hallways for lack of a common protocol. This drove the creation of RMF, now "Open-RMF," a global open-source standard managed by the Open-Source Robotics Alliance — solving the Interoperability Paradox and providing connective tissue for global smart cities. (A*STAR, 2025)<br><br>**Mapletree Business City — commercial-scale orchestration.** Singapore's first large-scale commercial RMF-based middleware deployment, spanning 268,600 sqm, lets AMRs handling cleaning, surveillance, and food delivery interoperate across floors without human lift operation.<br><br>**The ELEVATE Sandbox — hardware-in-the-loop conformance.** ELEVATE @ BCA Braddell Campus is Singapore's national sandbox for RMF, providing physical and virtual environments to validate robot–lift–building interaction using Open-RMF with reference to SS 713. |
| **Quality Certification & Procurement Standards** | These use cases address certifying probabilistic EAI models, validating continuous learning, and ensuring simulation-trained behaviours survive real-world physics.<br><br>**A*STAR, ARTC & Siemens - SSAM Catalyst.** The SSAM Catalyst sandbox by A*STAR and Siemens accelerates deployment readiness by serving as an advanced validation environment for industrial solutions. Using digital twins and simulation tools, it tests and refines manufacturing workflows prior to physical rollout, drastically shortening testing cycles and reducing capital risk. This rigorous process ensures all pilot technologies achieve structural consistency and align with international standards. |

# 05 EAI Standards & Research Initiatives

These local successes demonstrate that the technical hurdles of EAI are solvable. However, for these to be internationally relevant, they must be viewed through a wider lens of international alignment.

## OTHER GLOBAL EFFORTS TO ADDRESS STANDARDISATION GAPS

**Data Governance & Privacy Compliance**
**United States — Open X-Embodiment Consortium.** Led by U.S. institutions, including UC Berkeley, across 34 laboratories, Open X-Embodiment created one of the world's largest shared robotics datasets — nearly one million robot trajectories across 22 platforms. It addresses fragmented robotics data by enabling cross-platform learning while raising governance questions around provenance, ownership, privacy, and reuse, serving as a foundational example of how shared data ecosystems accelerate EAI while increasing pressure for standardised data governance and lineage frameworks.

**Learning from AI Safety & Reliability in Enterprise Systems**
**China — MIIT Humanoid Robot & Embodied Intelligence Standardisation Technical Committee.** Established in late 2025 to develop lifecycle standards for humanoid robots and EAI, focusing on safety, reliability, dataset governance, operational maintenance, and continuous evaluation. Under ISO/TC 299, China is also leading ISO/WD 26264-1 on humanoid robot datasets — one of the most comprehensive national attempts to build enterprise-grade assurance frameworks for adaptive embodied systems.

**Workplace Safety & Risk Management**
**South Korea — Hyundai & Boston Dynamics industrial humanoid deployments.** Hyundai and Boston Dynamics are testing humanoid robots and advanced mobile systems inside active manufacturing environments, exploring how autonomous robots safely coexist with workers around material handling, logistics, and inspection — highlighting the importance of HRI standards, cognitive workload management, and real-time safety assurance.

**Interoperability & Middleware Standards**
**Germany — Siemens industrial interoperability and digital-twin ecosystems.** Siemens has advanced interoperable industrial automation through open communication architectures, digital twins, and integrated orchestration layers connecting robots, machines, sensors, and enterprise systems — a reference model for how EAI systems may integrate into highly regulated industrial environments using standardised interfaces and middleware.

**Safety, Governance, Quality Certification & Procurement Standards**
**European Union & United Kingdom.** The U.K. is an active participant in ISO and IEC, and its work on AI safety — including the AI Safety Institute — provides a reference point for EAI governance. The European Commission has established one of the world's most comprehensive governance architectures for AI-enabled robotics through the AI Act, Machinery Regulation, and revised Product Liability Directive, creating structured pathways for conformity assessment, lifecycle accountability, risk classification, and post-market monitoring of intelligent systems.

# 06 A Call for Collaborators

SSC and ASTM International share a common interest in advancing EAI standardisation that is grounded in real-world deployment experience and designed to be practically useful for industry. Singapore's deployment-driven ecosystem, public-private coordination mechanisms, and sectoral testbeds provide a rich source of operational evidence, while ASTM International brings established robotics test methods through the F45 committee and its new committee on AI in manufacturing (F50). Launched in June 2026, F50 will provide a dedicated home for translating these gap areas into deployment-ready test methods and standards. ASTM's center-of-excellence model, which pairs pre-standardization research with standards development, offers a proven template for the research-to-standards pathway described here.

Early discussions between SSC and ASTM are expected to focus on dataset readiness and data governance for EAI systems. As Singapore develops its national efforts to build its EAI ecosystem, there may be opportunities to share practical experience and nascent frameworks on EAI data that could inform broader pre-standardisation discussions.

The standardisation gaps identified here are not unique to any single country or deployment context. They reflect challenges that the global EAI community is navigating collectively. Now is the time for industry, researchers, agencies, and standards bodies to align on the right questions and build the evidence needed to address these gaps.

