GSMA™

The Path to Industrial Physical AI at Scale

Connectivity, Computing and Ecosystem Collaboration

A Practical Maturity Model
for Industrial Physical AI Projects



Disclaimer

This report is provided for general informational purposes only and may include content derived from third-party sources or generated by Artificial Intelligence. GSMA makes no representation or warranty, whether express or implied, as to the accuracy, completeness or reliability of its contents and nothing herein shall constitute advice, recommendation or endorsement by GSMA. To the fullest extent permitted by applicable law, GSMA and its affiliates shall have no liability for any loss or damage arising out of or in connection with any reliance on, or use of, this report.

Anti-Trust Statement

GSMA is committed to compliance with all applicable competition and antitrust laws. Nothing in this report is intended to facilitate or encourage any agreement, coordination or exchange of competitively sensitive information between market participants or otherwise restrict independent commercial decision-making or competition. All organisations remain solely responsible for ensuring their own compliance with applicable competition and antitrust laws. For further information, please refer to https://www.gsma.com/about-us/legal/antitrust-policy/

Contributors

Hangzhou Qiantang

Scenario Opening | Policy Collaboration | Public Services | Verification Operations, and Industrial Ecosystem Organization

CAILINK

Industrial 5G/5G-A Independent Private Network | Computing and Network Integration | On-site Delivery and Cross-site Replication

intel

Terminal-edge heterogeneous computing | AI software stack | developer and ecosystem support

Rockwell Automation

Industrial control | digital twins | safety systems | real production line integration and validation

GALBOT

Continuous operation, real-world application | commercial verification of robots

DEXTELEOP 灵御智能

Real-machine data production | multimodal synchronization | data quality and cost efficiency

Foreword



Alex Sinclair
Chief Technology Officer
GSMA

Artificial intelligence is moving beyond the digital world and into the physical world. As AI enables machines to perceive their surroundings, perform tasks and participate in production, the way we measure its value must also evolve. We must look not only at what these systems can do, but at whether they can deliver value consistently, reliably and economically in real-world environments.

For industrial physical AI, moving from successful demonstrations to deployment at scale requires more than stronger models or more advanced robots. It requires connectivity, computing, industrial systems and production processes to work together. The real measure of progress is our ability to translate technological capability into measurable and sustained production outcomes.

The mobile industry has an important contribution to make. Connectivity and edge computing can support collaboration between devices, on-site data processing and distributed intelligence. But technology choices must be guided by production needs, and their value must ultimately be demonstrated through the performance of the complete system.

Achieving scale also requires collaboration across industries and institutions. Industry organisations, industrial enterprises, robotics developers, mobile operators and technology partners must work together to address real production needs. By opening up real-world industrial use cases, coordinating resources and undertaking joint validation, we can advance interoperability and develop reusable implementation approaches. Turning experience from individual projects into practices that can be applied across sites and scenarios will help reduce repeated integration work and accelerate the transition from pilots to deployment at scale.

This white paper brings together experience and perspectives from different fields. It offers a valuable reference for understanding the path from proof of concept to production readiness and deployment at scale. I hope it will foster wider dialogue and collaboration, help partners better understand the conditions needed for scale, and promote the sharing of practical experience across the ecosystem.

I would like to thank all the organisations and experts who have contributed to this white paper. This collective effort demonstrates the value of cross-industry collaboration and provides a foundation for further joint work. Together, we can help industrial physical AI move from individual proofs of concept to deployment at scale, delivering meaningful and lasting value for industrial transformation.

Executive Summary

Recent advances in AI have increased interest in physical AI for manufacturing, warehousing and logistics. Robots, intelligent equipment and autonomous systems are being tested in increasingly complex production environments, with some applications progressing towards real-world deployment. This does not imply that industrial physical AI has already achieved widespread, scalable commercial deployment.

However, there remains a significant gap between a robot completing a prescribed action and a system continuously creating value in a complex industrial environment. The production environment demands not only that robots possess perception, reasoning, and execution capabilities, but also that the entire system meets production, safety, and economic requirements for long-term operation. Therefore, the scaling of industrial physical AI cannot solely rely on model capabilities or robot body performance, but requires a collaborative closed loop of data, connectivity, computing, industrial control, safety, and operations.

