





INDUSTRIAL PHYSICAL AI IN HIGH-PRECISION MANUFACTURING

From Digital Engineering to Deployment Readiness

Main Contributors



Deloitte

SKAI
INTELLIGENCE

FOREWORD

In selected precision manufacturing scenarios where multi-variety production, small-batch manufacturing, high precision, rapid changeover, and quality consistency requirements coexist, manufacturers are entering the engineering exploration stage of Physical AI. As product iteration accelerates, production batches become more flexible, product forms and process requirements continue to evolve, and quality consistency requirements remain high, robotic systems are no longer expected merely to execute fixed paths. They increasingly need to perceive, locate, reason, and act in real physical environments.

In this context, robotic vision serves as an important bridge between AI perception and physical action. A robot needs to understand the position, pose, state, and task environment of a target object before it can reliably support industrial actions such as grasping, assembly, inspection, or handling. In relevant precision manufacturing scenarios, the challenge for robotic vision is not only about model performance. It also involves task definition, real-data coverage, vision variables, engineering validation, site feedback, and enterprise-level governance.

Traditional robotic vision engineering processes remain valuable in structured scenarios where product categories are stable, and process variation is limited. However, under new requirements such as rapid changeover, multi-variety production, high precision, excessive reliance on late-stage site validation, and concentrated real-data collection may lead to delayed validation windows, insufficient coverage of boundary cases, concentrated site-tuning pressure, and limited reuse of engineering experience.

Therefore, for precision manufacturing tasks in which product forms change quickly, real-data collection windows are limited, robotic actions depend heavily on vision-based judgment, and there is a need for future multi-station replication, Physical AI is becoming an engineering direction worth exploring.

This white paper summarizes this direction as Robotic AI Digital Engineering. Through Digital Twin, Synthetic Data, robotic vision data assets, and the Real2Sim2Real loop, part of the robotic vision validation work can be advanced to the digital engineering stage, while real validation and real feedback remain essential for deployment judgment.

The engineering implementation framework described in this white paper refers primarily to an industry-level methodology, POC organization process, Deployment Readiness assessment, and enterprise governance framework for manufacturing enterprises. It is not a station-level technical manual, system integration solution, or production acceptance standard. In actual projects, enterprises should conduct engineering validation based on their own products, tolerance requirements, process flows, equipment conditions, data availability, and site constraints.

This white paper proposes an industry-level engineering implementation framework. It begins with industry background, analyzes limitations in traditional processes, proposes a new paradigm and a proof of concept (POC) process, illustrates the methodology through a representative case, and further discusses the engineering expression of vision variables, enterprise-level governance, the business value framework, the implementation path, and industry collaboration.

This white paper emphasizes that digital-first does not mean replacing real validation, Synthetic Data does not mean replacing real data, and Physical AI is not an inevitable direction for all precision manufacturing scenarios. For scenarios where products are stable, workpieces are strongly constrained by tooling, cycle time is highly fixed, and rule-based vision can meet requirements, traditional automation, dedicated machines, fixed tooling, and rule-based vision may still offer stronger stability and better economics. The engineering implementation of Physical AI should proceed within an open, safe, cautious, and governable framework.

JOINT PUBLISHER

JOINT FOREWORD PERSPECTIVE / ABB ROBOTICS AND NVIDIA:

Robots need to see, and manufacturing needs earlier validation. Robotic vision is becoming a critical entry point for Physical AI in high precision manufacturing.



CRAIG MCDONNELL
Business Line Managing Director
Industries, ABB Robotics

“Manufacturers are no longer asking whether robots can move, they are asking whether robots can understand and adapt to changing production environments. The next stage of industrial automation requires connecting robotic vision, digital engineering, and real-world execution into a repeatable engineering process. Our goal is to help manufacturers identify risks earlier, validate robotic tasks before deployment, and build scalable automation capabilities with confidence.”

ABB Robotics



DEEPU TALLA
Vice President and General
Manager, Robotics & Edge AI
NVIDIA Robotics

“Physical AI is transforming how intelligent systems are developed and deployed in the physical world. By combining simulation, synthetic data, accelerated computing, and AI models, manufacturers can validate more scenarios, cover more edge cases, and accelerate innovation before deployment. The opportunity is not simply to build better models, but to create a continuous learning loop between the digital and physical worlds.”

 **NVIDIA**

MAIN CONTRIBUTORS



DR. OUYANG Ye
Chief Executive Officer, Chief
Technology Officer, AsiaInfo
Technologies, IEEE Fellow



“For Physical AI, being verifiable, reusable, governable, and scalable is not just a vision – it is a fundamental requirement for moving from the lab to the factory floor. At AsiaInfo, we are leveraging the Physical AI OS to seamlessly connect data, training, and execution in a closed-loop engineering workflow, transforming industry-specific physical expertise into deployable, intelligent engineering capabilities.”



NITIN MITTAL
Deloitte Global AI Leader

Deloitte.

“Deloitte’s strategic leadership in enterprise AI transformation emphasizes governable, business-led adoption of Physical AI, where value, risk, and scalability are aligned with corporate strategy and operating model innovation. This white paper reflects our dedication to advancing responsible, ROI-driven industrial AI that creates sustainable competitive advantage.”



MORGAN MAO
Global CEO
SKAI Intelligence

SKAI
INTELLIGENCE

“Our role is to make robotic vision evidence usable beyond the vision team. When visual variables, edge cases and real-world feedback are structured with traceability, they become a shared language for AI validation, robot execution, governance and ROI review – helping manufacturers decide what is ready, what needs refinement, and what should pause.”

