ROBOTICS AND ARTIFICIAL INTELLIGENCE IN GLOBAL SOUTH SCIENCE SYSTEMS

PATHWAYS, INFRASTRUCTURE AND INSTITUTIONAL CHOICES: TECHNOLOGY PROFILE ON ROBOTICS AND ARTIFICIAL INTELLIGENCE

CENTRE FOR SCIENCE FUTURES | International Science Council

To cite this report: International Science Council (May 2026). *Robotics and Artificial Intelligence in Global South Science Systems.*

Funding acknowledgement: This work was conducted with the aid of a grant from the International Development Research Centre (IDRC), Ottawa, Canada. The views expressed herein do not necessarily represent those of the IDRC or its Board of Governors.

Authors: Venkat Nadella, Chandan Nagarajappa

Reviewer: Slava Tykhonov

Project coordination: Felix Dijkstal, Vanessa McBride

Project chair: David Castle

Design: Mr Clinton

Cover photo: imaginima

About the International Science Council

The ISC is an international non-governmental organization with a unique global membership that brings together 250 international scientific unions and associations, national and regional scientific organizations, including science academies, research councils, regional scientific organizations, international federations and societies, and academies of young scientists and associations.

The ISC works at the global level to catalyse change by convening scientific expertise, advice and influence on issues of major importance to both science and society.

Executive summary

Robotics and artificial intelligence (AI) together constitute a distinct class of scientific capability. Their development depends on algorithms and data, as well as physical infrastructure, safety and governance systems, and human capability built through sustained hands-on engagement with real systems. In digital AI, progress scales with compute and data; while AI-enabled robotic systems must sense, decide and act in the physical world under constraints of energy, latency, reliability and safety. These systems introduce distinctive safety and governance considerations due to their operation in physical and often human-facing environments. For science systems in the Global South, this distinction has practical implications. The central challenge is institutional fit: how to build and sustain robotics and AI capability under material, organizational and regulatory constraints that differ from those assumed in digital AI policy frameworks.

This technology profile examines how science systems across Africa, Asia, and Latin America and the Caribbean address this challenge through four pathways:

- Flagship institutional depth;
- Distributed and context-responsive capacity building;
- Mission-mode integration and governance of autonomy; and
- Testbeds as scientific infrastructure for embodied experimentation.

Drawing on seven case studies and desk research, it identifies five enabling conditions that shape what is feasible:

- Organizational forms fit for purpose;
- Governance and safety as prerequisites;
- Well-governed testbeds that substitute for scale;
- Capability developed through sustained practice; and
- Funding continuity.

Key finding: *The analysis shows that the future impact of robotics and AI on Global South science systems will depend on strategic choices rather than technological inevitability. Progress does not follow a single path. By recognizing robotics as scientific infrastructure, supporting multiple pathways and aligning investments with institutional and material realities, science system leaders can strengthen research capacity while safeguarding legitimacy, sustainability and public trust.*

Introduction

Robotics translates sensing, modelling and decision-making into material action in real-world environments. Artificial intelligence (AI) enables perception, planning and adaptive behaviour. Therefore, robotics and AI together form embodied intelligent systems whose capabilities and constraints differ fundamentally from those of purely digital AI (Brooks, 1991; Pfeifer and Bongard, 2006; Siciliano and Khatib, 2016). Unlike digital systems, robotics research depends on physical infrastructure, safety and governance mechanisms, skilled technical teams and sustained access to experimental environments. These requirements determine not only what research is possible, but how it must be organized. For many science systems in the Global South, these demands reflect real resource constraints that call for context-sensitive responses, not proxies for deficit or delay on a single developmental path. They are conditions that call for purposeful, context-sensitive approaches that differ structurally from the scaling dynamics that characterize digital AI development (OECD, 2023; UNIDO, 2024).

Conceptual framing: AI in robotic systems

In the current discussions of AI in physical systems, there are three related but distinct terms. Understanding the difference is important for policy.

- **Embodied AI** is an established academic framing rooted in cognitive science and robotics theory (Brooks, 1991; Pfeifer and Bongard 2006). Intelligence in robotic systems arises from the interaction between algorithms, physical bodies, environments and human oversight. Capabilities such as autonomy, manipulation and real-time control cannot be fully separated from physical embodiment.
- **Physical AI** is an industry-equivalent term that refers to the same class of systems with less theoretical precision. It was popularized in 2024–25.
- **Edge AI** refers to a specific computational architecture, running AI on a device rather than in the cloud. It is one deployment option within robotic systems, not a synonym for embodied intelligence.

In the theoretical tradition, **embodied intelligence** arises from the integration of algorithms and physical substrate. For science system planning, the relevant criterion is whether the system acts in the physical world – the deployment architecture of the AI layer is treated as a separate design question. ‘Robotics and AI’ is used in this profile as the broad umbrella term for the field. ‘Embodied AI’ refers to the subset of systems within it where the defining characteristic is physical action in the world – the deployment architecture of the intelligence layer, whether it runs in the cloud or on the robot itself (edge device), is addressed as a distinct technical dimension in the profile.

In practice, intelligence in robotic systems operates across a continuum: from on-board edge processing (AI runs on the robot itself) to cloud-mediated control (AI runs on remote infrastructure and communicates with the robot). This profile covers this full range as it applies to robotic systems in the Global South – whether the intelligence is embedded in the device, distributed across local infrastructure or accessed via cloud services. What unifies the scope is not the architecture of the AI, but the fact that the system acts in the physical world. This shifts attention from technological frontier performance toward the conditions under which scientific capability can be built, sustained and governed (Stilgoe, J. et al., 2013). Figure 1 maps key distinctions and their implications for science system planning.

Scope

This profile focuses on AI as applied to robotics. It treats AI as an enabling layer for perception, control, planning and autonomy in physical systems. It encompasses both research on robotics (advancing robotic systems, algorithms, hardware and control methods) and research that uses robotics as scientific instruments for discovery in other domains (autonomous field monitoring, laboratory automation or assistive care). The infrastructure, governance and capability requirements for these two orientations differ, and the case studies deliberately span both. The profile does not survey all AI applications in science or address digital transformation except where these intersect with embodied research. Geographically, it covers Africa, Asia (excluding China), and Latin America and the Caribbean, spanning university laboratories, national and mission-oriented research institutions, regional networks and selected innovation platforms. The objective is not to rank technologies or countries, but to identify pathways and enabling conditions that make AI-enabled robotics research feasible under constraints.

This profile does not provide cost or investment benchmarks. Establishment and operating costs vary substantially across Global South contexts depending on institutional type, national economy and funding source (government, non-governmental or private-sector organizations), and the institutions consulted did not share financial data. Prescribing investment levels would in any case contradict the profile's central finding that the cases examined yield no universally applicable model, neither organizational nor financial nor technical.