**Sharing knowledge and building common ground**
Creating spaces for knowledge exchange allows those working on EAI deployment, research, and regulation to surface shared challenges and exchange emerging practices — which is often where the most useful pre-standardisation insights originate.

**Turning experience into standards-ready outputs**
Industry players and research institutions are well placed to translate deployment experience and targeted investigation into guidance documents, technical frameworks, and shared benchmarks that can anchor standards development.

**Testing ideas in real-world contexts**
Sector-specific pilots and cross-border testbed collaborations generate the concrete, reproducible evidence that standards development requires, while surfacing the variation in conditions that internationally relevant standards must accommodate.

**Contributing to the global standards conversation**
International partners and standards bodies are warmly welcomed to contribute and share deployment experience and regulatory perspectives, as well as explore where co-participation in international standards development could help translate this collective work into globally coherent outcomes.

**GET INVOLVED**
Organisations interested in contributing to this work are encouraged to reach out to the Singapore Standards Council or ASTM International to explore how they can participate

# 07 References


1. Agency for Science, Technology and Research (2025). Singapore Champions Robotics Adoption And Talent Development With New Open-source Initiatives
2. AI Verify Foundation (2026). Global AI Assurance Sandbox.
3. Anthropic. (2024). Introducing the Model Context Protocol.
4. Huang et al. (2024). DaDu-Corki: Algorithm-Architecture Co-Design for Embodied AI-powered Robotic Manipulation. arXiv: 2407.04292.
5. CNRS@CREATE. (2026). Embodied AI for human-robot collaborative task execution.
6. Fan, S., Shi, A., & Lee, C. (2026). Soft Responsive Materials Enhance Humanoid Safety. Human-Centred Robotics Lab, National University of Singapore.
7. Google. (2025). Agent2Agent (A2A): Horizontal communication and peer-to-peer negotiation for autonomous fleets. Linux Foundation.
8. Infocomm Media Development Authority (2026) Model AI Governance Framework for Agentic AI.
9. ISO. (2025). ISO 10218-1:2025: Robotics — Safety requirements — Part 1: Industrial robots.
10. ISO. (2014). ISO 13482:2014: Robots and robotic devices — Safety requirements for personal care robots.
11. ISO (2016) ISO 15066:2016: Robots And Robotic Devices - Collaborative Robots.
12. Land Transport Authority (2025). New Autonomous Shuttle Services to be Progressively Deployed in Punggol
13. Li, C., Zhang, R., Wong, J., Gokmen, C., Srivastava, S., & Savarese, S. (2024). BEHAVIOR-1K: A Human-Centered, Embodied AI Benchmark with 1,000 Everyday Activities and Realistic Simulation. arXiv:2403.09227.
14. Liu, Y. (2026). Process Analysis Toolkit (PAT): Automated verification of concurrent and probabilistic systems. Nanyang Technological University.
15. Morgan Stanley. (2025). Humanoids: A $5 Trillion Market.
16. National Quantum Strategy. (2024). Quantum in Singapore: Opportunities for collaboration to drive mutual growth.
17. NIST. (2025). NIST AI RMF 2025 Updates: Toward Maturity-Based Certification of Embodied AI.
18. Open Source Robotics Alliance. (2024). Open-RMF development & next-generation orchestration: Maturity from local pilot to global standard. OSRA.
19. Poskitt, C. M., Sun, J., & Roychoudhury, A. (2024). REDriver: Runtime Enforcement for Autonomous Vehicles. IEEE/ACM 46th International Conference on Software Engineering (ICSE). arXiv:2401.02253.
20. Singapore Standards Council. (2021) TR 68-1:2021 Autonomous vehicles – Part 1 : Basic behaviour
21. Singapore Standards Council. (2021) TR 68-2:2021 Autonomous vehicles – Part 2 : Safety
22. Singapore Standards Council. (2021) TR 68-3:2021 Autonomous vehicles – Part 3 : Cybersecurity principles and assessment framework
23. Singapore Standards Council. (2021) TR 68-4:2021 Autonomous vehicles – Part 4 : Vehicular data types and formats
24. Singapore Standards Council. (2025). SS 713:2025 Data exchange between robots, lifts and automated doorways
25. Singapore Standards Council. (2025) TR 130:2025 Interoperability between central command and/or robotics middleware and autonomous mobile robots (AMRs) or AMR fleet management (AMR-FM)
26. SMU. (2025). Robots that rethink: A SMU project on self-adaptive embodied AI.
27. Workplace Safety and Health Council. (2026). Strengthening Safe Human-EAI Collaboration.

Singapore Standards Council

ASTM INTERNATIONAL

ADVANCING STANDARDS
TRANSFORMING MARKETS

ENTERPRISE SINGAPORE HEADQUARTERS

230 VICTORIA STREET #09-00
BUGIS JUNCTION TOWERS
SINGAPORE 188024

Tel +65 6898 1800

Web www.enterprisesg.gov.sg

ASTM INTERNATIONAL HEADQUARTERS

100 BARR HARBOR DRIVE
P.O. BOX C700
WEST CONSHOHOCKEN, PA 19428-2959
USA

Tel +1.610.832.9500

Web go.astm.org