TABLE OF CONTENTS

08	Executive Summary: Moving Vision Validation Toward Deployment Readiness	23	KPI, ROI and Deployment Readiness
10	Why Physical AI Now	26	Implementation Roadmap and Ecosystem Collaboration
12	Digital Engineering Reference Architecture and Loop Robotic Digital Engineering	32	Appendix A Glossary
14	Robotic Digital Engineering Verification	33	Appendix B Case A: 3DoF Localization POC Evidence Package
16	Synthetic Data and the Real-World Feedback Loop	34	Appendix C Case B: Offline 6D Pose Validation Evidence Package
18	Application Scenario Matrix	35	Appendix D Reference Figures and Usage Notes
20	Case Evidence Summary: 3DoF and 6D Pose	36	Appendix E Key Takeaways
		37	Appendix F Contributors

Executive Summary

Moving Vision Validation Toward Deployment Readiness

Autonomous Versatile Robotics not only need to move. They need to see and decide. According to Deloitte’s report, Physical AI: The moment of acceleration, 41% of organizations cite cost and resource requirements as the primary barrier to Physical AI adoption¹, highlighting the growing need for scalable validation and deployment readiness. The automation bottleneck in high precision manufacturing is gradually shifting from robot hardware and fixed-path execution toward visual perception, data coverage, validation timing, and deployment readiness. In scenarios where multi-variety production, small batches, high precision, and rapid changeover requirements coexist, robots need to understand the position, pose, state, and task environment of target objects before AI perception can be translated into executable industrial actions.

But in the real world, parts shift, lighting, materials and camera views all vary. These parameters are critical; small variations can have a big impact.

The core method introduced in this white paper is to move part of robotic vision risk identification into the digital engineering stage through digital twins, task-relevant synthetic data, AI validation, robotic digital engineering verification, and real-world feedback loops. The real world defines the problem, the digital environment expands coverage, and real-world feedback continuously calibrates the loop. This creates


traceable, verifiable, and reusable data assets and engineering judgments.

The value should not be measured by model accuracy alone. More important questions include whether vision risks are identified earlier, whether edge cases are covered, whether real-world feedback enters the loop, whether robotic tasks are ready to enter Pilot Cell validation, and whether data assets can be reused across products, cells, and projects.

This white paper also emphasizes clear boundaries: digital-first validation does not replace real-world validation; offline model metrics do not equal production capability; and proof of concept (POC) success does not equal deployment success. It is not a commitment to any specific production deployment result, but a framework to help manufacturers assess Physical AI deployment readiness.

For manufacturers, a practical starting point is to select vision localization, inspection, or offline 6D Pose evaluation scenarios with clear business value, low safety risk, well-defined task boundaries, accessible digital assets, and feasible real-reference validation.

A staged path is recommended:
Digital POC → Pilot Cell → Multi-cell Scale-up.



¹ Deloitte, Physical AI: The moment of acceleration, June 2026.

WHY NOW


Visual perception, data coverage, validation timing and deployment readiness are becoming new automation bottlenecks.

METHOD

Digital Twin → Synthetic Data → AI Validation → Robotic Digital Engineering → Feedback

DECISION PATH

Digital POC → Pilot Cell → Multi-cell Scale-up



Why Physical AI Now

In high precision manufacturing, high-mix production, frequent new product introductions, and multiple SKUs running on the same line make traditional automation difficult to scale. At the same time, the production environment itself is not stable: lighting conditions vary, parts come in different shapes and tolerances, and their positions are not always fixed.

Every change requires reprogramming, re-teaching, and significant engineering effort, meaning automation becomes costly and difficult to scale.

The challenge is not that traditional automation does not work. The challenge is that it does not adapt fast enough.

High precision manufacturing is therefore moving from fixed automation toward perception-driven automation. In scenarios where multi-variety production, small batches, high precision, and rapid changeover requirements coexist, robots are no longer expected only to execute fixed paths. They need to understand the position, pose, state, and task environment of target objects.

Robotic vision has therefore become a critical entry point for Physical AI. Only when robots can reliably see and understand target objects can AI perception be translated into executable industrial actions such as grasping, assembly, inspection, and handling.

Traditional robotic vision engineering workflows remain valuable, especially in scenarios where product variation is limited, process boundaries are clear, and site conditions are stable. However, in high-variation precision manufacturing environments, traditional workflows often rely on physical samples, real-image collection, and late-stage on-site tuning. Many vision risks are therefore exposed too late in the engineering process.

The issue with traditional workflows is not lack of value, but late validation. The value of digital-first validation is to move part of vision risk identification into the digital engineering stage.


Instead of trying to control every detail of the environment, we enable the system to adapt to variability.

Through digital twins, task-relevant synthetic data, and early AI validation, manufacturers can cover visual variables, edge cases, and potential failure modes more systematically before real deployment. This makes subsequent real-world validation more focused and more traceable.



Validation moves Earlier

From Fixed Automation to Physical AI



The goal is not to skip real validation, but to identify risk earlier.

01 From fixed automation to physical AI - moving the validation window earlier

FROM LATE-STAGE PHYSICAL VALIDATION TO DIGITAL-FIRST VALIDATION	
Traditional Workflow	Digital-First Workflow
Design → Physical Samples → Real Image Collection → On-site Tuning	Design → Digital Twin → Task-Relevant Synthetic Data → Early Validation → Real-world Feedback Calibration
Risks often surface late on site	Risks enter digital engineering review earlier
Depends on samples, workcells and on-site windows	Enables partial validation before all physical conditions are ready
Higher pressure during late-stage tuning	Real-world validation becomes more focused and targeted

Digital Engineering Reference Architecture and Loop


From Vision Validation to Deployment Readiness

The deployment of Physical AI in high precision manufacturing is not about directly placing an AI model into a factory. It is not a single software, robot, or dataset problem. Autonomy is not just about intelligence; it also depends on accurate modelling, training, and validation in the physical world. It requires an engineering loop that connects manufacturing tasks, digital assets, synthetic data, AI validation, robotic digital engineering verification, and real-world feedback.