This technology profile supports strategic planning and decision-making by science system leaders, policy-makers and research funders. It does not suggest which robotics technologies to adopt. Instead, it asks how science systems can organize infrastructure, institutions and governance to build and sustain robotics and AI capability over time, and what opportunities that sustained capability enables for science and society.

Technology overview: AI-Enabled robotics for science system planning

From digital AI to AI-enabled robotics

Contemporary AI policy is shaped by developments in digital AI, where progress is driven by scale – large datasets, high compute and increasingly general-purpose models (Kaplan et al., 2020; Bommasani et al., 2022). Robotics and AI-enabled systems operate fundamentally under different constraints: intelligence must sense, decide and act in the physical world, often in real time and under safety-critical conditions. Progress depends less on scaling compute and more on integration between algorithms and hardware, perception and control, simulation and physical testing, and autonomy and human oversight (Brooks 1991; Siciliano and Khatib, 2016). Policies effective for digital AI – short funding cycles, software-first metrics, rapid-deployment incentives – do not translate directly to robotics and AI research. Figure 1 summarizes these distinctions across seven planning-relevant dimensions.

Key technical dimensions with policy relevance

Four technical dimensions directly related to science system planning.

- **Levels of autonomy** span a continuum from teleoperation to autonomous decision-making; many scientific applications rely on hybrid or human-in-the-loop configurations, particularly where safety, interpretability or regulatory constraints apply (Shahriari and Shahriari, 2017; Sheridan, 2021).

- **Forms of embodiment.** There are many forms of embodiment, including ground vehicles, aerial platforms (e.g. drones for field monitoring), marine systems (e.g. autonomous underwater vehicles for ocean research), fixed-base robotic arms (e.g. laboratory automation and surgical systems), and mobile manipulators (platforms combining a mobile base with articulated arms for manipulation tasks, such as assistive robots or inspection platforms). Each form of embodiment imposes different infrastructure and expertise requirements. For example, aerial platforms require airspace access and outdoor testing protocols; marine systems depend on waterproof hardware and field logistics; fixed-base manipulators rely on controlled environments and precision tooling. This diversity reinforces the need for plural research pathways (Siciliano and Khatib, 2016).

- **Placement of intelligence** is where the AI enabling layer runs. This can be on-board the robot (edge), on local infrastructure, in the cloud or across hybrid configurations. The placement affects latency, reliability, data governance and dependence on external systems (Latent AI, 2025). A field-monitoring drone in a low-connectivity region may require on-board edge processing; a laboratory automation platform may rely on local compute; a social robotics application may offload language processing to cloud services.

For science system planning, the critical question is not only where inference runs but who controls the ability to update the model: a system performing on-board inference may still depend on a foreign provider for retraining or model updates, creating a hidden external dependency that introduces the same vulnerabilities as full cloud dependence.

- **Safety and assurance are non-optional:** verification, validation and constrained experimentation are integral to embodied research design (International Organization for Standardization, 2014).

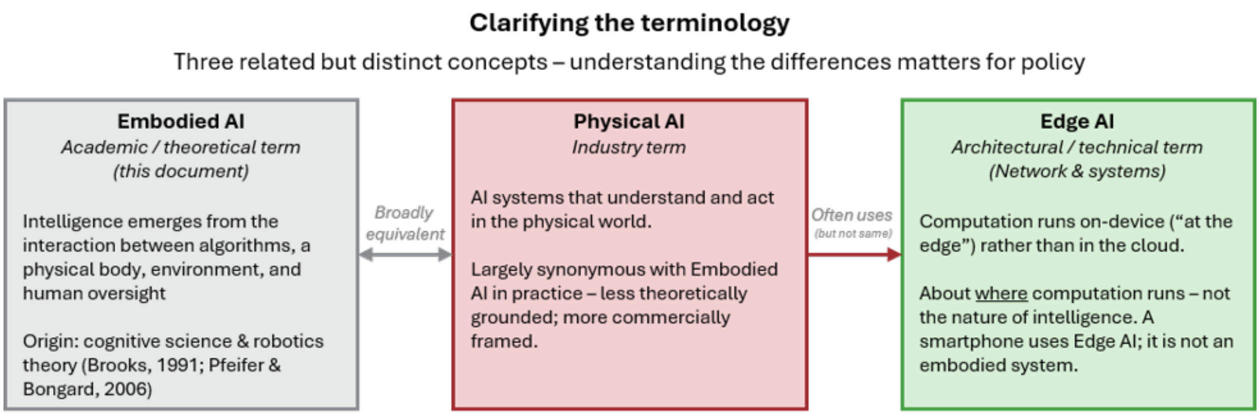

Key distinctions for science system planning

Digital AI e.g. large language models, foundation models	Dimension	Embodied AI / Physical AI e.g. robotics, autonomous field instruments Edge AI often used for computation placement
Scale: larger datasets, more compute, general-purpose model reuse	**Progress driver**	Integration: algorithms + hardware + environment + human oversight
Cloud servers, data pipelines, software frameworks	**Key infrastructure**	Physical labs, field sites, testbeds, safety & governance systems
Cloud-first, centralized servers, high bandwidth assumed	**Computation placement**	Often edge AI (on-device) for low latency & real-time operation, or hybrid cloud-edge depending on context
Data quality, compute costs, IP and bias	**Critical constraints**	Energy, latency, reliability, safety, real-time operation, maintenance
Data privacy, fairness, accountability for predictions	**Governance needs**	Safety, liability, verification & validation, human-in-the-loop oversight
Short-cycle, outcome-driven, rapid deployment incentives	**Funding fit**	Long-cycle, infrastructure investment, skill retention over time
Replication at scale improvements generalise across contexts	**Capability logic**	Context-specific embodied practice, improvements depend on local infrastructure

Policy implication: strategies effective for digital AI – short funding cycles, software-first metrics, rapid-deployment incentives do not translate to embodied/physical AI. Edge AI is one architectural option within embodied systems, not a substitute for physical infrastructure

Figure 1: Digital AI vs embodied or physical AI: concepts, terminology and key distinctions for science system planning. Edge AI (computation on-device) is an architectural option within embodied systems, not a synonym for embodied intelligence.

Robotics as scientific infrastructure

Robotics and AI research depend on access to experimental environments where physical interaction, failure and iteration are possible. For science systems operating under resource, regulatory or geographic constraints, the critical issue is not access to frontier hardware, but access to legitimate experimentation. This means experimentation conducted with adequate safety oversight, ethical accountability and institutional governance. It is distinct from ad hoc testing that lacks these structural guarantees.