This report presents three core observations.



The competitive focus of industrial physical AI is shifting from single-task capability to production-level continuous operation capability. The technical value ultimately needs to be verified through task performance, continuous operation, safety, and economy in real workstations.



The value of infrastructures such as 5G and edge computing should not be judged solely based on peak rates, single-shot latency, or computing performance, but rather measured through their actual support for mobile operations, multi-robot collaboration, multi-modal data processing, and on-site operations. Critical closed loops involving personnel safety, equipment protection, and high-frequency motion control should still be prioritized and retained within the robot itself or reliable on-site systems.



The scaling of industrial physical AI is a systems-level challenge spanning multiple technologies and industries. Manufacturing enterprises, robotics companies, connectivity and computing providers, industrial automation enterprises, system integrators, and regional industrial organizations need to collaborate around real production tasks, gradually establishing open architectures, maturity models, evaluation methods, and replicable implementation mechanisms.

Contents

Foreword

01 Industrial Physical AI: From Validation to Production 02

1.1 From digital intelligence to the physical world 02

1.2 The verification value in industrial scenarios 02

02 Core Challenges of Scalable Implementation 03

2.1 Data and model closed loop 03

2.2 Connectivity and computing collaboration 04

2.3 Industrial system integration 04

2.4 Safety and commercial feasibility 04

03 Connectivity and Edge Intelligence Architecture 05

3.1 Task division and network coordination among endpoints, edges, and clouds 05

3.2 The role of 5G and 5G-Advanced 06

3.3 Synergy of multiple technologies 06

04 A Proposed Maturity Model for Production-Ready Industrial Physical AI 07

4.1 M0: Proof of concept 07

4.2 M1: Connectable 08

4.3 M2: Closed-loop capability 08

4.4 M3: Production-ready 08

4.5 M4: Replicable 08

05 Ecosystem Collaboration: A Proposed Implementation Blueprint for Qiantang 09

5.1 From joint publishing to joint verification 09

5.2 The value of Qiantang's regional scenario platform 09

5.3 From scenario verification to reproducible outcomes 09

06 Practical Steps for Project Teams 10

Conclusion 11

01 Industrial Physical AI: From Validation to Production

1.1 From digital intelligence to the physical world

Physical AI brings artificial intelligence into the physical world. It combines perception, reasoning, decision-making, and action, so that robots, machines, and intelligent devices can interpret on-site information, decide what to do, and act through motors, robotic arms, tools and other actuators. Unlike AI that primarily generates or analyses digital content, physical AI operates within real environments where movement, timing, physical constraints, and safety directly affect the outcome.

In industry, physical AI combines artificial intelligence with sensors, connectivity, computing, industrial control, and execution systems. For example, an autonomous mobile robot can detect people and obstacles as it moves materials through a factory; a vision-guided robotic arm can recognise components presented in different positions and adjust its grip; and an intelligent inspection system can identify a possible defect, reposition a camera or product for a closer view, and route the item for further checks. These systems do more than analyse data: they use it to influence a physical process.

Embodied intelligent robots are important examples, but industrial physical AI is broader than humanoid robots. It also includes mobile robots, machine vision systems, intelligent production equipment, autonomous vehicles and other autonomous systems capable of real-time interaction with the industrial environment.

The industrial site is subject to multiple constraints, including timing requirements, space, equipment, processes, materials, energy availability and efficiency, and the safety of people, equipment and operations. Just because a model can complete tasks under experimental conditions does not mean that the entire system has reached a production-ready level. Industrial physical AI is not simply about connecting robots to large models, but rather establishing a complete closed loop between perception, decision-making, control, feedback, and learning.

1.2 The verification value in industrial scenarios

Industrial production tasks can typically be broken down into specific processes, and their effectiveness can be measured through indicators such as quality, takt time[1], success rate, safety, and cost. Industrial scenarios can therefore simultaneously undertake functions such as technology verification, production quality validation, data production, and business value verification.

Industrial physical AI is not a simple replacement for traditional automation. Traditional automation has accumulated long-term experience in industrial control, functional safety, and engineering governance; artificial intelligence, on the other hand, enhances the system's capabilities in perception, reasoning, task planning, and environmental adaptation. To integrate industrial physical AI into production environments, it is necessary to combine the two, improving production flexibility while maintaining safety and predictable operation.