In this loop, the real world defines the problem, the digital environment expands validation coverage, AI models are validated earlier with task-relevant data, and robotic digital engineering environments further assess whether vision outputs can enter robot targets, paths, and station logic. Pilot Cell validation and real-world feedback then continuously calibrate digital twins, data assets, model versions, and deployment decisions.


For management teams, the key output of this loop is not a single model accuracy number. It is deployment readiness evidence. Manufacturers need to assess whether vision risks are identified earlier, edge cases are covered, model outputs can support robotic tasks, real-world feedback can update data assets, and the capability can be reused across more workcells.

Loop path:
Manufacturing Tasks & Real-World Inputs → Digital Twins → Task-Relevant Synthetic Data → AI Validation → Robotic Digital Engineering Verification → Pilot Cell / Deployment Readiness → Real-World Feedback → Digital Asset Update



Physical AI Digital Engineering Loop

From Vision Validation to Deployment Readiness




The real world defines the problem, the digital environment expands coverage, and real-world feedback continuously calibrates the loop.

02 Physical AI digital engineering reference architecture and loop

SEVEN CAPABILITY LAYERS IN THE REFERENCE ARCHITECTURE

Capability Layer	Management Question
Manufacturing task and process constraints	What task needs to be solved, and what are the quality and cycle-time boundaries?
Digital assets and digital twins	Which product, work cell and robotic objects need to be digitally represented?
Task-relevant synthetic data	Which visual variables, edge cases and failure modes need early coverage?
AI model and vision validation	Can the model support task-relevant perception?
Robotic digital engineering verification	Can vision outputs enter robot targets, paths and station logic?
Pilot Cell and real-world feedback	Can digital validation be calibrated under real or near-real conditions?
Governance and scale-up	Are data, models, workflows and responsibilities traceable, reusable and auditable?



Robotic Digital Engineering Verification

From Vision Outputs to Robotic Tasks

Robotic vision cannot exist in isolation from robotic execution. A vision model may complete image-level recognition, localization, or pose estimation, but this does not necessarily mean that its outputs can be reliably translated into robot target points, paths, station logic, cycle-time constraints, and quality loops.

Therefore, after AI model validation, project teams need to verify whether vision outputs are executable and deployment-ready within a robotic digital engineering environment. For management teams, the question is not only whether the model recognizes correctly, but whether model outputs can support real robotic tasks.

Within ABB Robotics’ robotic digital engineering system, RobotStudio® HyperReality can serve as an engineering implementation example at this layer. It connects robotic digital workcells, vision data, AI validation, robotic task verification, and physical deployment preparation. Its role is not to replace real-world tuning, but to help engineering teams evaluate whether vision outputs can support robotic task execution before on-site commissioning.


In this workflow, NVIDIA’s simulation, rendering, and AI computing ecosystem can support task-relevant digital scene representation and AI validation paths. Task-relevant synthetic data and auto annotation then serve as inputs for AI validation and robotic task verification.

THREE MANAGEMENT QUESTIONS

Question	Management Meaning
Can vision outputs be translated into robot target points?	Determines whether AI perception can enter robotic task definition
Are target points and poses executable?	Assesses whether path, reachability, station logic and cycle time can be validated
Does the result support Pilot Cell readiness?	Determines whether the project has evidence for the next validation stage


From Vision Outputs to Robotic Tasks

Robotic Digital Engineering Verification



Model recognition is only the starting point; deployment readiness depends on robotic execution.

03 From vision outputs to robotic tasks



Synthetic Data and the Real-World Feedback Loop

From Visual Variables to Deployment Readiness

Synthetic data should not be understood simply as generating more images. For robotic vision, what matters is whether the data covers task-relevant variables and whether it supports model validation, robotic task verification, and deployment readiness decisions.

In high precision manufacturing scenarios, many vision risks are difficult to collect sufficiently from real data at an early stage. These include reflection, glare, low contrast, side lighting, pose shift, partial occlusion, surface interference, batch variation, and low-frequency abnormal samples. The value of synthetic data is to systematically express these variables that are difficult to capture early in real environments and convert them into traceable, verifiable, and reusable data assets.

These data should not be treated as arbitrary image augmentation. They should be constructed around manufacturing tasks, vision risks, and engineering constraints. An effective synthetic data system should record variable sources, version changes, validation results, and failure samples, so that the data can support model iteration, robotic digital engineering verification, and Pilot Cell review.

In this white paper, this continuous calibration mechanism between the real world, the digital environment, and real-world feedback is referred to as Real2Sim2Real.

Real2Sim uses real-world observations and operational data to continuously improve the digital environment. Sim2Real transfers validated models and policies developed in simulation to physical systems for deployment. Together, they form a closed-loop engineering workflow that continuously improves both simulation fidelity and real-world performance.

Its core logic is: the real world defines the problem → the digital environment expands coverage → real-world feedback continuously calibrates the loop. Real-reference images, failure cases, and site feedback can continuously update digital twins, synthetic data strategies, model versions, and deployment decisions. In this sense, synthetic data is not only a model-training input. It becomes a reusable engineering asset in robotic vision digital engineering.

THREE MANAGEMENT-LEVEL VALUES OF SYNTHETIC DATA

Value Dimension: Earlier coverage of risk variables.

Description: Builds validation scenarios for reflection, low contrast, occlusion, pose variation and edge cases before physical samples or workcells are fully ready.

Value Dimension: Support for deployment readiness review.

Description: Helps assess whether model outputs are ready to enter robotic task verification and Pilot Cell validation.

Value Dimension: Reusable data asset creation.


Description: Converts variables, versions, failure samples and real-world feedback into traceable and reusable data assets.