Testbeds, simulation-first workflows and shared platforms enable research to proceed safely and reproducibly without requiring large hardware fleets. Importantly, these infrastructures are not neutral technical assets – their design and governance determine who can experiment, under what conditions and with what forms of accountability. Robotics platforms, testbeds and shared laboratories function as science system infrastructure in the same sense as libraries, field stations or computing facilities: decisions about access, safety standards and accountability are embedded in their design, not separate from it. Robotics infrastructure must therefore be understood as part of the science system itself, not merely as equipment.

The risk profile of embodied AI research – where the defining criterion is that the system acts in the physical world, regardless of where its intelligence runs – differs qualitatively from digital AI. This difference is not in absolute magnitude, and not by comparison with other scientific domains such as nuclear research, genomics or synthetic biology, each of which carries its own distinct risk profile – but specifically in its physical immediacy relative to purely digital systems.

A digital AI system that fails produces a wrong output, such as misclassifications, biased predictions or hallucinated responses. It remains idle in a server room. An embodied system that fails acts wrongly in the world: it moves, exerts force or makes a real-time decision in a physical environment from which it cannot simply be paused. There is no safe idle state. Failure modes are environment-dependent, difficult to anticipate fully in advance and physically consequential in ways that digital failures are not, which is why the margin for error is far narrower. This exposure is especially acute when systems interact with people or operate in safety-critical settings. This is why governance and safety capacity function as prerequisites for legitimate research – not compliance burdens applied after technical development, but structural conditions that determine whether research can proceed credibly and sustainably in the first place.

Technology landscape: AI-enabled robotics in Global South science systems

Robotics and AI capability in the Global South develops through four distinct pathways, each representing a different organizing logic through which AI-enabled robotics advances under constraint. These pathways are not stages of development or policy prescriptions; they reflect how different institutional contexts, resource environments and scientific objectives direct what is feasible and effective. The pathways are not mutually exclusive, nor are they fixed institutional categories where a case is assigned to one. This reflects the dominant organizing logic observed at the time of study, not a fixed institutional category. A recurring finding across all pathways is that enabling conditions – organizational fit, governance capacity (encompassing both the formal mechanisms for oversight, ethics review and safety assurance, and the institutional expertise to implement them), access to testbeds and funding continuity – are more decisive than technological ambition alone. Mapping these pathways is itself a contribution to equitable science system development: it avoids privileging a single model and instead identifies what is actually feasible under different resource and institutional conditions. It is important to recognize that building robotics and AI capability under constraint is not a deficit to be corrected but a different problem requiring different solutions. Figure 2 maps the four pathways, their organizing logics and design insights.

Four pathways: Building robotics & AI capability in Global South science systems

Pathway I **Flagship labs & deep capability building**	**Pathway II** **Distributed capacity through networks & field labs**	**Pathway III** **Mission-mode integration & governance of autonomy**	**Pathway IV** **Testbeds as scientific infrastructure**
Organising logic Intentional institutional design, long-term funding alignment, conservative engineering choices that allow capability to accumulate before scale or market pull.	**Organising logic** Distributed, context-responsive arrangements, expand participation and visibility by anchoring robotics to locally salient challenges.	**Organising logic** Government-funded agencies with defined national mandates integrate robotic autonomy as an enabling capability driven by operational demands, not research agendas.	**Organising logic** Shared, governed physical environments compensate for lack of scale; enable controlled, repeatable, safe experimentation under resource constraints.
Case studies • India: CAIR, IIT-Mandi • South Africa: Electronic Systems Lab, Stellenbosch University • India: ARTPARK, IISc	**Case studies** • Kenya: Social Robotics Lab, JKUAT • Latin America: LACORO Summer School • Africa: AFRON / AfROB	**Case studies** • India: Indian Space Research Organisation (ISRO)	**Case studies** • Global South: context-embedded lab & course testbeds • USA: Robotarium, Georgia Tech (benchmark reference)

Design insights **Institutional**: centre-based or platform-level design with pedagogy-first pipelines **Technical**: hybrid AI-control architectures, safety–critical conservatism **Policy**: mission alignment stabilizes long-horizon public funding	**Design insights** **Institutional**: small labs & regional networks can introduce new subfields **Technical**: commercial platforms & cloud AI lower entry barriers **Policy**: ethics & governance capacity essential for sensitive domains	**Design insights** **Institutional**: mission institution embed autonomy without dedicated robotics centres **Technical**: safety requirements favour hybrid, classical + AI architectures **Policy**: mission authority & accountability govern robotics & AI under high risk	**Design insights** **Institutional**: shared testbeds expand access without hardware duplication **Technical**: safety & verification mechanisms are enabling, not just safeguards **Policy**: operational capacity as critical hardware investment

Enabling conditions:	Organisational forms fit for purpose	Governance & safety as prerequisites	Well-governed testbeds substitute for scale	Capability through sustained practice	Funding continuity over funding amount

Figure 2: Four pathways through which robotics and AI capability is built, distributed, governed and enabled within science systems in the Global South. The enabling conditions at the base apply across all pathways.

Seven case studies across India, South Africa, Kenya, Latin America and Africa illustrate how these pathways operate in practice. Cases were selected to represent institutional diversity, not to rank performance. Three reference cases (ARTPARK, IISc; ISRO, India; and the Robotarium at Georgia Tech, USA) are drawn from desk research and included to illustrate platform-level design (the integration of research, startups, testbeds, and industry collaboration under a single institutional umbrella anchored at a flagship university); national mission institutional design; and testbed infrastructure, respectively. These cases are assessed for institutional design, not organizational performance. Figure 3 provides a structured comparison across all cases.

Pathway	Institution / Country	Organisational Form	Technical Approach	Critical Enabling Condition	Policy Insight
Pathway I: Flagship Labs & Deep Capability	**CAIR, IIT-Mandi / India**	University research centre; centre-based organisation	Simulation-first, edge-executed navigation; hybrid AI-control; hardware built in-house to reduce dependency	Pedagogy as pipeline: FDP feeds students into research, missions & ventures; curriculum-integrated capacity	Centre-based design + pedagogy-led pipelines enable deep capability before scale or market pull; mission alignment stabilises funding
	Electronic Systems Lab, Stellenbosch Univ. / South Africa	Department-embedded laboratory; disciplinary and industry-linked	Simulation-first, hardware-validated; classical control + selective AI for perception; no end-to-end	Industry engagement + regulatory environment shaping; clear agenda toward	Industry-linked funding substitutes for venture capital; regulatory environments drive validation over frontier performance; success

