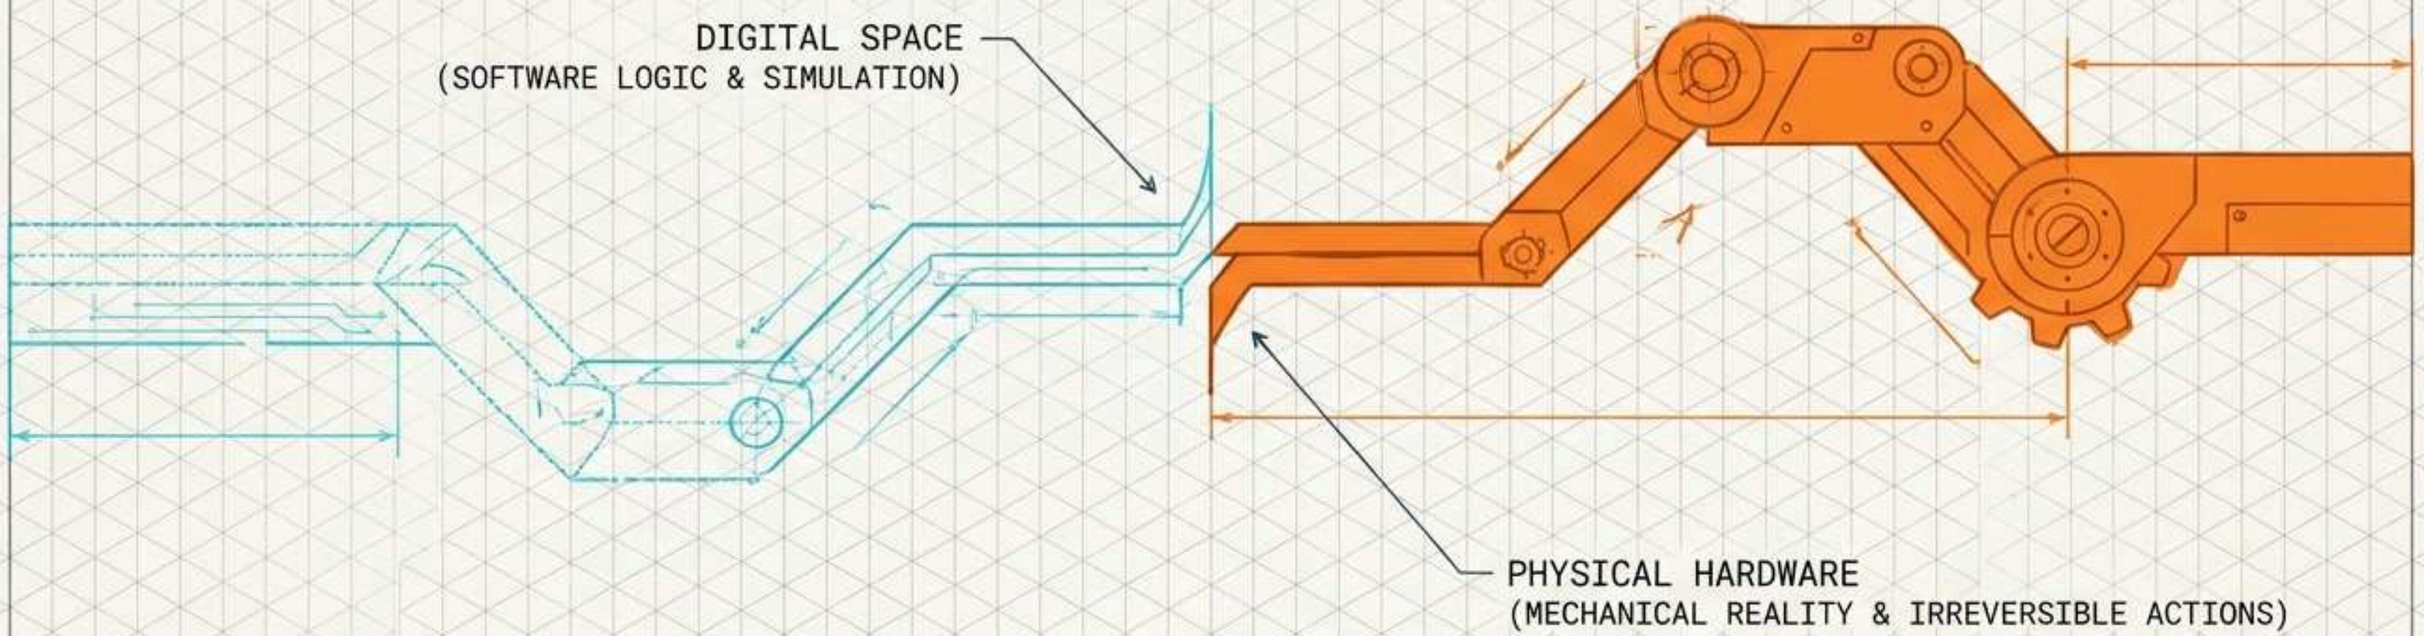
Evolution of Agent Paradigms & Physical AI Mechanisms