DIGITAL STATION / ROBOTSTUDIO SIMULATION
Robot / station / controller logic



PRODUCT & VISION DIGITAL TWIN
Photorealistic assets and sensor-ready scenes



PHYSICAL ROBOT / REAL-WORLD DEPLOYMENT
Production execution and feedback




04 Real2Sim2Real Architecture for Robotic AI Deployment

Application Scenario Matrix

Where to start Digital Proof of Concept (POC)

Physical AI should not be understood as a universal solution for all manufacturing environments. Manufacturers should select Digital POC scenarios based on business value, task boundaries, vision challenges, digital asset availability, and real-world validation conditions. When selecting the first scenario, management teams should prioritize clear task definition, controllable risk, and assessable value.

Application Direction	Typical Tasks	Key Vision Capabilities	Physical AI Value	Caution Conditions
High precision part localization	Loading, assembly, pre-inspection localization	2D / 3DoF localization, center-point and rotation-angle estimation	Early validation of localization stability and angle robustness	Extremely narrow process tolerance and insufficient real-reference validation
High-mix assembly	Multi-SKU flexible assembly, frequent changeover cells	Pre-grasp localization, part-state recognition, pose confirmation	Reduced changeover validation and lower on-site trial uncertainty	Fixtures and process definitions are frequently redesigned
Visual inspection and recheck	Surface quality inspection, assembly confirmation, NG recheck	Defect detection, abnormal sample recheck, edge-case recognition	Edge-case generation and failure-sample feedback loop	Defect mechanisms are difficult to represent digitally
Spatial grasping and assembly guidance	Robotic grasping, spatial assembly, pose guidance	6D Pose, RGB-D offline evaluation, spatial pose robustness assessment	Earlier screening of spatial localization and pose-estimation paths	Depth sensor stability or field-condition variance remains insufficient



PRECISION PART LOCALIZATION

Typical tasks: Loading / assembly / pre-inspection

Value: Validate localization stability early

Caution: Narrow tolerance without real-reference validation

2D / 3DoF / Rz

HIGH-MIX ASSEMBLY

Typical tasks: Multi-SKU assembly / changeover

Value: Reduce validation and on-site trial uncertainty

Caution: Frequent fixture or process changes

Multi-SKU / Changeover

VISUAL INSPECTION AND RECHECK

Typical tasks: Surface inspection / NG recheck

Value: Build edge-case and failure-sample loop

Caution: Defects hard to represent digitally

Inspection / NG

SPATIAL GRASPING AND ASSEMBLY GUIDANCE

Typical tasks: Grasping / assembly / pose guidance

Value: Screen pose-estimation paths earlier

Caution: Depth sensor stability or field variance risk

6D Pose / RGB-D

Vertical industry examples: Electronics assembly, Medical devices, Optical components, Automotive components, Semiconductor & precision equipment

The first Digital POC should prioritize clear tasks, assessable value and controllable risk.

05 Application Scenario Matrix

VERTICAL INDUSTRY EXAMPLES		
Vertical Industry	Suitable Initial Exploration Tasks	Caution Points
Electronics assembly	Small-part localization, assembly confirmation, pose guidance	High reflection, low texture, rapid changeover
Medical devices	Precision assembly, pre-inspection localization, quality recheck	Compliance requirements, traceability responsibility, quality boundaries
Optical components	Surface inspection, pose localization, pre-assembly confirmation	Transparent / reflective materials, imaging stability
Automotive components	Grasping, assembly, quality inspection	Size variation, cycle time, safety requirements
Semiconductor equipment / precision equipment	Spatial localization, assembly guidance, pose confirmation	Narrow tolerance and high real-world validation requirements

Case Evidence Summary: 3DoF and 6D Pose

Stage-gate evidence before Pilot Cell

This white paper uses two representative robotic vision validation cases to illustrate how digital twins, synthetic data, AI validation, and robotic digital engineering loops can support deployment readiness decisions. The two cases focus on different validation questions: the 3DoF Localization POC focuses on planar position and rotation-angle localization, while the Offline 6D Pose Validation focuses on spatial position and pose estimation.

For management teams, the key questions are whether synthetic data can support task-relevant variable coverage, whether model outputs can generate reviewable validation evidence, and whether the current evidence is sufficient for the next stage or still requires real-cell validation.

CASE A: 3DOF LOCALIZATION POC

This POC focuses on a precision-part localization task. The objective is to train a 3DoF localization model with synthetic data so that it can output the target object’s x–y planar position and Rz rotation angle.

Dimension	Executive Summary
Task	x–y planar localization + Rz rotation angle estimation
Data	1,000 synthetic samples covering background, exposure, field of view, surface interference and pose variation
Training samples	120 representative training samples selected from the synthetic dataset
Acceptance criteria	Absolute accuracy: x–y ≤ 2 pixels, Rz ≤ 0.2°; repeatability: x–y ≤ 1 pixel, Rz ≤ 0.1°
Result summary	In two real-reference validation sets, x–y localization met the pixel-level boundary; Rz achieved approximately 88% / 77% coverage under the stringent angular threshold, and approximately 92% / 84.7% under the near-threshold Rz ≤ 0.25° check
Abnormal samples	10 abnormal samples produced localization outputs acceptable under qualitative human review
Deployment judgment	Serves as stage-gate evidence for vision localization capability; before entering Pilot Cell, image-level error still needs to be mapped to process tolerance, robotic execution linkage and quality windows

CASE B: OFFLINE 6D POSE VALIDATION

This case evaluates how synthetic RGB-D data and FoundationPose can support 6D Pose technical path screening. Unlike 3DoF localization, 6D Pose estimation focuses on both position and orientation in 3D space.