Early explorations of industrial embodied intelligence typically focused on whether robots could complete tasks such as grasping, handling, sorting, assembling, or inspecting. These verifications can prove technical feasibility, but a single success is not sufficient to demonstrate production value. The real breakthrough lies in enabling complete systems to complete valuable production tasks in a long-term, safe, stable, and economical manner, and gradually replicate from one workstation to more production lines and factories.

[1]Takt time: the maximum allowable time interval between completion of one unit of production to meet customer demand. It is used to synchronize production pace with customer demand rate.

02 Core Challenges of Scalable Implementation

Once industrial physical AI enters real production environments, it needs to complete perception, decision-making, and action in continuously changing working conditions, while meeting production, safety, cost, and responsibility management requirements. The main bottleneck of scalability is usually not whether individual technologies are available, but whether different technologies can be effectively integrated and continuously operated.

2.1 Data and model closed loop

Industrial physical AI relies on robots to generate and utilize multimodal data in real-world environments, encompassing images, positions, actions, torques, equipment status, task outcomes, and records of human intervention. This data is often closely tied to specific equipment, workstations, materials, and process conditions, with high acquisition costs and limited direct reuse.

Data scale does not directly equate to data value. Data that can be used for training, evaluation, and continuous optimization should have context related to tasks, time, environment, and version, and be able to trace execution results. Failed samples and abnormal records are equally important, as production systems not only aim to improve average success rates but also to identify low-frequency risks that may have significant impacts.

From training to production deployment, the model needs to undergo evaluation, verification, controlled release, and production observation. Data, equipment, and model versions should be traceable and have clear scope of application and rollback conditions. Cross-enterprise projects should also clarify data ownership, authorization scope, and reuse conditions at an early stage.

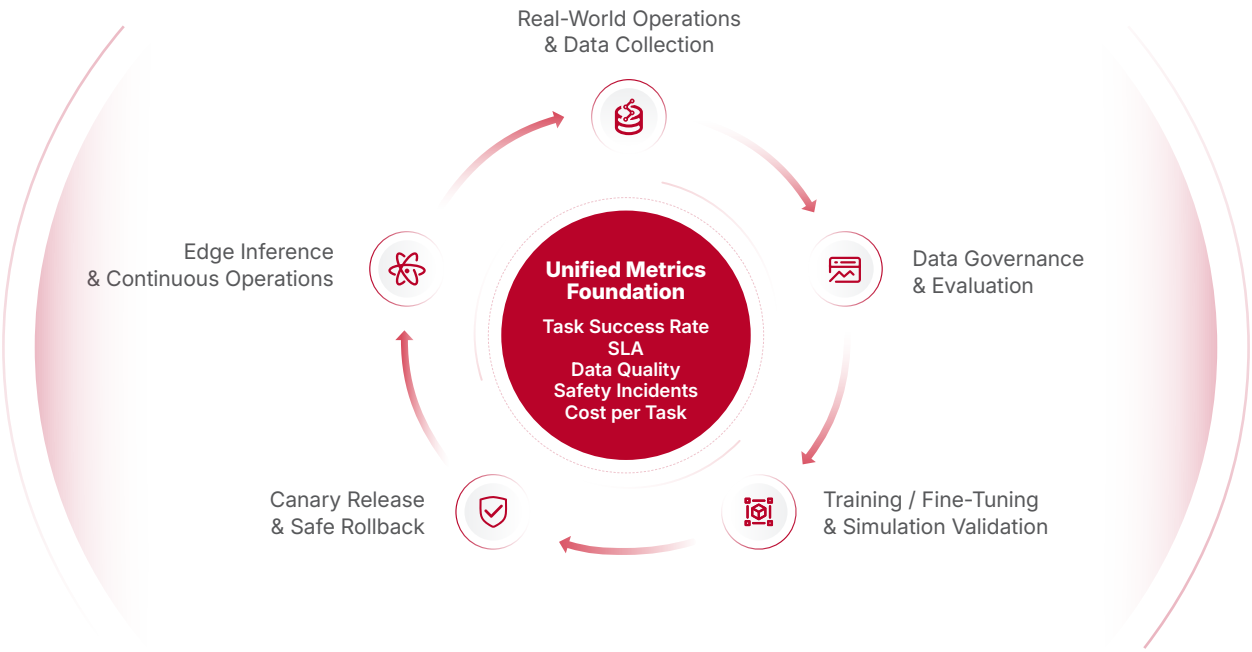

Figure 2-1 Data, Model, and Production Operation Closed Loop of Industrial Physical AI