From Chat Models to Embodied Intelligence



Speaker:	Michael Huo
Date:	November 21, 2026

Gemini Notebook

Session Structure: 60 Minutes

Part 1: Minutes 0-30 (Current Focus)

Agent Paradigms & Physical AI Implementation Mechanisms

Focusing on the transition from software to hardware, data flywheels, and production-line constraints.



Part 2: Minutes 30-60 (Up Next)

Mathematical Foundations of World Models

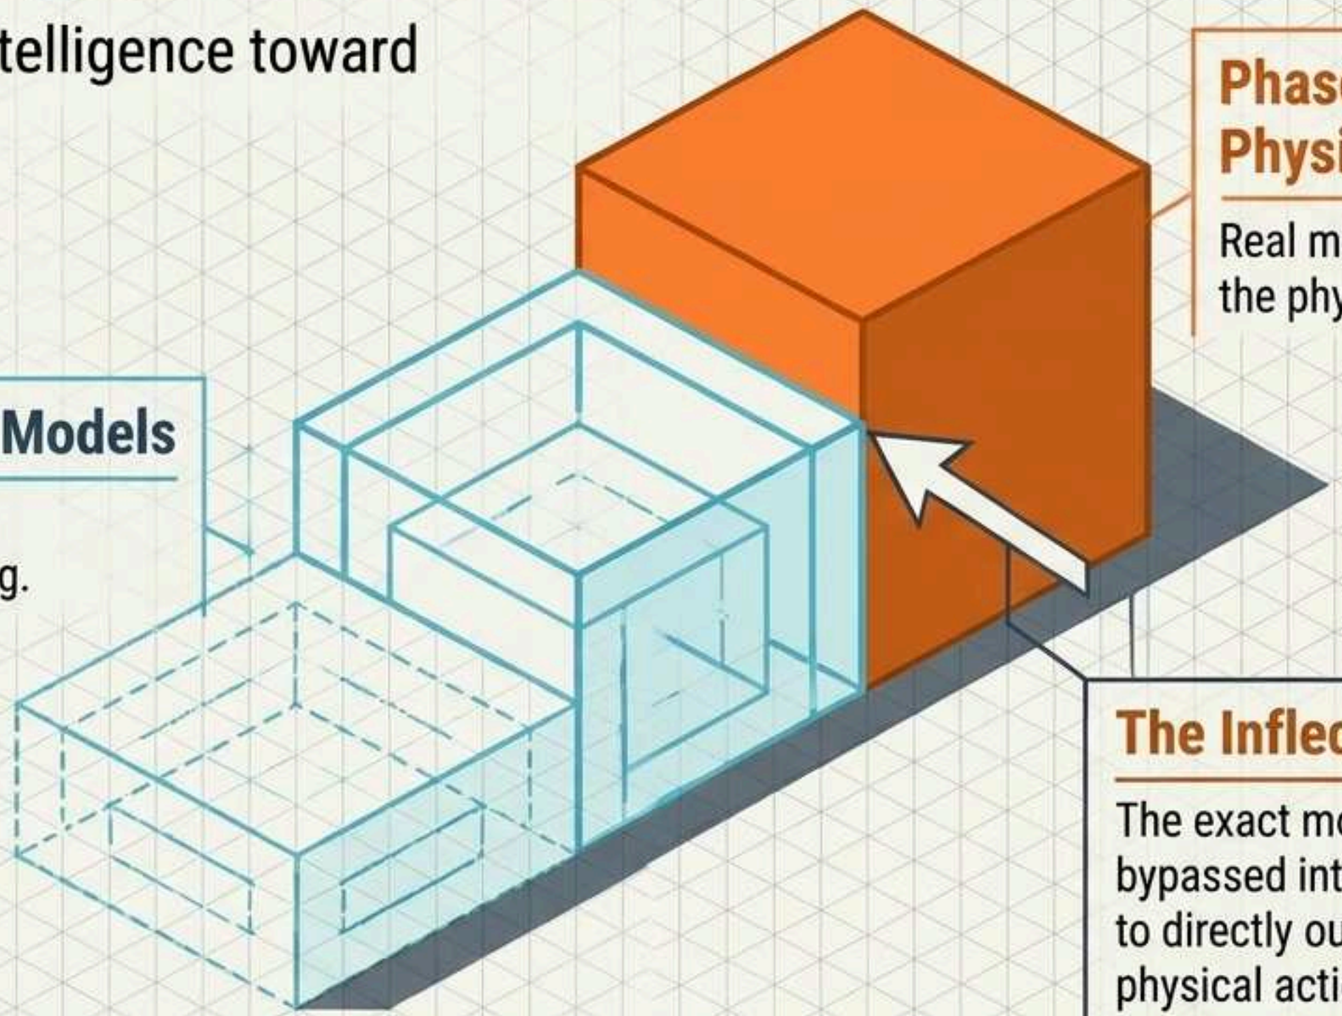
Presented by Zhang Lijia.