Dimension	Executive Summary
Task	Spatial position and pose estimation of the target object
Data	Synthetic RGB-D data covering translation, rotation, depth variation, exposure variation and surface disturbance
Model	FoundationPose, used to evaluate the offline 6D Pose path under RGB-D input conditions
Result summary	Under controlled synthetic RGB-D offline evaluation conditions, the deviation between model outputs and synthetic ground truth remained within a low range, with no obvious outlier failure case observed
Boundary	Results do not represent real RGB-D sensor accuracy, real calibration-chain accuracy, robot TCP positioning, or production deployment capability
Next validation step	Before entering Pilot Cell, real workcell variables, sensor stability, calibration drift, occlusion, reflective materials, cycle-time constraints and robotic execution linkage still need validation



CASE A: 3DOF LOCALIZATION POC
Planar position + Rz angle estimation

1,000
synthetic samples

120
representative training samples

Acceptance criteria
x–y ≤ 2 pixels • Rz ≤ 0.2°

Stable x–y localization
Pixel-level boundary met in two real-reference sets

Strict angular-threshold coverage
Rz ≤ 0.2°: ~88% / ~77%

Near-threshold coverage
Rz ≤ 0.25°: ~92% ~84.7%

10 abnormal samples
acceptable under qualitative human review

Before Pilot Cell: Map vision error to process tolerance, robotic execution linkage, and quality windows.



CASE B: OFFLINE 6D POSE VALIDATION
Synthetic RGB-D + FoundationPose

Task
Spatial position and pose estimation

Model
FoundationPose

Synthetic RGB-D Data
Translation / rotation / depth / exposure / surface disturbance

Low-range deviation
Relative to synthetic ground truth.

No obvious outlier failure observed
Under controlled synthetic RGB-D offline conditions

Technical path screening
Not full real-cell validation


Boundary
Not real RGB-D sensor accuracy, real-calibration chain accuracy, robot TCP positioning or production deployment capability

Next validation: real workcell variables, sensor stability, calibration drift, occlusion, reflective materials, cycle time and robotic execution linkage

Stage-gate evidence

Common conclusion:
Both cases support moving validation earlier, but Pilot Cell entry still requires real workcell, cycle-time, quality and robotic execution validation.

06 3DoF / 6D Pose Case Evidence Summary



KPI, ROI and Deployment Readiness

From Model Metrics to Management Decisions

The business value of the Physical AI digital engineering loop should not be measured by model accuracy alone. For manufacturers, more important questions include whether vision risks are identified earlier, whether on-site trial uncertainty is reduced, whether edge cases are covered, whether data assets can be reused, and whether robotic tasks are ready to enter Pilot Cell validation.

Therefore, KPI and ROI assessment should support management decision-making rather than create return guarantees. Technical validation results should be translated into a combined assessment of engineering efficiency, operational stability, quality risk, asset reuse, and governance maturity.

KPI ASSESSMENT DIMENSIONS

Indicator Category	Management Focus
Engineering efficiency	Whether sample waiting, real-image collection, on-site tuning and changeover validation pressure can be reduced
Production operations	Whether mis-grasp, missed detection, false detection, rework, downtime and cycle-time impact can be improved
Quality and risk	Whether low-frequency abnormalities, edge cases, failure samples and recheck mechanisms are covered
Asset reuse	Whether digital twins, synthetic data, models and validation workflows can be reused across workcells
Governance maturity	Whether data, models, versions, approvals, rollback and change-triggered revalidation are traceable

CORE DEPLOYMENT READINESS QUESTIONS

Readiness Dimension	Management Question
Task definition	Are the target object, vision output, robotic action and quality boundary clear?
Data coverage	Do synthetic data and real-reference inputs cover key variables, edge cases and failure samples?
Robotic execution linkage	Can vision outputs be translated into robot targets, paths and station logic?
Quality and abnormal handling	Are misjudgment, missed detection, rework and recheck responsibilities integrated into the quality system?
Governance mechanism	Are data, models, versions, approvals, rollback and revalidation traceable?

NON-COMMITTAL ROI CALCULATION FRAMEWORK

$B_{total} = B_{tuning} + B_{commissioning} + B_{quality} + B_{reuse}$, including avoided on-site tuning cost, opportunity value from shorter commissioning, quality / downtime risk reduction and multi-cell reuse value.

Estimated ROI = $(B_{total} - C_{total}) / C_{total} \times 100\%$. This formula is intended to support internal business case development, budget discussion and stage-gate review. It should not be interpreted as a return commitment.


$C_{total} = C_{software} + C_{hardware} + C_{integration} + C_{data} + C_{model_ops} + C_{governance}$, including software, hardware, system integration, data operations, model maintenance, governance and training costs.

If critical readiness items are not yet satisfied, the project should enter Refine or additional validation rather than move directly to the next stage. Actual ROI depends on workcell complexity, sample availability, real-world validation plans, replication scale, data governance maturity and deployment context.

KPI, ROI and Deployment Readiness Framework

From technical evidence to management stage-gate decisions

Technical evidence input KPI / ROI / readiness judgment should be derived from the evidence package, not from model metrics alone.



Key point: KPI and ROI should translate technical evidence into management decisions. Model metrics are necessary but not sufficient; deployment readiness requires engineering, production, quality, reuse and governance evidence, evaluated through enterprise-specific assumptions and stage-gate review.

Implementation Roadmap and Ecosystem Collaboration

From Digital POC to Scale-Up

Manufacturers advancing Physical AI and robotic vision digital engineering should not start by targeting large-scale deployment immediately. A staged path is more appropriate: first use Digital POC to validate the task, data, and model path; then use Pilot Cell to validate robotic execution, cycle time, quality, and site workflow under real or near-real conditions; and finally consider multi-cell scale-up, extending reusable assets, governance mechanisms, and validation workflows across more products, workcells, or factories.