			learning (AI controls full system but verifiable steps) in safety-critical loops	reliability and validation	measured by industry-ready graduates
	ARTPARK, IISc / India	Platform-level national innovation infrastructure; multi-stakeholder: startup & public actors	System integration and validation focus; deployment readiness over isolated technical outputs	Long-horizon public funding + shared testbeds + governance arrangements spanning academia, industry, state	Platform-level entities bridge research, deployment and entrepreneurship without fragmenting capability or public accountability; requires stable state commitment
Pathway II: Distributed Networks & Field Labs	**Social Robotics Labs, JKUAT / Kenya**	Small university-embedded lab that introduces robotics as a new research field without mature systems; MSc project-driven	Commercial humanoid platforms; cloud-mediated AI; speech recognition, dialogue mgmt, perception; no hardware dev	Ethics review & responsible computing treated as enabling infrastructure, not procedural overhead	Context-responsive labs introduce new subfields without large centres; MSc projects build researcher community and disciplinary presence in a new field; ethics & governance capacity essential when targeting sensitive populations; ceiling without dedicated staffing
	LACORO Summer School / Latin America	Recurring regional summer school; rotating governance across host universities	Conceptual grounding and research orientation; no hardware or prototype building; short intensive exposure	Quality & diversity of instructors; rotating host model distributes ownership & resilience	Periodic convening strengthens regional capacity when permanent infrastructure is uneven; impact realised after event through collaborations, mentorships, research directions
	AFRON / AfROB / Africa	Continent-wide lightweight coordination network; no central laboratory or institute	Coordination, visibility & connection; ecosystem mapping; workshops at global conferences; no direct instruction	Dedicated operational support beyond volunteer effort; sustained funding for coordination infrastructure	Where formal institutions are sparse, coordination networks sustain community presence and researcher orientation — a real but partial intervention that cannot substitute for the infrastructure; most valuable as a bridge toward deeper institutional development, not an endpoint; sustainability requires explicit funding beyond volunteer effort
Pathway III: Mission-Mode Integration	**Indian Space Research Org. (ISRO) / India**	National mission institution; formal R&D under Dept of Space	Conservative hybrid architectures; classical control; rule-based logic + limited AI perception; rigorous inspection & certification	Mission authority to define requirements, orchestrate supply chain & retain system-level validation responsibility	Advanced robotic autonomy governed within mission institutions without dedicated robotics centres; mission authority & accountability decisive under high-risk, safety-critical conditions
Pathway IV: Testbeds as Infrastructure	**Robotarium Georgia Tech**	Shared remotely accessible research facility;	Multi-robot coordination; user-supplied	Centralised operation, maintenance,	Shared testbeds expand access to experimentation without duplicating

	(Benchmark) / USA	centralised management	algorithms; formal safety mechanisms (barrier certificates) enforce collision avoidance	scheduling & governance; stable institutional & funding support	hardware; remote access + embedded safety mechanisms (barrier certificates) enable responsibility; organisational & operational capacity as critical as hardware — key design lessons for Global South testbed planners

Figure 3: Case study comparison across seven cases plus one Global North benchmark: institutional form, technical approach, critical enabling conditions and policy insights. Pathway colours correspond to Figure 2.

The Robotarium at Georgia Tech – included as a Global North benchmark drawn from desk research – offers transferable design lessons rather than a replication model. Its core contribution is demonstrating that shared remote access, formally embedded safety mechanisms (barrier certificates for collision avoidance), and common experimental conditions together enable reproducible multi-robot research without each team requiring dedicated hardware. For Global South testbed planners, the critical implication is that organizational and operational capacity – scheduling, governance, maintenance, safety oversight – is as determinative of research access as the hardware investment itself.

Enabling conditions, opportunities and constraints

Across the cases examined, five structural conditions emerge as consistently more decisive for robotics and AI research than technological ambition alone. These are not best practices or prescriptions but structural conditions that shape feasibility, legitimacy and long-term impact. They are not a ranked checklist: all five interact, and weakness in any one can constrain capability even where the others are well-developed. Where these conditions are misaligned, robotics research tends to stall at the prototype stage or short-term demonstration, irrespective of the technical expertise invested. Figure 4 maps the five conditions and their implications for science systems.

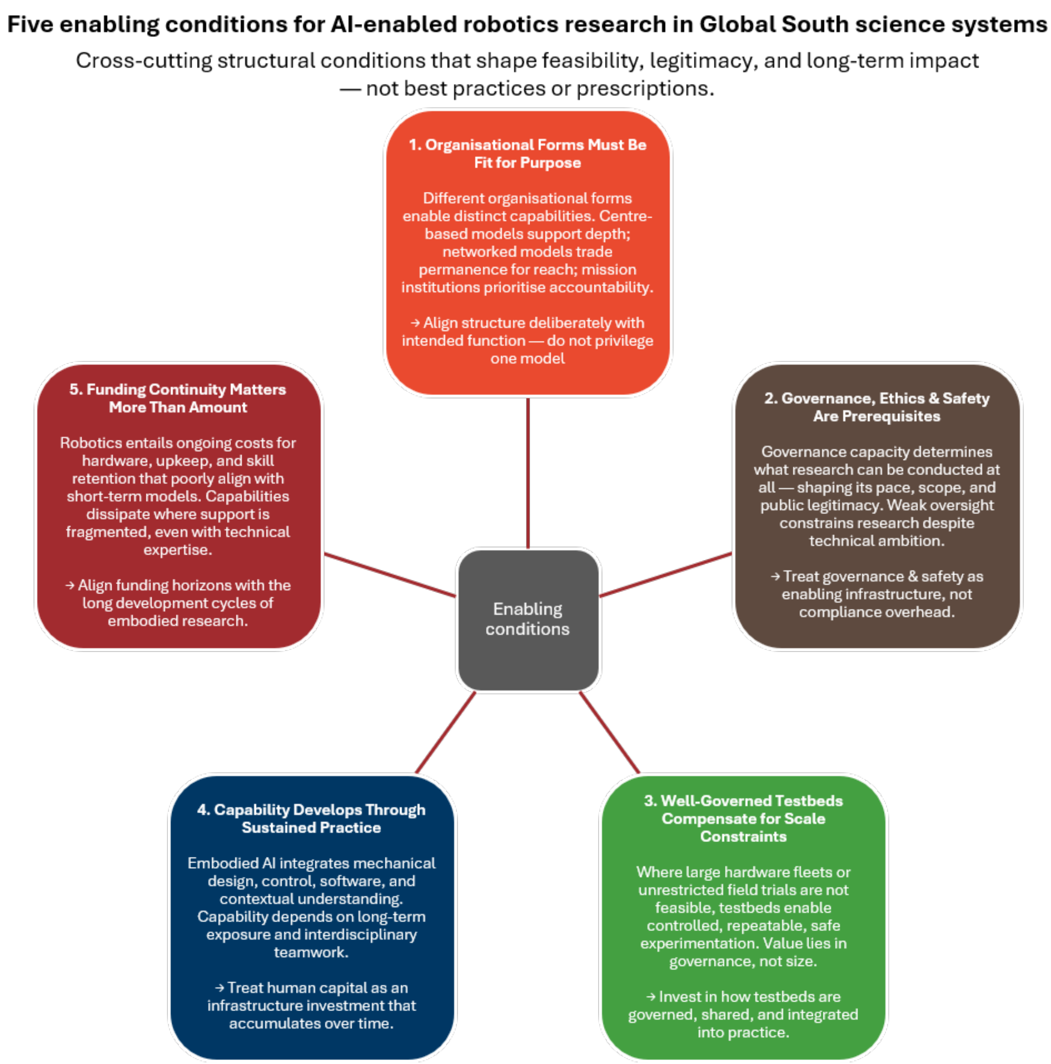

Figure 4: Five enabling conditions for AI-enabled robotics research in Global South science systems. The arrow in each card indicates the implication for science system planning.