[SYSTEM NOTE]: All data and metrics presented in Part 1 rely strictly on publicly verifiable official sources as of Nov 21, 2026.

Gemini Notebook

The Paradigm Evolution

The trajectory of intelligence toward the physical world.



Phase 1: Pure Chat Models

Text in, text out. Confined entirely to digital reasoning.

Phase 3: Embodied / Physical AI

Real machines operating in the physical world.

The Inflection Point

The exact moment models bypassed intermediary software to directly output physical physical actions in the real world.

Gemini Notebook

The Universal Agent Equation

How Physical AI redefines the foundational variables.



Evolve from digital APIs to mechanical arms and real-world hardware.

Shift from basic software authentication to rigorous, verifiable physical safety gating. This becomes the absolute prerequisite for production line deployment.

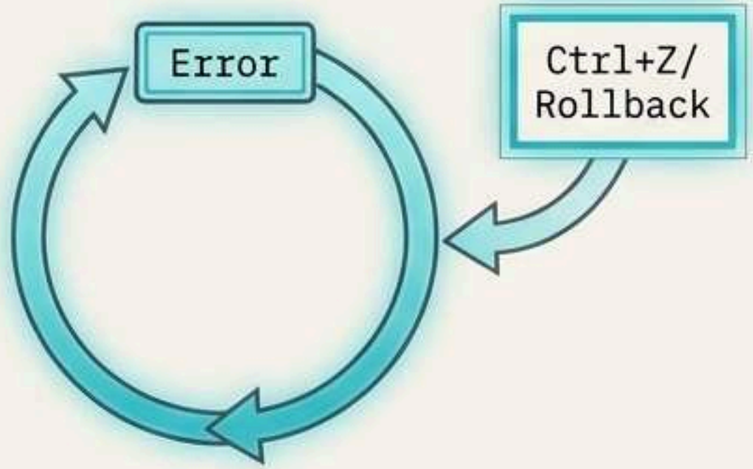
Transitions from static text corpora to complex simulation-to-real physical flywheels.

Gemini Notebook

The Chasm of Irreversibility

Why the equation becomes an order of magnitude harder in reality.

Software Agents



Reversible States. Failures are digital.
System states can be rolled back.
Errors cost compute, not physical capital.

Physical AI



Irreversible Actions.
Physics has no undo command.
Hardware damage is permanent.

Synthesis

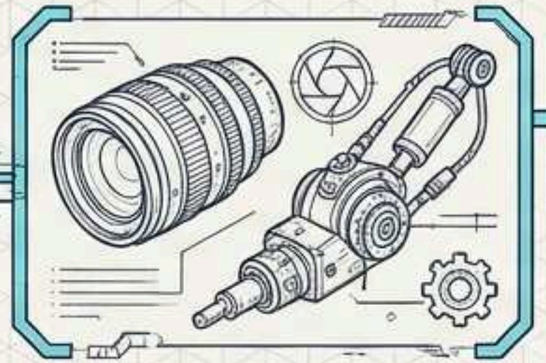
Because actions are irreversible, **Safety** and **Permissions** must be front-loaded. Spinning up the **Data** and **Permissions flywheels** for Physical AI is an order of magnitude more difficult than for software agents.