Each stage should produce reviewable evidence, not only a model demonstration result. Digital POC should at least generate a data coverage summary, model validation results, failure samples, and an initial deployment readiness judgment. Pilot Cell should further generate real-world feedback, cycle-time validation, quality risk assessment, and robotic execution linkage validation. Scale-up should focus on data asset reuse, governance mechanisms, replication cost, and continuous feedback loops.

THREE-STAGE ROADMAP

Stage	Stage Objective	Management Decision Question
Digital POC	Validate task definition, digital assets, synthetic data coverage, AI model behavior and initial deployment readiness	Is the project ready to enter Pilot Cell?
Pilot Cell	Validate the fit among robot, vision, process, quality and site workflow in a real or near-real workcell	Go / Refine / No-Go
Multi-cell Scale-up	Extend validated data assets, model validation systems, governance mechanisms and engineering workflows to more products, cells or factories	Is the capability ready for broader replication or continuous optimization?

ECOSYSTEM COLLABORATION

The deployment of Physical AI in high precision manufacturing cannot be achieved by a single model, a single software tool, a single robot, or a single company. The joint publishers and main contributors contribute complementary methodology perspectives to support the engineering loop. This collaboration should not be interpreted as a single commercial solution or a fixed implementation path.

- **ABB Robotics:** industrial robotics, automation engineering, robotic digital engineering verification and physical deployment practice;

- **NVIDIA:** simulation, AI acceleration, the Physical AI computing ecosystem and foundation-model capabilities represented by FoundationPose;
- **AsialInfo:** enterprise AI platforms, data governance, model lifecycle management and industry application integration perspectives;
- **Deloitte:** industry transformation, operating model design, ROI assessment and scale-up methodology;
- **SKAI Intelligence:** digital twin assets, synthetic data, auto annotation, robotic vision data infrastructure and Real2Sim2Real technical experience.

FOUR PRINCIPLES FOR IMPLEMENTATION


- Digital-first**
Not replacing real-world validation
- Data-driven**
Grounded in engineering constraints
- Feedback-loop driven**
Gated by deployment readiness review
- Ecosystem collaboration**
Not a single commercial solution



Implementation Roadmap and Ecosystem Collaboration

Build evidence stage by stage from Digital POC to Pilot Cell and Multi-cell Scale-up


Roadmap principle: Start with task, data and model evidence: validate real engineering constraints in Pilot Cell; scale only when assets, workflows and governance are reusable.



Key point: Implementation should build evidence gradually, not expand deployment immediately. Digital POC validates the task, data and model path; Pilot Cell validates real engineering constraints; multi-cell scale-up depends on reusable assets, governed workflows and clear collaboration roles.


04 STAGE-GATE EVIDENCE

Minimum evidence formed before moving to the next stage




05 ECOSYSTEM COLLABORATION AS SUPPORTING METHODOLOGY

Complementary capabilities support the loop; this is not a fixed commercial implementation path



**SEE EARLIER.
TRAIN SMARTER.
VALIDATE DIGITALLY.
DEPLOY WITH
FEEDBACK-BACKED
READINESS.**



Appendix A

Glossary

This glossary aligns the core concepts used in this white paper. Some terms are retained in English because they are common in robotics, AI, digital engineering and manufacturing contexts.

Term	Definition	Term	Definition
Physical AI	AI capability for physical-world tasks, emphasizing the perception, judgment, execution and feedback loop.	Robotic Vision	Vision capability that helps robots understand object position, pose, state and environmental conditions.
Digital Twin	Engineering representation of products, vision, stations, robots and workflows in a digital environment.	Product Digital Twin	Attributes such as geometry, material, surface, color, tolerance and product version.
Vision Digital Twin	Camera, lens, lighting, exposure, background, occlusion and imaging-response representation.	Station / Robot Digital Twin	Robot, tool, fixture, layout, path, reachability and safety-boundary representation.
Synthetic Data	Data generated in a digital environment, with task relevance and traceability.	Task-Relevant Synthetic Data	Data built around a task, visual variables, edge cases and failure modes.
Real2Sim	Converting real products, stations, cameras, lighting and reference images into digital assets.	Sim2Real	Bringing digital validation results, model outputs and robot targets into physical validation.
Real2Sim2Real	Continuous calibration between the real world, digital environment and real feedback.	AI Validation	Validating whether model outputs meet task-relevant metrics.
Robotic Digital Engineering Verification	Validating whether vision outputs can enter targets, paths and station logic.	Pilot Cell	A trial workcell for validating real or near-real conditions.
Digital POC	Pre-deployment evaluation through digital twins, synthetic data, AI and robotic verification.	Multi-cell Scale-up	Extending validated assets, models and governance to more cells.
Deployment Readiness	Evidence basis for entering the next stage.	3DoF	Three-degree localization, mainly x–y planar position and Rz angle in this paper.
6D Pose	Six-dimensional pose estimation: three-axis translation and three-axis rotation.	RGB-D	RGB image plus depth data, commonly used for spatial pose estimation.
FoundationPose	Model capability used for 6D Pose technical path screening.	Foundation Model	A model trained at scale and adaptable to multiple tasks
Edge Case	A low-frequency or near-failure sample that challenges robustness.	Failure Sample	A sample with misjudgment, missed detection, abnormal output or unsupported robot execution.
Go / Refine / No-Go	Stage decision: proceed, strengthen validation, or pause / redirect.	RACI	Responsibility model: Responsible, Accountable, Consulted and Informed.
TCO	Total cost of ownership, including software, hardware, integration, operations, maintenance and governance.		

Appendix B

Case A: 3DoF Localization POC Evidence Package

Case A provides stage-gate evidence for 3DoF vision localization. It should be read as image-level and early engineering evidence, not as direct proof of process tolerance or full robotic execution capability.