2.2 Connectivity and computing collaboration

Different tasks have varying requirements for connectivity and computation. Safety protection and high-frequency motion control emphasize determinism and local availability; complex perception and multi-robot collaboration may require shared edge computing power; large-scale training and cross-scenario analysis are more suitable for completion in private clouds or centralized computing resources.

Connectivity and computing power ultimately need to serve complete production tasks. Therefore, when evaluating, it is important not to only compare peak rates, single-time delay, or chip computing power, but also to focus on the end-to-end performance from data collection, transmission, and processing to task execution, as well as the recovery and degradation capabilities under abnormal conditions.

2.3 Industrial system integration

Real factories typically require industrial physical AI systems to adapt to existing production systems. Robots and their connectivity, computing, and control systems need to coordinate with production equipment, industrial control systems, manufacturing execution systems, warehousing systems, quality management systems, and safety systems, and comply with existing process flows and responsibility boundaries.

Industrial equipment has a long lifecycle, and new models and software need to operate alongside equipment from different eras, suppliers, and protocols. Completely replacing existing systems often incurs high costs and may increase production risks. Therefore, open interfaces, modular deployment, and phased transformation are generally more feasible than reconstructing the entire technical system at once.

Industrial automation and digital twin capabilities can assist in defining process and safety constraints, supporting simulation verification, and facilitating integration with real production lines. Robots, networks, computing systems, models, and industrial systems must also adopt unified task identifiers, time benchmarks, and operational logs to support cross-system fault localization.

2.4 Safety and commercial feasibility

Industrial physical AI requires the simultaneous management of personnel and equipment safety, network and data security, model update risks, as well as supply chain and remote operation and maintenance risks. Security management should be integrated throughout the entire process of requirements, development, testing, deployment, operation, and updating.

Technical reliability does not equate to commercial viability. Projects require a comprehensive evaluation of production performance, continuous operation, human intervention, and risk mitigation, as well as the calculation of full lifecycle costs including on-site modification, connectivity, computing power, software, data, and operation and maintenance. These costs should be compared on an equal footing with those incurred through manual or traditional automation methods. Only when the production value, risk, and total cost are jointly verified can technological achievements be transformed into large-scale investments.

03 Connectivity and Edge Intelligence Architecture

Industrial physical AI requires a reasonable allocation of tasks between the physical device, such as a robot, on-site network, edge computing, and central cloud. The focus of architecture design is not to concentrate all data and computation on a single layer, but to establish a reasonable division of labour based on security, real-time performance, data governance, and cost requirements.