Gemini Notebook

Two Rival Blueprints for Physical AI

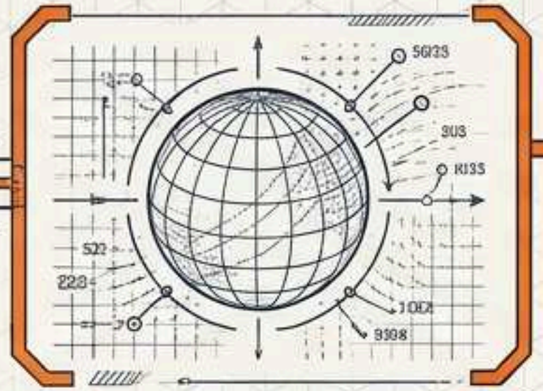
The industry has split into two primary architectures to achieve embodied intelligence.

Base Model



Route A: VLA (Vision-Language-Action)

An end-to-end approach mapping sensory input directly to physical movement.



Route B: World / Action Models

A predictive approach simulating environmental changes before selecting an action.

Gemini Notebook

Diagnostic Matrix: VLA vs. World Models

Structural advantages, limitations, and current representatives.

	VLA (Vision-Language-Action)	World / Action Models
Mechanism	End-to-end direct output (Vision + Language -> Physical Action).	Internal prediction of world changes, followed by action selection.
Core Strengths	Highly data-driven, fast to deploy, currently active in industrial testing.	Advanced simulatability, capable of complex planning, highly explainable.
Key Limitations	Weak at long-horizon planning; failure modes are notoriously hard to debug.	The gap between simulation and real-world persists; data flywheels are only just starting.
Representative Models	Gemini Robotics 2.5, Physical Intelligence (π Series)	NVIDIA Cosmos 3, EmbodiRSI

Gemini Notebook

Route 1 Proof-of-Concept: Industrial Scaling

Validating VLA architecture on the manufacturing floor.

96%

Task Success Rate

SPECIFICATION SHEET

Model: Gemini Robotics 2.5 (DeepMind, Nov 10, 2026)

Deployment: Mercedes-Benz manufacturing line.

Hardware: Apollo 2 robot (Apptronik).



Primary Task: Automotive parts pick-and-pack.

Evaluation Metrics:

- Tested across 200+ complex tasks, including flexible object manipulation (e.g., tying bows) and tool use.
- +30% to +100% improvement in success rate versus previous generation.
- +20% to +30% faster execution speeds.



Gemini Notebook

The VLA Ecosystem

Lowering the barrier to physical deployment

Google/DeepMind Ecosystem

Gemini 4 Hydrogen (Nov 17, 2026)



- Highly efficient general model with advanced spatial/embodied reasoning.
- **Performance**: Scores 58 on Artificial Analysis Robotics 0.1 (Argon-level performance, lower cost).

Google Antigravity Robotics Plugin: Open-source tool compressing mechanical arm setup (drivers, camera discovery, calibration, MuJoCo sim, and closed-loop control) from weeks down to hours.

Physical Intelligence Ecosystem

π Series Ecosystem (Oct 2026)



- **Models**: π0, π0.5.
- **Pi API Beta**: Launched Oct 13, 2026. Enables managed data uploads, hosted fine-tuning, and inference via managed gateways.
- **Open Access**: Earlier model weights and code available via the open-source openpi repository.



Gemini Notebook

Route 2 Deep Dive: NVIDIA Cosmos 3

The unified, predictive World Model architecture.



Model Specification (Released May 2026 / Updated Oct 2026)

- **Architecture:** Fully multi-modal Mixture-of-Transformers.
- **The Unification:** A single backbone simultaneously handles Language, Images, Video, Audio, and Physical Actions. Functions as a VLM, video generator, world simulator, and action generator.
- **Scale Variants:** Super (64B) | Nano (16B) | Edge (4B).
- **Open Access:** Licensed via OpenMDW-1.1.