Beyond the enabling conditions above, several **structural risks** warrant explicit attention in planning for Global South science systems.

- **Platform and supply-chain dependency.** Reliance on proprietary foreign hardware, cloud infrastructure or closed software ecosystems creates fragility: access can be disrupted by commercial decisions, export controls or supplier pricing changes. Building in-house hardware competence and favouring open or interoperable platforms where feasible reduces this exposure.

- **Research agenda and governance capture.** Industry funding in robotics and AI research can shape what questions are asked, which safety standards are treated as sufficient, and what governance positions are adopted – orienting institutions toward funder interests rather than public science system goals. This risk is distinct from hardware dependency: institutions may retain full equipment ownership while their research priorities, evaluation criteria and governance frameworks are systematically shaped by external commercial interests.

- **Capability dispersal.** Trained researchers and technical staff are mobile; end-of-project funding cycles can rapidly erode hard-won institutional knowledge. In resource-constrained settings where replacement recruitment is slow, this poses a greater threat to sustained capability than in well-resourced ecosystems.

- **Governance and regulatory gaps.** The absence of national regulatory frameworks for autonomous systems creates uncertainty that can deter investment, stall legitimate experimentation and leave institutions exposed when incidents occur. Proactive governance development – including ethics review processes and liability frameworks – reduces this uncertainty rather than constituting an additional burden.

- **Asymmetric international partnerships.** Collaborations with well-resourced Global North institutions can extract scientific value – data, publications, trained graduates – without building lasting local capacity, governance infrastructure or intellectual property. Structuring partnerships around explicit capacity-transfer obligations and shared governance mitigates this risk.

For most Global South science systems, the binding constraint is not the organizational model but capital access – hardware, maintenance, safety infrastructure and long-horizon funding. This makes distributed activities such as regional networks, summer schools and department-embedded labs the realistic and primary mechanisms for sustaining research presence in a capital-intensive field. These are not transitional arrangements awaiting replacement by more resource-intensive models; they are durable adaptations to structural resource conditions, and the policy question is how to strengthen and sustain them, not how to bypass them.

Future directions: Considerations for science system planning

The considerations below span technical, infrastructural, organizational and policy dimensions.

World models, simulation and their limits. Recent advances in world models, spatial intelligence and large-scale simulation can reduce the cost and risk of embodied learning by shifting experimentation into virtual environments, accelerating iteration and broadening participation where access to physical infrastructure is limited (Bommasani et al., 2022). However, they do not eliminate the need for physical interaction, safety validation and context-specific testing and serve as complements rather than substitutes for real-world experimentation (Zhao et al., 2020; Aljalbout et al., 2025). Simulation-led research without pathways for real-world validation risks producing systems that are difficult to interpret, certify or deploy responsibly.

Testbeds as strategic infrastructure. Shared, modular or domain-specific testbeds should be understood as strategic infrastructure choices that shape which robotics and AI research is feasible under constraints. They support controlled experimentation, skill formation and reproducibility, particularly where large facilities or unrestricted field trials are unavailable. The key issue for planning is not whether testbeds are adopted, but how they are scoped, governed and sustained within stable institutional and funding arrangements (Mokhtarian et al., 2024).

Startups as bounded experimental pathways. Startups and entrepreneurial initiatives increasingly contribute to robotics and AI through rapid prototyping and exploratory experimentation. Their value to science systems is greatest when they operate as bounded experimental pathways that complement academic and public research through clearly defined interfaces. Without such boundaries, engagement risks prioritizing short-term deployability over long-term capability formation and public accountability – a dynamic visible in the case studies, where industry-linked funding consistently oriented research toward validation and deployment readiness over frontier capability development (Figure 3).

Mission institutions as diffusion mechanisms. Mission-mode institutions – space agencies, national defence research bodies and public operators of autonomous systems – concentrate capital and technical demands at a scale that can generate diffusion effects across the broader science ecosystem. Whether and how these effects reach civilian and academic research depends on institutional design choices and not on the nature of the missions alone. The same mission-mode investment can be structured to build reusable domestic capacity or leave diffusion to chance. Making diffusion an explicit goal – rather than treating it as an incidental byproduct – is an underexplored dimension of science system planning in contexts where mission institutions are among the largest actors in the robotics and AI landscape (Mazzucato, 2018; Mazzucato, 2021).

Each of these considerations reinforces a central finding of the profile. **Advances in robotics and AI do not automatically translate into scientific capacity.** Overreliance on simulation, without corresponding investment in real-world experimentation and governance, can narrow rather than expand what science systems can do – producing outputs that are difficult to validate, certify or deploy responsibly in real-world settings (Amodei et al., 2016).

Conclusion

This technology profile argues that robotics and AI-enabled systems constitute a distinct class of scientific capability, shaped as much by physical operation, safety and institutional design as by algorithms or data. Progress depends on sustained access to infrastructure, governance arrangements that enable legitimate experimentation and human capability developed through prolonged engagement with real systems – regardless of whether the AI layer runs on-board the robot, on local infrastructure or via the cloud. For Global South science systems, these are not disadvantages to overcome but conditions that require intentional, context-sensitive pathways for capability building. Across the cases examined, no single organizational model, funding instrument or technical trajectory emerges as universally appropriate. Enabling conditions such as institutional fit, governance capacity, access to infrastructure and funding continuity are consistently more decisive than technological ambition alone.

A direct implication follows: policy instruments calibrated for digital AI – short funding cycles, software-first metrics and rapid-deployment incentives – are structurally misaligned with robotics and AI research, which requires physical infrastructure, sustained safety governance and long-horizon capability development. Reorienting funding and evaluation frameworks to reflect these realities matters as much as any technology choice.

Recommendations

The recommendations in Figure 5 support four stakeholder audiences in making such choices without prescribing uniform solutions.

For decision-makers, these recommendations translate into three immediate priorities:

- Invest in shared and governed testbeds;
- Integrate safety and governance into research design; and
- Align funding with long-term capability development.