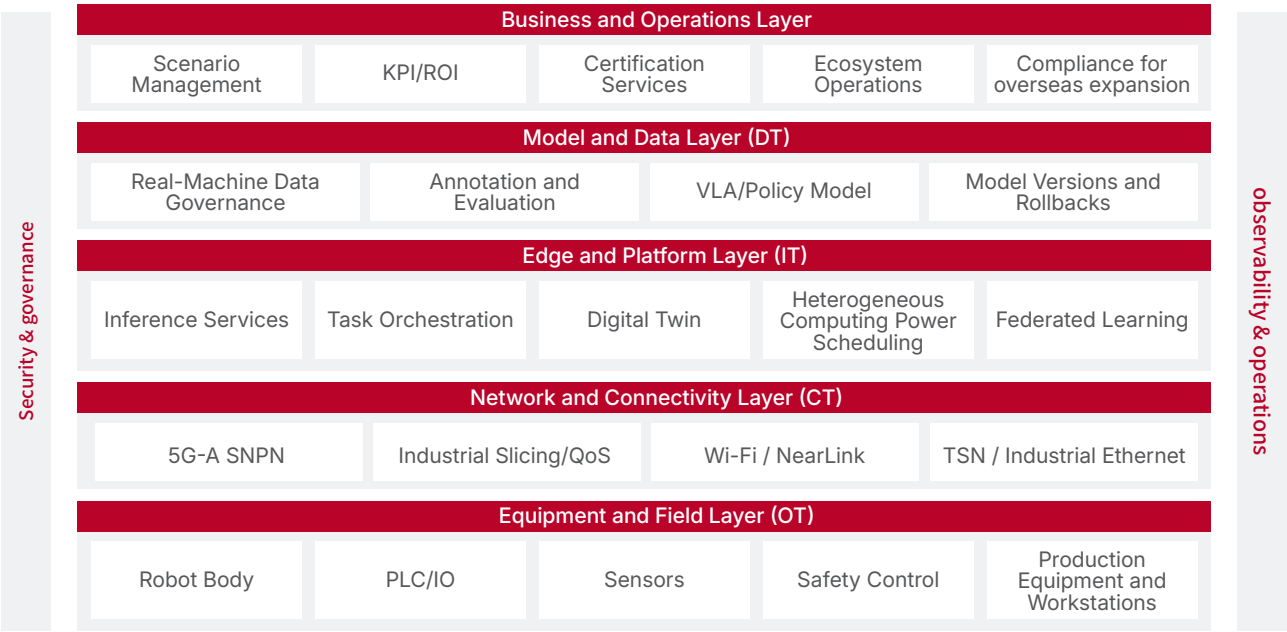

Figure 3-1 Reference Architecture of Industrial Physical AI System

3.1 Task division and network coordination among endpoints, edges, and clouds

Industrial physical AI requires the allocation of computation and control tasks among the robot body, on-site edge, and central cloud, and achieves coordination of data, models, and operational status through networks. Specific deployment should be determined based on the requirements of tasks for safety, real-time performance, computing power, data governance, and cost.

The robot body and on-site control system should prioritize emergency stop, motion control, state estimation, and other safety-critical and high-real-time tasks. When anomalies occur in the network, edge nodes, or upper-level software, the robot should be able to safely shut down or switch to a pre-set limited operation mode.

The on-site edge is suitable for undertaking tasks such as multimodal data preprocessing, complex perception and reasoning, multi-robot task scheduling, and shared maps. The value of edge computing lies not only in providing more computing power, but also in heterogeneous orchestration of different types of computing resources based on perception,

reasoning, planning, and control tasks, as well as supporting resource isolation, version management, and operational monitoring.

Private clouds or centralized computing resources are more suitable for undertaking model training, large-scale simulation, cross-scenario analysis, and global model management. The decision on whether to deploy specific tasks on the terminal, edge, or cloud should be inferred from production tasks, rather than solely based on the peak performance of a single computing node.

Location	Tasks suitable for undertaking	Main constraints	Basic Principles
Robot body and on-site control system	Emergency stop, motion control, state estimation, safety-critical tasks	Real-time performance, power consumption, functional safety	Maintain safety when the network connection is lost or the upper-level system malfunctions
On-site edge	Multimodal data preprocessing, complex perception and reasoning, multi-robot scheduling, shared map	Task latency, computing power competition, version management	Support resource isolation, priority management, and operation monitoring
Private cloud or centralized computing resources	Model training, large-scale simulation, cross-scenario analysis, and global model management	Data governance, bandwidth, cost	Determine the deployment location based on production tasks and data requirements

Table 3-1 Task Division of Industrial Physical AI: Endpoint, Edge, and Cloud

3.2 The role of 5G and 5G-Advanced

5G and 5G-Advanced can play a significant role in continuous connectivity for mobile operations, multi-robot collaboration, multi-modal data upload, edge intelligence access, and unified connection management.

Autonomous mobile robots may traverse workstations, production lines, and storage areas, necessitating business continuity during their movement; multi-robot collaboration requires the exchange of task status, maps, and scheduling information; video, point cloud, and equipment operation data may also generate continuous uplink demands. The combination of 5G and edge processing can strike a balance between data transmission, on-site preprocessing, and event-triggered uploads.