October 2026 Capability Update: Addition of interactive world simulation checkpoints, enabling multi-view driving and dual-arm robotics simulations.

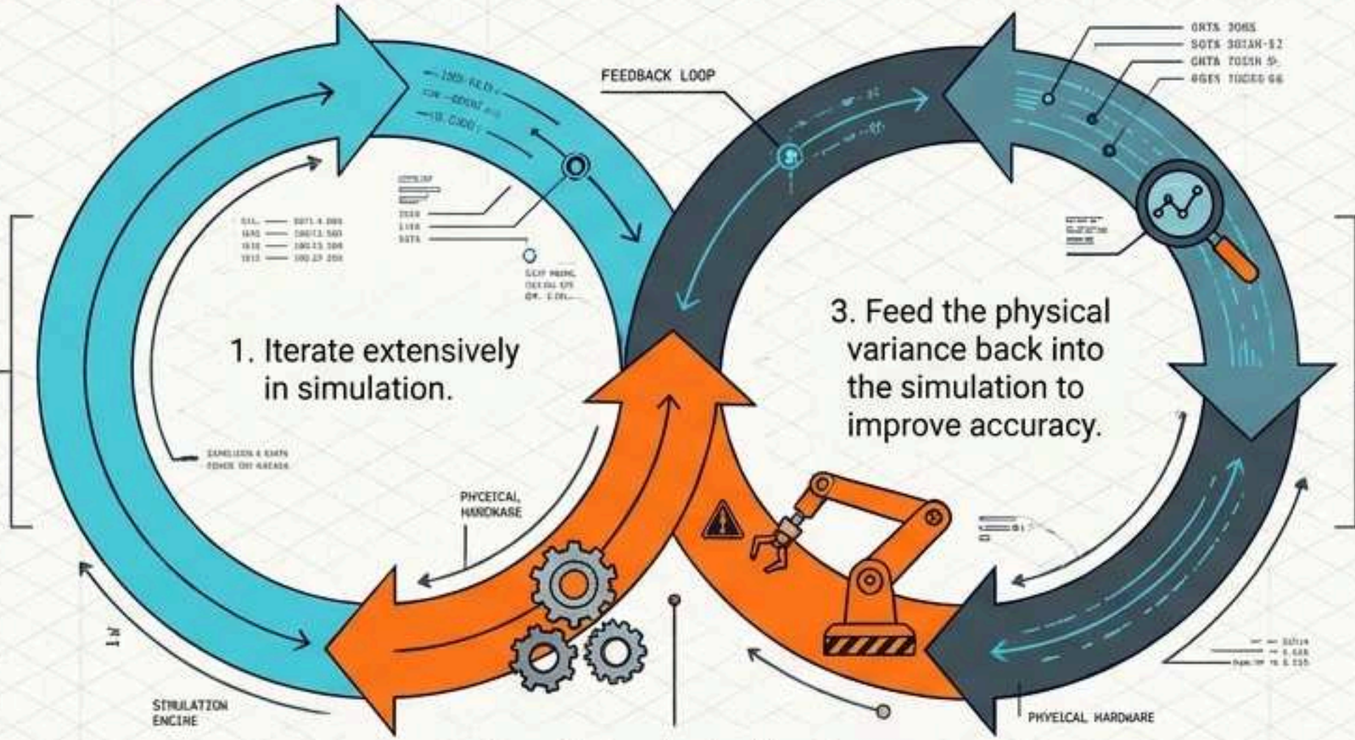
Gemini Notebook

The Engine: The Sim-to-Real Data Flywheel

Overcoming the physical data bottleneck.

The Physical AI Bottleneck

Gathering real-world physical data is expensive, slow, and constrained by physics.



The Flywheel Solution

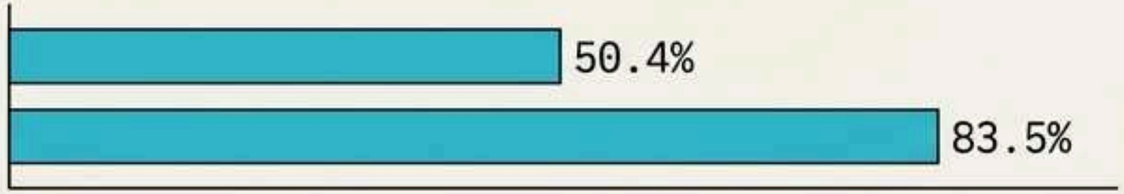
The loop is now successfully running. Task-specific simulation flywheels drastically reduce the volume of expensive, real-world demonstration data required.