Evidence item	Current record
Task	x–y planar position and Rz rotation-angle estimation
Data scale	1,000 synthetic images and 120 representative training samples
Synthetic variables	Background, exposure, field of view, surface disturbance and rotation angle
Real-reference validation	Two real-reference image validation sets
Absolute accuracy criteria	x–y ≤ 2 pixels; Rz ≤ 0.2 degrees
Repeatability criteria	x–y ≤ 1 pixel; Rz ≤ 0.1 degrees
x–y result	Both real-reference validation sets met the pixel-level boundary
Rz result	Strict angular-threshold coverage: approximately 88% / 76.7%; near-threshold observation: approximately 92% / 84.7%
Abnormal-sample recheck	10 abnormal images generated localization outputs suitable for qualitative human review
Verified scope	Synthetic-data training path, x–y stability, Rz coverage, repeatability and abnormal-sample recheck
Evidence to add before Pilot Cell	Pixel-to-mm mapping, FOV, process tolerance, hand-eye calibration, TCP positioning, robot targets and quality window
Stage judgment	Useful as pre-Pilot Cell stage-gate evidence for vision localization; error-chain mapping remains required

How to read the result

- x–y pixel performance should be mapped to engineering units before process judgment.
- Rz coverage should be interpreted under a stringent angular threshold, not as a generic pass rate.
- Pilot Cell readiness requires robotic execution and quality-window mapping.

Appendix C

Case B: Offline 6D Pose Validation Evidence Package

Case B provides technical path-screening evidence for 6D Pose. It should be interpreted as controlled synthetic RGB-D offline evaluation relative to

synthetic ground truth, not as real sensor, real calibration-chain, robot TCP or production deployment performance.

Evidence item	Current record
Task	3D spatial position and pose estimation
Input	Synthetic RGB-D data
Model path	FoundationPose
Randomized variables	Translation, rotation, depth, exposure, scratches and noise
Ground truth	Synthetic ground truth
Result meaning	Model-output deviation relative to synthetic ground truth under controlled synthetic RGB-D offline evaluation
Result summary	Translation and rotation deviations were in a low range; no obvious outlier failure observed under current conditions
Verified scope	6D Pose technical path-screening potential
Should not be extrapolated as	Real RGB-D sensor accuracy, real calibration-chain accuracy, robot TCP positioning accuracy or production deployment capability
Evidence to add	Real sensor stability, calibration drift, occlusion, reflective materials, cycle time, robotic execution linkage and quality recheck
Stage judgment	Can support next-stage real-cell or Pilot Cell validation design; cannot serve as a deployment conclusion

CASE EVIDENCE PACKAGE OUTPUT FORMAT	
Output	Recommended use
Case summary	Task, test conditions, aggregate results, boundary and next validation step
Boundary table	Verified scope, unverified conditions and Pilot Cell validation items
Stage decision	Go / Refine / No-Go and supporting evidence

- Key takeaways**
- Case B is useful for technical path screening.
 - Synthetic-ground-truth results must not be over-extended to real work cell performance.
 - Pilot Cell validation should focus on sensor, calibration, robotic execution, cycle time and quality mechanisms.

Appendix D

Reference Figures and Usage Notes

Figure Number	Title	Primary Use	Page
Figure 1	Validation Moves Earlier	Explain why robotic vision risk should enter digital engineering review earlier.	11
Figure 2	Physical AI Digital Engineering Reference Architecture and Loop	Show the relationship between reference architecture, digital engineering loop, AI validation, robotic verification and feedback.	13
Figure 3	From Vision Outputs to Robotic Tasks	Show how model outputs enter robot targets, paths, reachability and station logic.	15
Figure 4	Real2Sim2Real Architecture for Robotic AI Deployment	Show the loop from real inputs, digital twins, synthetic data and AI validation to robotic verification and feedback.	17
Figure 5	Application Scenario Matrix	Help enterprises select Digital POC entry points and caution conditions.	19
Figure 6	3DoF / 6D Pose Case Evidence Summary	Summarize case tasks, data, results, boundaries and next validation steps.	21
Figure 7	KPI, ROI and Deployment Readiness Framework	Translate technical evidence into KPI, ROI and deployment-readiness decision logic.	24
Figure 8	Implementation Roadmap and Ecosystem Collaboration	Show digital POC, pilot Cell, multi-cell scale-up and ecosystem collaboration roles.	28