3.3 Synergy of multiple technologies

Industrial Ethernet, Wi-Fi/NearLink, 5G and 5G-Advanced, and short-range wireless technologies will coexist in the industrial field for a long time. Fixed equipment, safety-critical control, mobile operations, and internal connections of robots have different requirements, and no single technology can meet all tasks at optimal cost.

The selection of connectivity technology should be based on factors such as device mobility, data volume, real-time requirements, security standards, coverage range, and lifecycle costs. The relevant architecture should also integrate tasks, data, models, and operation and maintenance information, and support anomaly localization through unified identification, time benchmarking, and version recording.

Looking ahead, we can further identify the requirements of 6G in terms of Integrated Sensing and Communication (ISAC), network intelligence, and multi-robot collaboration based on the data accumulated from real-world scenarios of 5G-Advanced.

04 A Proposed Maturity Model for Production-Ready Industrial Physical AI

The value of industrial physical AI cannot be judged solely by whether a single demonstration is successful, but rather depends on whether the system can consistently, safely, and economically complete tasks under real production conditions.

Drawing on the contributors' practical perspectives, this paper presents a proposed five-stage maturity model as a practical project-management tool for industrial physical AI projects. It provides a common structure for defining objectives, aligning project participants, planning evidence and testing requirements, assessing progress, and considering readiness for production deployment and replication. The model is intended as practical guidance rather than a formal national or industry standard, and it does not correspond to automation levels used in domains such as autonomous driving.