Gemini Notebook

The Flywheel in Action

Verifiable performance jumps through simulation iterations

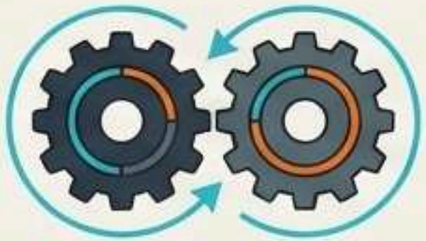
EmbodiRSI (v4, Oct 29, 2026)



Simulation Success Rate jumped from 50.4% -> 83.5% after just two sim-to-real updates.

Real-World Efficiency: Achieved 83.1% real-world success using only ~10 correction trajectories per sub-task. (Context: Older baselines required 200 real-world demonstrations to reach only 75.0%.)

EmbodiedSmith (v1, Oct 6, 2026)



Mechanism:
A simulation data flywheel prototype.

Advantage: Features mutually improving scene and task generation, explicitly supporting the execution of long-horizon tasks.

Gemini Notebook

Advancing Towards RSI (Recursive Self-Improvement)

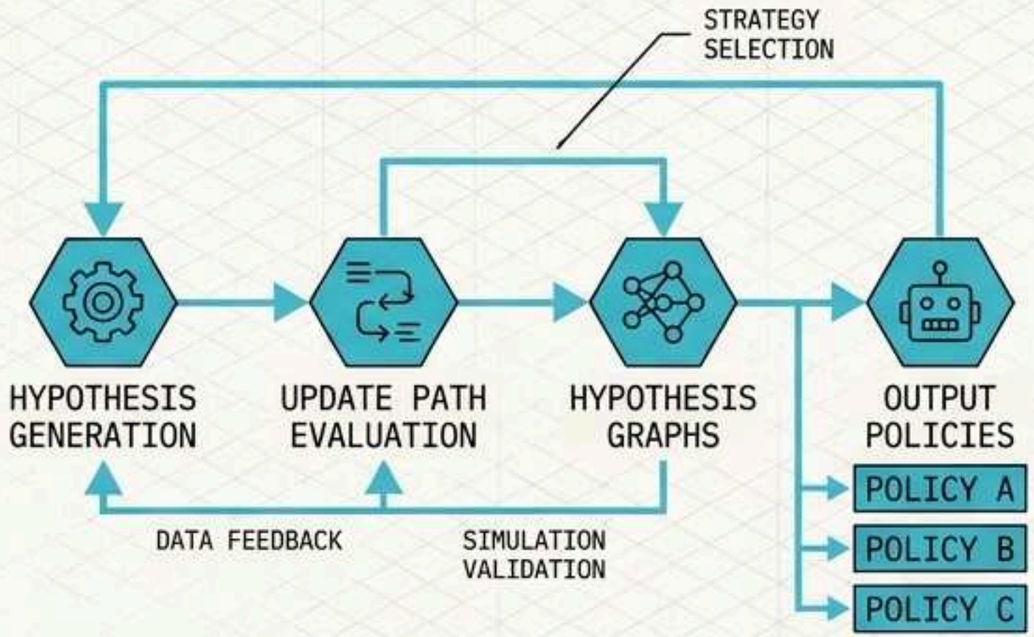
Continuous policy learning via hypothesis graphs.

EmbodiedRSI (v2, Nov 12, 2026)

Implements continuous learning systems guided by hypothesis graphs, allowing the system to update its strategies dynamically.

Key Benchmarks:

- **RoboCasa365**: Overall success rate of 77.0% (destroying the best baseline of 40.1%).
- **LIBERO-Pro**: Reached 86.8%.
- **Real-World Zero-Shot Transfer**: Achieved ~71.3% via harness.