Audience	**Infrastructure & Organisational Design** How to structure and invest in physical and institutional architecture	**Governance, Safety & Ethics** *How to embed accountability, oversight and risk capacity*	**Capability, Continuity & Collaboration** *Sustaining human, financial and partnership capacity over time*
National Policymakers & Science Ministries	1.1 Treat robotics as scientific infrastructure — not short-cycle innovation projects *Plan for Robotics & AI as infrastructure requiring testbeds, technical staff, safety capacity, and ongoing maintenance rather than project-based or outcome-driven funding.*	1.2 Design for multiple institutional pathways — do not standardise organisational models *Enable a portfolio of forms — centres, embedded labs, networks, mission institutions — each aligned with distinct objectives such as depth, diffusion, or accountability.*	1.3 Align funding horizons with the material realities of Robotics research *Use instruments that accommodate long development cycles, infrastructure maintenance, and skill retention rather than privileging rapid deployment or short-term demonstrators.*
Public Research Funders & Funding Agencies	2.1 Invest in shared, modular testbeds as access-enabling infrastructure *Prioritise testbeds that are governed, shareable, and safe — reducing hardware duplication while expanding who can conduct legitimate experimentation.*	2.2 Require governance & safety capacity as part of research design — not post hoc compliance *Embed validation, oversight, and risk containment expectations directly into funding calls, particularly for human-facing or public-domain systems.*	2.3 Evaluate continuity & capability formation alongside publications and prototypes *Include infrastructure utilisation, team stability, and collaborative capacity as progress indicators when assessing Robotics & AI research programmes.*
University & Research Institution Leaders	3.1 Align organisational structure explicitly with intended research function *Be clear about whether initiatives aim for deep capability, education and diffusion, mission-oriented deployment, or regional coordination — and design governance accordingly.*	3.2 Structure industry & startup engagement to safeguard scientific priorities *Engage non-academic actors through clearly defined interfaces enabling experimentation and translation without displacing long-term research goals or accountability.*	3.3 Protect continuity of technical capability — not just individual training *Plan for continuity in staffing, technical roles, and long-term engagement with systems; capability is institutional and cumulative despite technical expertise.*

Regional & International Science Organisations	4.1 Enable South-South collaboration through coordination & shared infrastructure *Support lightweight networks, shared platforms, and collaborative challenges that expand reach without requiring permanent organisational integration.*	4.2 Facilitate exchange on governance, safety & testbed design — not only on technology *Prioritise knowledge sharing on institutional and governance practices that enable responsible research under diverse resource and regulatory constraints.*	4.3 Fund coordination mechanisms sustainably — beyond volunteer-driven models *Sustaining continental or regional networks (e.g. AFRON) requires dedicated operational support; continuity of coordination is itself a science system investment.*

These recommendations prioritise legitimacy and sustainability over frontier performance. They do not prescribe uniform solutions — they support deliberate, context-sensitive choices.

Figure 5: Recommendations for science system leaders and policy-makers, organized by audience and three cross-cutting themes: infrastructure and organizational design; governance, safety and ethics; and capability, continuity and collaboration.

Glossary of terms

- **Adaptive behaviour:** The capacity of a robotic or AI-enabled system to modify its actions in response to changes in its environment, enabled by AI perception, planning, and learning methods.

- **Aerial platforms / drones:** Robotic platforms operating in the air, including drones used for field monitoring. Identified in this profile as an embodiment form requiring airspace access and outdoor testing protocols. See also Forms of embodiment.

- **ARTPARK / IISc:** AI & Robotics Technology Park at the Indian Institute of Science - the profile's primary reference case for platform-level institutional design, integrating research, startups, testbeds, and industry collaboration under a single university umbrella. Cited as a Pathway I (Flagship Institutional Depth) reference case drawn from desk research.

- **Asymmetric international partnerships:** Research collaborations in which scientific value is extracted from partner institutions without corresponding development of local capability, governance infrastructure, or long-term institutional benefit. Identified in this profile as a structural risk for robotics and AI research.

- **Autonomous systems:** Robotic or AI-enabled systems capable of operating and making decisions with varying degrees of independence from continuous human control, spanning a continuum from human-supervised teleoperation to autonomous decision-making with reduced human supervision.

- **Barrier certificates:** Formal mathematical safety mechanisms - exemplified in the Robotarium at Georgia Tech - used in multi-robot testbeds to enforce specified safety constraints such as collision avoidance under defined assumptions. They formally enforce particular constraints rather than guaranteeing overall safe operation of the broader experiment.

- **Capability dispersal:** The erosion of accumulated institutional knowledge and technical skill that occurs when trained researchers and staff leave, particularly after end-of-project funding cycles in resource-constrained settings.

- **Diffusion effects:** The capacity of mission-mode institutions - space agencies, defence research bodies, public operators of autonomous systems - to generate spillover benefits for civilian and academic research.

- **Digital AI:** AI systems that operate entirely in computational environments - processing data, generating predictions or outputs - without sensing, deciding, or acting in the physical world. Policy frameworks calibrated for digital AI (short funding cycles, software-first metrics, rapid-deployment incentives) do not translate directly to embodied robotics and AI.

- **Distributed and Context-Responsive Capacity Building (Pathway II):** An institutional pathway in which robotics and AI capability is developed across multiple institutions, regional networks, and department-embedded labs, with methods adapted to local material, organisational, and regulatory conditions rather than imported from well-resourced contexts. Exemplified by the Social Robotics Lab at JKUAT (Kenya), the LACORO Summer School (Latin America), and AFRON/AfROB (Africa).

- **Edge AI:** A computational architecture in which AI inference runs on the device itself - on-board the robot - rather than in a remote cloud. One deployment option within embodied systems, not a synonym for embodied intelligence; the profile distinguishes the two explicitly.

- **Embodied AI:** In this profile: AI systems whose defining characteristic is physical action in the world - regardless of where the intelligence layer runs (on-device, on local infrastructure, or in the cloud), and distinct from Edge AI, which describes only the deployment architecture. The underlying academic framing, rooted in cognitive science and robotics, holds that intelligence in robotic systems arises from the interaction between algorithms, physical bodies, environments, and human oversight.

- **Embodied experimentation:** Research and testing involving robotic or AI-enabled systems operating in physical environments, where sensing, movement, interaction, and safety constraints shape experimental design.

- **Enabling conditions:** The five structural conditions identified in this profile as more decisive for robotics and AI capability than technological ambition alone: organisational forms fit for purpose (see Institutional fit / Organisational fit); governance and safety as prerequisites (see Governance capacity, Ethics review, Safety assurance); well-governed testbeds (see Testbed); capability developed through sustained practice (see Sustained practice); and funding continuity (see entry). The conditions interact, and weakness in any one constrains capability even where others are well-developed.