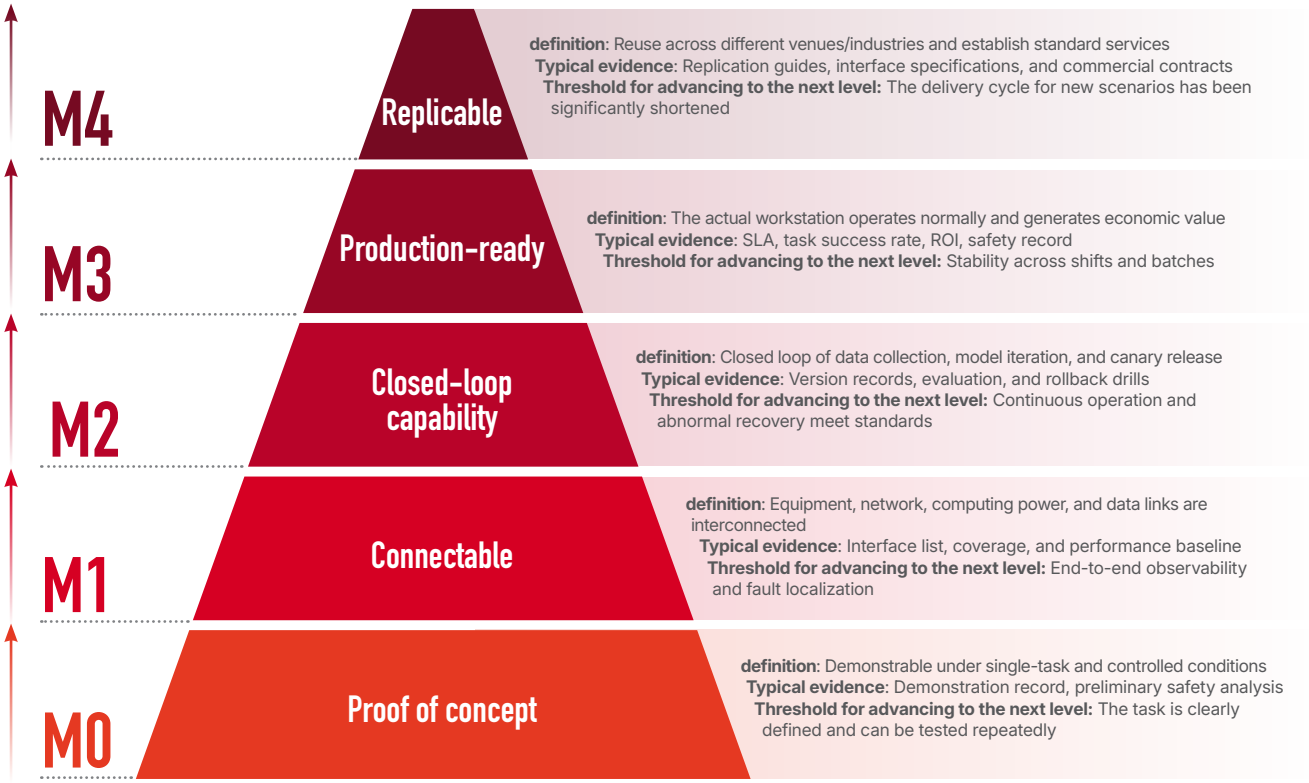

Table 4-1 Proposed Maturity Model for Industrial Physical AI Projects

4.1 M0: Proof of concept

The system is capable of completing single or limited tasks under controlled conditions, providing preliminary evidence that basic technological paths such as robotics, modeling, connectivity, and computing are feasible.

At this stage, it is necessary to clarify the task definition, testing conditions, equipment and model versions, sample size, and failure scenarios, to avoid directly equating a single successful demonstration with capability maturity.

4.2 M1: Connectable

Robots, networks, computing power, data, and necessary industrial system interfaces can be connected, and basic end-to-end data and task links can operate.

In this stage, not only should it be proven that the equipment can be networked, but also basic identity, interface, time synchronization, and operation records should be established, so that tasks, networks, and computing states can be observed, and anomalies can be preliminarily located.

4.3 M2: Closed-loop capability

The system is capable of forming a basic closed loop encompassing data collection, model evaluation, controlled release, production observation, and exception rollback. It can also repeatedly execute tasks and handle common changes within a certain timeframe.

In addition to task success rate, attention should also be paid to continuous operation, manual takeover, abnormal recovery, and safe degradation capabilities. When anomalies occur in the network, edge nodes, or upper-layer software, the system should enter a safe state or degrade its operation according to a predetermined method.

4.4 M3: Production-ready

The system is capable of continuous operation in real workstations and meets specific requirements for quality, takt time, safety, cost, and operation and maintenance.

The evaluation should encompass robots, scenario modification, connectivity, computing power, software, data, system integration, and operation and maintenance, rather than just the robot itself. The project should also be compared with manual or traditional automation methods to demonstrate its production value and commercial feasibility.

4.5 M4: Replicable

The system can be migrated to new workstations, production lines, or factories with relatively predictable costs and timelines, and develop reusable scene descriptions, interfaces, testing, deployment, and operation and maintenance methods.

If redevelopment and integration are required for each new scenario, the project remains a highly customized endeavour rather than a scalable and generalizable solution.

Production-level evaluation should cover task performance and manual takeover, continuous operation and abnormal recovery, end-to-end connectivity and computational performance, data and model traceability, as well as safety, unit task cost, and replication cycle. All results should specify test conditions, sample size, system version, manual intervention methods, and applicable boundaries.

Dimension	First-level indicator	Explanation
Task	Success rate, takt time, quality, manual takeover rate	With the completion of business tasks as the ultimate indicator, it is necessary to disclose the sample size and operating conditions
Continuity	Continuous operation duration, MTBF, recovery time	Distinguish between planned downtime, abnormal downtime, and external causes
Network	End-to-end latency quantile, jitter, availability, and upstream throughput	Make statistics separately based on business flow and coverage area
Computing power	Inference latency, throughput, resource utilization, and energy efficiency	Test with actual model and input scale
Data	Efficiency, completeness rate, synchronization error, reuse rate, unit cost	Simultaneously, the failure samples and labour input are recorded
Safety	Dangerous events, protection triggers, patch timeliness, and rollback success rate	Zero incidents do not equate to zero risk; exposure needs to be recorded
Economy	Single task cost, TCO, payback period, utilization rate	Comparison with manual/traditional automation baseline on an equal footing

Table 4-2 Evaluation Dimensions for Applying Maturity Model

05 Ecosystem Collaboration: A Proposed Implementation Blueprint for Qiantang

This chapter sets out a proposed implementation blueprint for Hangzhou Qiantang, China, bringing together regional scenario coordination, connectivity, computing, industrial automation, robotics and real-world data production. This chapter connects the capabilities and architecture discussed in Chapters 2 and 3 with the proposed maturity model in Chapter 4, explaining how partners, scenarios, validation activities and reusable outputs could be organised within a regional implementation programme.