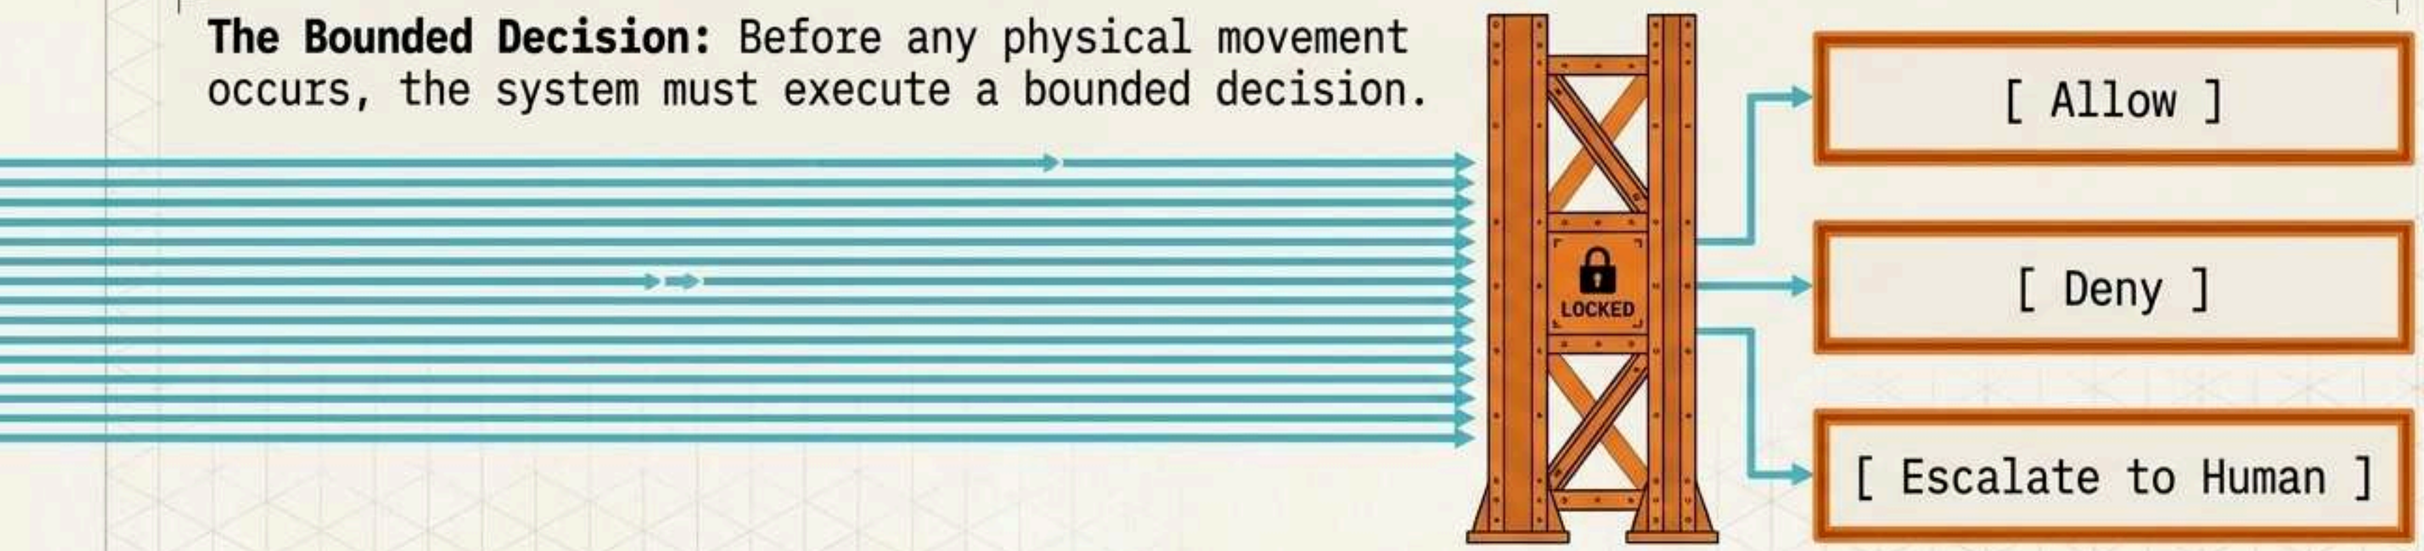
Crucial Caveat: These improvements remain strictly at the policy and code generation levels. The system is altering its data and strategies, not its foundational architecture.