- **Ethics review:** A formal institutional process for assessing the ethical implications of proposed research involving autonomous or robotic systems. Treated in this profile as a prerequisite for legitimate experimentation rather than an optional compliance step applied after technical development.

- **Flagship Institutional Depth (Pathway I):** An institutional pathway organised around a major research centre or university laboratory with long-term funding, conservative engineering choices, and safety-critical quality standards - designed to accumulate deep capability before scaling or commercialising. Exemplified by CAIR at IIT-Mandi (India) and the Electronic Systems Lab at Stellenbosch (South Africa), with ARTPARK at IISc (India) cited as a platform-level reference case (see entry).

- **Forms of embodiment:** The physical configurations through which robotic systems operate - a planning-relevant dimension in this profile. Includes ground vehicles, aerial platforms (see Aerial platforms / drones), marine systems, fixed-base robotic arms (including laboratory automation, see entry), and mobile manipulators (see entry). Each form imposes distinct infrastructure and expertise requirements.

- **Funding continuity:** The structural condition in which infrastructure and capability funding is sustained across project cycles - through institutional mandates, mission funding, or long-horizon grants - rather than depending on short cycles that strand hardware, safety infrastructure, and trained staff when projects end. Identified in this profile as more decisive for robotics and AI than funding volume.

- **Global South:** A geopolitical shorthand used throughout this profile to refer to countries in Africa, Asia (excluding China), and Latin America and the Caribbean, whose science systems operate under material, organisational, and regulatory constraints that differ structurally from those assumed in digital AI policy frameworks.

- **Governance capacity:** The formal mechanisms for oversight, ethics review, and safety assurance, together with the institutional expertise required to implement them. Identified in this profile as a prerequisite for legitimate robotics and AI research, not a compliance function applied after the fact.

- **Governance capture:** A risk in which industry or funder interests systematically shape research questions, evaluation criteria, safety thresholds, and governance positions - orienting institutions toward funder priorities rather than public science-system goals. Distinct from hardware or supply-chain dependency: institutions may retain full equipment ownership while their research agenda is captured.

- **Human-in-the-loop:** A system configuration in which a human operator retains meaningful oversight, decision authority, or intervention capacity within an autonomous system's operation. Required in many scientific and safety-critical applications, particularly where interpretability or regulatory constraints apply.

- **Hybrid configurations:** Robotics and AI system designs that combine multiple modes - automated and AI-enabled functions with classical control, on-device (edge) and cloud computation, or autonomous operation and human-in-the-loop oversight. In this profile, hybrid configurations are highlighted where safety, interpretability, or regulatory constraints make a single approach insufficient.

- **Institutional fit:** The alignment between an organisation's form, mandate, resources, and governance structures and the type of robotics and AI capability it is trying to build. Identified in this profile as the central challenge for Global South science systems building robotics and AI capability under constraint. The profile uses 'institutional fit' (introductory framing) and 'organisational forms fit for purpose' (Figure 2 and the enabling conditions discussion) as parallel terms for the same concept; see also Organisational fit.

- **Laboratory automation:** The use of robotic systems as scientific instruments to perform laboratory tasks (such as sample handling, assay running, or instrument operation) - distinguished in this profile from research on robotics. Falls within the fixed-base robotic arm category and appears in case studies including the Electronic Systems Lab and ARTPARK at IISc. See also Forms of embodiment.

- **Latency:** The delay between a robotic system's sensor input, computation, and control response. Low latency is critical where systems must react in real time; the placement of AI inference - on-board (edge), local, or cloud - directly affects latency.

- **Legitimate experimentation:** Experimentation conducted under conditions of adequate safety oversight, ethical accountability, and institutional governance, as distinct from ad hoc testing that lacks these structural guarantees. Framed in this profile as the central condition under which robotics and AI research can proceed credibly and sustainably.

- **Levels of autonomy (autonomy continuum): A** planning-relevant dimension spanning the range from teleoperation to autonomous decision-making. Many scientific applications rely on hybrid or human-in-the-loop configurations, particularly where safety, interpretability, or regulatory constraints apply.

- **Mission institutions:** Government-funded organisations - space agencies, national defence research bodies, public operators of autonomous systems - that concentrate capital and technical demand at a scale capable of generating diffusion effects (see entry) across the broader science ecosystem. Whether such diffusion reaches civilian and academic research depends on institutional design choices, not the missions themselves.

- **Mission-Mode Integration (Pathway III):** An institutional pathway in which government-funded agencies with defined national mandates - such as space agencies or defence research bodies - integrate robotic autonomy as an enabling capability driven by operational requirements rather than research agendas. Exemplified in this profile by the Indian Space Research Organisation (ISRO).

- **Mobile manipulators:** Robotic platforms combining a mobile base with one or more articulated arms for manipulation tasks, such as assistive robots or inspection platforms. Require distinct infrastructure and expertise from fixed-base, aerial, or marine systems. See also Forms of embodiment.

- **Organisational fit:** See Institutional fit. (The profile uses both terms for the same concept.)

- Physical AI: An industry term (popularised circa 2024-25) referring to AI systems that act in the physical world. In this profile, it overlaps with embodied AI but is treated as less theoretically precise than the academic term.

- **Placement of intelligence:** A technical dimension describing where the AI enabling layer runs - on-board the robot (edge), on local infrastructure, in the cloud, or across hybrid configurations. Affects latency, reliability, data governance, and external dependency; this profile flags that control over model updates can create hidden external dependencies even when inference runs locally.

- **Platform and supply-chain dependency:** A structural risk arising from reliance on proprietary foreign hardware, closed software ecosystems, commercial cloud services, or external suppliers - creating fragility when access is disrupted by commercial decisions, export controls, or pricing changes. Mitigated by building in-house hardware competence and favouring open or interoperable platforms where feasible.

- **Reality gap:** The performance discrepancy between a robotic system validated in simulation and its behaviour when deployed in the physical world. A core challenge in simulation-first robotics research, addressed through sim-to-real transfer methods and pathways for physical validation.

- **Regulatory gaps:** The absence of national regulatory frameworks for autonomous systems, identified in this profile as a structural risk that creates uncertainty deterring investment, stalling legitimate experimentation, and exposing institutions when incidents occur. The profile argues that proactive governance development - including ethics-review processes and liability frameworks - reduces this uncertainty rather than constituting an additional burden. Distinct from Governance capacity (presence of capability) by addressing the absence of frameworks.

- **Reproducibility:** The capacity of experimental results to be reliably replicated under consistent conditions. In robotics, reproducibility is supported by shared testbeds with common experimental conditions, formal safety mechanisms, and standardised platforms.

- **Robotarium:** A shared, remotely accessible multi-robot research facility - the Georgia Tech Robotarium serving as the benchmark reference case in this profile - that enables reproducible multi-robot experiments without each team requiring dedicated hardware. Its design embeds barrier-certificate safety mechanisms (see entry) and standardised experimental conditions.