5.1 From joint publishing to joint verification

The proposed collaboration model brings together complementary capabilities: regional scenario coordination and public services; industrial 5G and edge integration; heterogeneous computing and AI deployment; industrial control and digital twins; continuous robot operation; and real-world robot data production and quality. These capabilities are intended to support shared tasks, metrics, testing and delivery rather than joint announcements alone. Partners would need to agree on data use, software and model versions, interfaces, testing, publication of results, intellectual property and accountability before joint implementation.

5.2 The value of Qiantang's regional scenario platform

Qiantang's manufacturing base, including automotive and components, life sciences and intelligent equipment, provides settings in which industrial needs could be matched with technical capabilities.

The blueprint proposes a networked arrangement combining an operations hub, training and data collection, simulation and controlled validation, and access to real industrial workstations. Regional bodies would coordinate scenarios, resources and public services, while technology and industrial partners would contribute connectivity, computing, robotics, control, data and integration capabilities. This structure is intended to connect controlled trials with sustained production rather than concentrate all activities in a showcase facility.

5.3 From scenario verification to reproducible outcomes

The proposed initial scenarios cover automotive and components, life sciences, equipment manufacturing, real-world robot data production, precision assembly, and hazardous or specialised operations. Implementation would start with two or three clearly bounded workstations to establish a minimum operational loop, then extend to multi-robot and cross-workstation operation before developing replication packages. Each project should produce reusable scenario templates, data rules, interface specifications, testing methods and deployment guidance. Progress should be assessed through sustained operational performance, verified business value, continued enterprise participation and replication across scenarios.

The proposed M0-M4 model helps project teams identify the validation activities and evidence needed at each stage: M0 addresses task feasibility under controlled conditions; M1 connectivity and observability; M2 the data-model-operation loop and safe recovery; M3 sustained production performance and commercial viability; and M4 replication across workstations, production lines or factories. This proposed implementation approach is intended to reduce project risk, strengthen evidence from trials and support repeatable deployment.

06 Practical Steps for Project Teams

Drawing on the practical perspectives of the contributors and the approach discussed in this paper, the following steps could help organisations undertaking industrial physical AI projects. They provide a starting point for use-case selection, technical design, testing, governance and capturing reusable results, rather than formal GSMA recommendations.

01 Start with tasks where the production value is measurable, on-site conditions are adaptable, data is accessible, and risks are controllable, and initially form a minimal verification closed loop.

02 Determine the configuration of robots, networks, computing power, and data based on production, safety, and economic goals, avoiding the pursuit of extreme parameter values without considering the specific scenario.

03 Establish an open architecture with a reasonable division of labour among the terminal, edge, and cloud, enabling key components to be tested, upgraded, or replaced.

04 Unify task definitions, testing conditions, data, and statistical standards, and clarify data usage, version changes, result publishing, and responsibility for anomalies.

05 Distil the scenario templates, interface specifications, testing methods, and deployment experiences formed by the project into replicable outcomes, laying the foundation for subsequent standardization and cross-scene promotion.

Conclusion

The key to the scaling of industrial physical AI lies not in enabling robots to complete a more complex demonstration, but rather in forming a closed loop among data, connectivity, computing, industrial control, safety, and operations, continuously creating value in real-world production.

5G and edge intelligence can expand mobile operations, multi-robot collaboration, and on-site data processing capabilities, but their value ultimately needs to be verified through production results. The proposed Qiantang blueprint illustrates how complementary capabilities could be organised around shared scenarios, evaluation methods and a phased implementation roadmap to support the transition from proof of concept to production application.

Looking ahead to the next stage, all parties involved in the industry need to be driven by real production tasks, establish a collaborative mechanism centered around scenario openness, joint verification, hierarchical evaluation, and continuous operation, and turn the verification results into operational, comparable, and reproducible methods, promoting the large-scale application of mature capabilities in more factories and production scenarios.





GSMA
1 Angel Lane
London
EC4R 3AB
United Kingdom

Tel: +44(0)20 7356 0600
Fax: +44(0)20 7356 0601