Appendix E

Key Takeaways

Executive Summary: Moving Vision Validation Toward Deployment Readiness	The core value of Physical AI is to move part of robotic vision risk identification into the digital engineering stage.	Synthetic data and digital twins should support deployment readiness, not only model training.	Manufacturers should use stage-gate evidence to support Go / Refine / No-Go decisions.	
Why Physical AI Now	High precision manufacturing is pushing robots from fixed execution toward perception-driven automation.	Robotic vision is a critical entry point for Physical AI.	Traditional workflows remain valuable, but validation often happens too late in high-variation environments.	The core value of digital-first validation is earlier risk identification, not replacing real-world validation.
Digital Engineering Reference Architecture and Loop	Physical AI deployment requires a reference architecture, not a stack of individual tools.	The core output of the digital engineering loop is deployment readiness evidence, not a single model accuracy metric.	Real-world feedback must flow back into digital twins, data assets and the next validation cycle.	
Robotic Digital Engineering Verification	Robotic vision must support robotic tasks, not only model-level metrics.	After AI validation, vision outputs must be verified against robot targets, paths and station logic.	RobotStudio® HyperReality should appear as an engineering implementation example, not as a standalone product narrative.	
Synthetic Data and the Real-World Feedback Loop	Synthetic data is not image stacking; it is task-relevant data asset creation.	The value of synthetic data is to express low-frequency variables, boundary conditions and failure modes systematically, not to replace real data.	Real2Sim2Real can be understood as a continuous calibration loop between the real world, the digital environment and real-world feedback.	An effective synthetic data system should support deployment readiness, not only model training.
Application Scenario Matrix	Physical AI should not be treated as a universal solution for all manufacturing scenarios.	The first Digital POC should prioritize clear task boundaries, clear value and controllable risk.	6D Pose should be positioned as a key vision capability within spatial grasping and assembly guidance, not as a standalone manufacturing scenario.	Scenario selection directly affects data preparation, validation path, ROI assessment and deployment readiness.
Case Evidence Summary: 3DoF and 6D Pose	The 3DoF Localization POC shows that synthetic data can support early validation and abnormal sample recheck for planar localization tasks.	The Offline 6D Pose Validation shows that FoundationPose and synthetic RGB-D data can support early screening of spatial pose-estimation paths.	Both cases support moving the validation window earlier, but neither should be interpreted as proof of production deployment.	Before entering Pilot Cell, real work cell, cycle-time, quality and robotic execution linkage still need validation.
KPI, ROI and Deployment Readiness	KPIs should expand from model metrics to deployment readiness and operational indicators.	The ROI formula is intended for internal estimation and stage-gate review, not as a return commitment.	Deployment readiness requires joint assessment of tasks, data, robotic execution, quality and governance.	If critical items are insufficient, the project should enter Refine rather than move directly to the next stage.
Implementation Roadmap and Ecosystem Collaboration	POC success does not equal production deployment success.	Manufacturers should build deployment readiness through digital POC, pilot Cell and multi-cell scale-up.	Physical AI deployment requires robotics, AI, data, engineering, quality and governance systems to work together.	The goal of ecosystem collaboration is to build verifiable, reusable, governable and scalable robotic AI engineering capability.

Appendix F

Contributors

ABB Robotics	NVIDIA Robotics	AsiaInfo	Deloitte	SKAI Intelligence
Steven Chen Senior Director, Electronics	Shri Sundaram Director of Product Management, Isaac Robotics Platform	Xiangde Meng Vice President, Executive GM, R&I Center (2030 Lab)	Scott Xiaozhen Sun Deloitte Asia Pacific Physical AI Centre of Excellence Leader, Deloitte China	Jaechool Lee CEO
Leon Yan Technical Architect Manager, Electronics	Brian Klobucher Senior Manager, Omniverse Ecosystem	Changshui Liu Head, Intelligent Manufacturing STF	Geliang Gong Deloitte China Supply Chain and Network Operations Offering Leader, Deloitte China	Jaehyuck Shin Partner
Louis Liu Strategic Account Manager, Electronics	Staffan Hellsvik Senior Solution Architect, Manufacturing, Sweden/Nordics	Guo Chen Deputy GM, R&I Center (2030 Lab)	Mike Dajiang Guo Deloitte Asia Pacific Physical AI Centre of Excellence Partner, Deloitte China	Josh Liu CTO
Celia Xu Technical Program Manager, Electronic	Andrew Dahdouh Senior Developer Relations Manager, Robotics	Yang Cao GM, Physical AI Innovation Center	Alain Yilei Sun Deloitte China Supply Chain and Network Operations Director, Deloitte China	Alex Kong Director of Physical AI Dataflow & Synthetic Reality Pipeline
Ina Yao Eco-System Partner Manager, Electronics	Jonathan Tremblay Senior Research Scientist, Computer Vision/Robotics	Ying Ding Chief Planning Expert, Physical AI Innovation Center	Joe Zhuo Deng Deloitte China Supply Chain and Network Operations Manager, Deloitte China	Shawn Ding Director of Physical AI Engineering & Systems Integration
Dalla Yang Product Marketing Specialist, Electronics	Hugo Hadfield Senior Robotics Research Software Engineer	Dr. Yucheng Wang Product Director, Physical AI Innovation Center		Benjamin Zhou Physics-Consistent Neural Materials Architect
Rachel Wang Communication Manager, Electronics	Benedetta Delfino Solutions Architect, Robotics	Qingping Peng Technical Director, Intelligent Manufacturing STF		Leon Wang Simulation-Aware Material Systems Architect
Sherry Liu Strategic Program Manager, Electronics	Heather McDiarmid Senior Manager, Physical AI Marketing, Robotics Partners	Dr. Feiyu Zhu Technical Director, Physical AI Innovation Center		Leo Yang Physical Lighting & Sensor Simulation Architect
Darren Hung Robotics Product Senior Manager, Electronics		Boyu Qi Algorithm Engineer, Intelligent Manufacturing STF		Zhe Tan Physical AI Solution Delivery Architect
Zoe Huang Robotics Product Specialist, Electronics		Jian Chen Algorithm Engineer, Intelligent Manufacturing STF		Nojun Kwak Independent Non-Executive Director
Michelle Wang Product Portfolio Senior Manager, Electronics				Dongkyun Han Vice President
Josh Wu Product Portfolio Manager, Electronics				Ying Xiao Head of Investment
				Celine Lo Head of Marketing





ABB Robotics

Hagenholzstrasse 60
8050 Zürich
Switzerland

Additional information

We reserve the right to make technical changes or modify the contents of this document without prior notice. With regard to purchase orders, the agreed particulars shall prevail. ABB AG does not accept any responsibility whatsoever for potential errors or possible lack of information in this document.

We reserve all rights in this document and in the subject matter and illustrations contained therein. Any reproduction, disclosure to third parties or utilization of its contents – in whole or in parts – is forbidden without prior written consent of ABB AG.

© Copyright 2026 ABB. All rights reserved.
Specifications subject to change without notice.



**ABB
Robotics**



**Electronics
Segment**