- **Safety assurance:** The set of verification, validation, and constrained-experimentation processes required to confirm that a robotic system behaves safely under defined conditions. Integral to embodied AI research design rather than a compliance add-on applied afterwards.

- **Science systems:** The institutions, funding arrangements, governance structures, infrastructures, and human networks through which scientific research and innovation are organised within a country or region.

- **Sim-to-real transfer:** The process of deploying a robotic system, model, or control strategy trained or validated in simulation into real-world environments. A key technical challenge because simulation cannot fully replicate physical-world variability.

- **Simulation-first workflows:** A research methodology in which robotic system development and validation begin in virtual environments before physical deployment, reducing cost and risk. This profile flags that simulation-first work without pathways for real-world validation risks producing systems that are difficult to interpret, certify, or deploy responsibly.

- **Startups as bounded experimental pathways**: A concept in this profile describing the productive role of startups when they operate as well-defined complements to academic and public research, contributing rapid prototyping and exploratory experimentation through clearly defined interfaces.

- **Structural risks:** The five risks identified in this profile as warranting explicit attention in planning for Global South science systems: platform and supply-chain dependency; research agenda and governance capture (see Governance capture); capability dispersal; governance and regulatory gaps (see Regulatory gaps); and asymmetric international partnerships. See individual entries.

- **Sustained practice:** Long-term, hands-on engagement with real robotic systems through which capability in embodied AI is developed and retained. Identified in this profile as an enabling condition; cannot be replicated through simulation alone or through short-term, project-based engagement.

- **Teleoperation:** Remote human control of a robotic system, representing one end of the autonomy continuum; contrasted with autonomous decision-making in which the system acts without direct human instruction.

- **Testbed:** A governed physical or simulated environment providing shared, controlled, and repeatable conditions for robotics and AI experimentation. Identified in this profile as strategic scientific infrastructure that enables safe research without requiring large hardware fleets.

- **Testbeds as Scientific Infrastructure for Embodied Experimentation (Pathway IV)**: An institutional pathway in which shared and governed testbeds - physical or simulated - substitute for the large hardware fleets that individual institutions cannot maintain, enabling controlled, safe, and reproducible robotics and AI experimentation. The Robotarium at Georgia Tech is the principal design reference; context-embedded laboratory and course testbeds in Global South settings are also represented in the case studies.

- **Verification and validation (V&V)**: The processes used to confirm that a robotic system meets its specifications (verification) and fulfils its intended purpose safely in real-world conditions (validation). Non-optional in embodied AI research design.

- **World models:** AI models that build internal representations of physical environments, enabling robotic systems to reason about conditions and predict consequences of actions. Recent advances can reduce the cost and risk of embodied learning by supporting large-scale simulation - though this profile cautions that they complement rather than substitute for real-world experimentation.

References

Aljalbout, E. et al. 2025. The Reality Gap in Robotics: Challenges, Solutions, and Best Practices. *ArXiv*. https://arxiv.org/abs/2510.20808.

Amodei, D. et al. 2016. Concrete Problems in AI Safety. *ArXiv*. https://arxiv.org/abs/1606.06565.

Bommasani, R. et al., 2022. On the Opportunities and Risks of Foundation Models. *ArXiv*. https://arxiv.org/abs/2108.07258.

Brooks, R. A. 1991. Intelligence Without Representation. *Artificial Intelligence,* Vol. 47, No. 1, pp. 139–59. https://doi.org/10.1016/0004-3702(91)90053-M.

International Organization for Standardization. 2014. *ISO 13482:2014 Robots and Robotic Devices – Safety Requirements for Personal Care Robots*. Switzerland, International Organization for Standardization. https://www.iso.org/obp/ui/en/#iso:std:iso:13482:ed-1:v1:en.

Kaplan, J. et al. 2020. Scaling Laws for Neural Language Models. *ArXiv*. https://arxiv.org/abs/2001.08361.

Latent AI. 2025. *AI's Edge Continuum: A New* Look *at the Cloud Computing Role in Edge AI*. Princeton, New Jersey, Latent AI. https://latentai.com/white-paper/ai-edge-continuum/.

Mazzucato, M. 2018. Mission-Oriented Innovation Policies: Challenges and Opportunities. *Industrial and Corporate Change*, Vol. 27, No. 5, pp. 803–15. https://doi.org/10.1093/icc/dty034.

Mazzucato, M. 2021. Mission Economy: A Moonshot Guide to Changing Capitalism. New York, Harper Business.

Mokhtarian, A. et al. 2024. A Survey on Small-Scale Testbeds for Connected and Automated Vehicles and Robot Swarms. *ArXiv*. https://arxiv.org/abs/2408.14199.

OECD. 2023. Artificial Intelligence in Science: Challenges, Opportunities and the Future of Research. Paris, Organisation for Economic Co-operation and Development (OECD) Publishing. https://doi.org/10.1787/a8d820bd-en.

Pfeifer, R., and Bongard, J. 2006. *How the Body Shapes the Way We Think: A New View of Intelligence*. MIT Press. https://doi.org/10.7551/mitpress/3585.001.0001.

Shahriari, K and Shahriari, M. 2017. IEEE standard review – Ethically aligned design: A vision for prioritizing human wellbeing with artificial intelligence and autonomous systems, *2017 IEEE Canada International Humanitarian Technology Conference (IHTC)*. Toronto, Institute of Electrical and Electronics Engineers (IEEE), pp. 197–201. https://ieeexplore.ieee.org/document/8058187.

Sheridan, T. B. 2021. Human Supervisory Control Of Automation. In: Salvendy G. and Karwowski, W. (eds), *Handbook of Human Factors and Ergonomics*, John Wiley & Sons, Ltd, pp. 736–60. https://doi.org/10.1002/9781119636113.ch28.

Siciliano, B. and O. Khatib. 2016. *Springer Handbook of Robotics*. Springer Handbooks, Springer.

Stilgoe, J. et al. 2013. Developing a Framework for Responsible Innovation. *Research Policy,* Vol. 42, No. 9, pp. 1568–80. https://doi.org/10.1016/j.respol.2013.05.008.

UNIDO. 2024. *Industrial* Development *Report 2024: Turning Challenges into Sustainable Solutions: The New Era of Industrial Policy*. United Nations Industrial Development Organization (UNIDO). https://doi.org/10.18356/9789211066647.

Zhao, W. et al. 2020. Sim-to-Real Transfer in Deep Reinforcement Learning for Robotics: a Survey, *2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, Canberra, Institute of Electrical and Electronics Engineers (IEEE), pp. 737–744, https://doi.org/10.1109/SSCI47803.2020.9308468.