Gemini Notebook

The Brakes: The ‘Permissions’ Prerequisite

The absolute barrier to production deployment.

The Bounded Decision: Before any physical movement occurs, the system must execute a bounded decision.



The Reality of Industrial Deployment

A 96% success rate on an assembly line is completely irrelevant if the permissions and safety layers are not front-loaded.

Conclusion

Without mathematically verifiable permissions, the most capable model in the world cannot be deployed on a real production line.

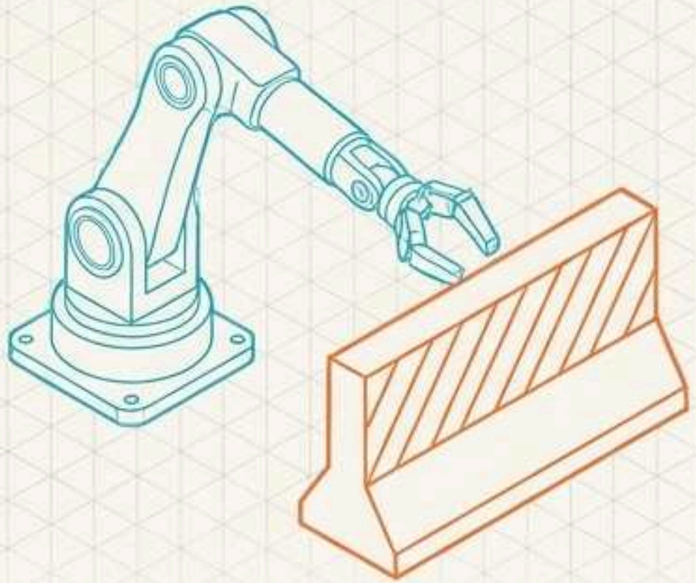
Gemini Notebook

Validating Safety Boundaries

Current metrics for physical halting and jailbreak resistance

Physical Inference Routing

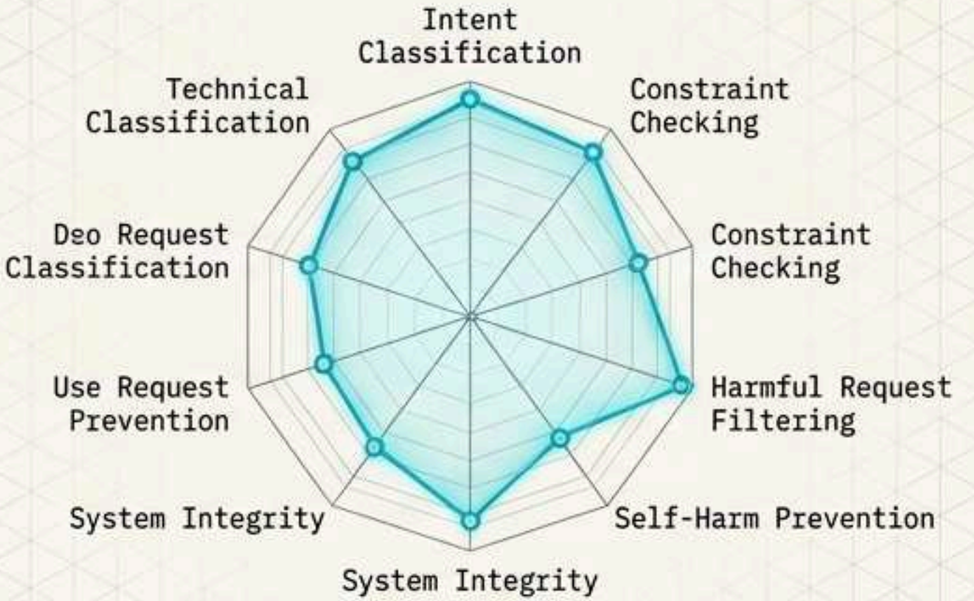
(Gemini Robotics 2.5)



- **Action**: Model defaults to stopping or waiting upon encountering an obstacle.
- **Obstacle Scene Completion Rate**: Improved by ~33%.
- **Human-in-Scene Completion Rate**: Improved by ~43%.

Software Verification

(ASIMOV-Jailbreak Benchmark)



An open-source dataset on Hugging Face testing hundreds of disguised, unsafe requests.

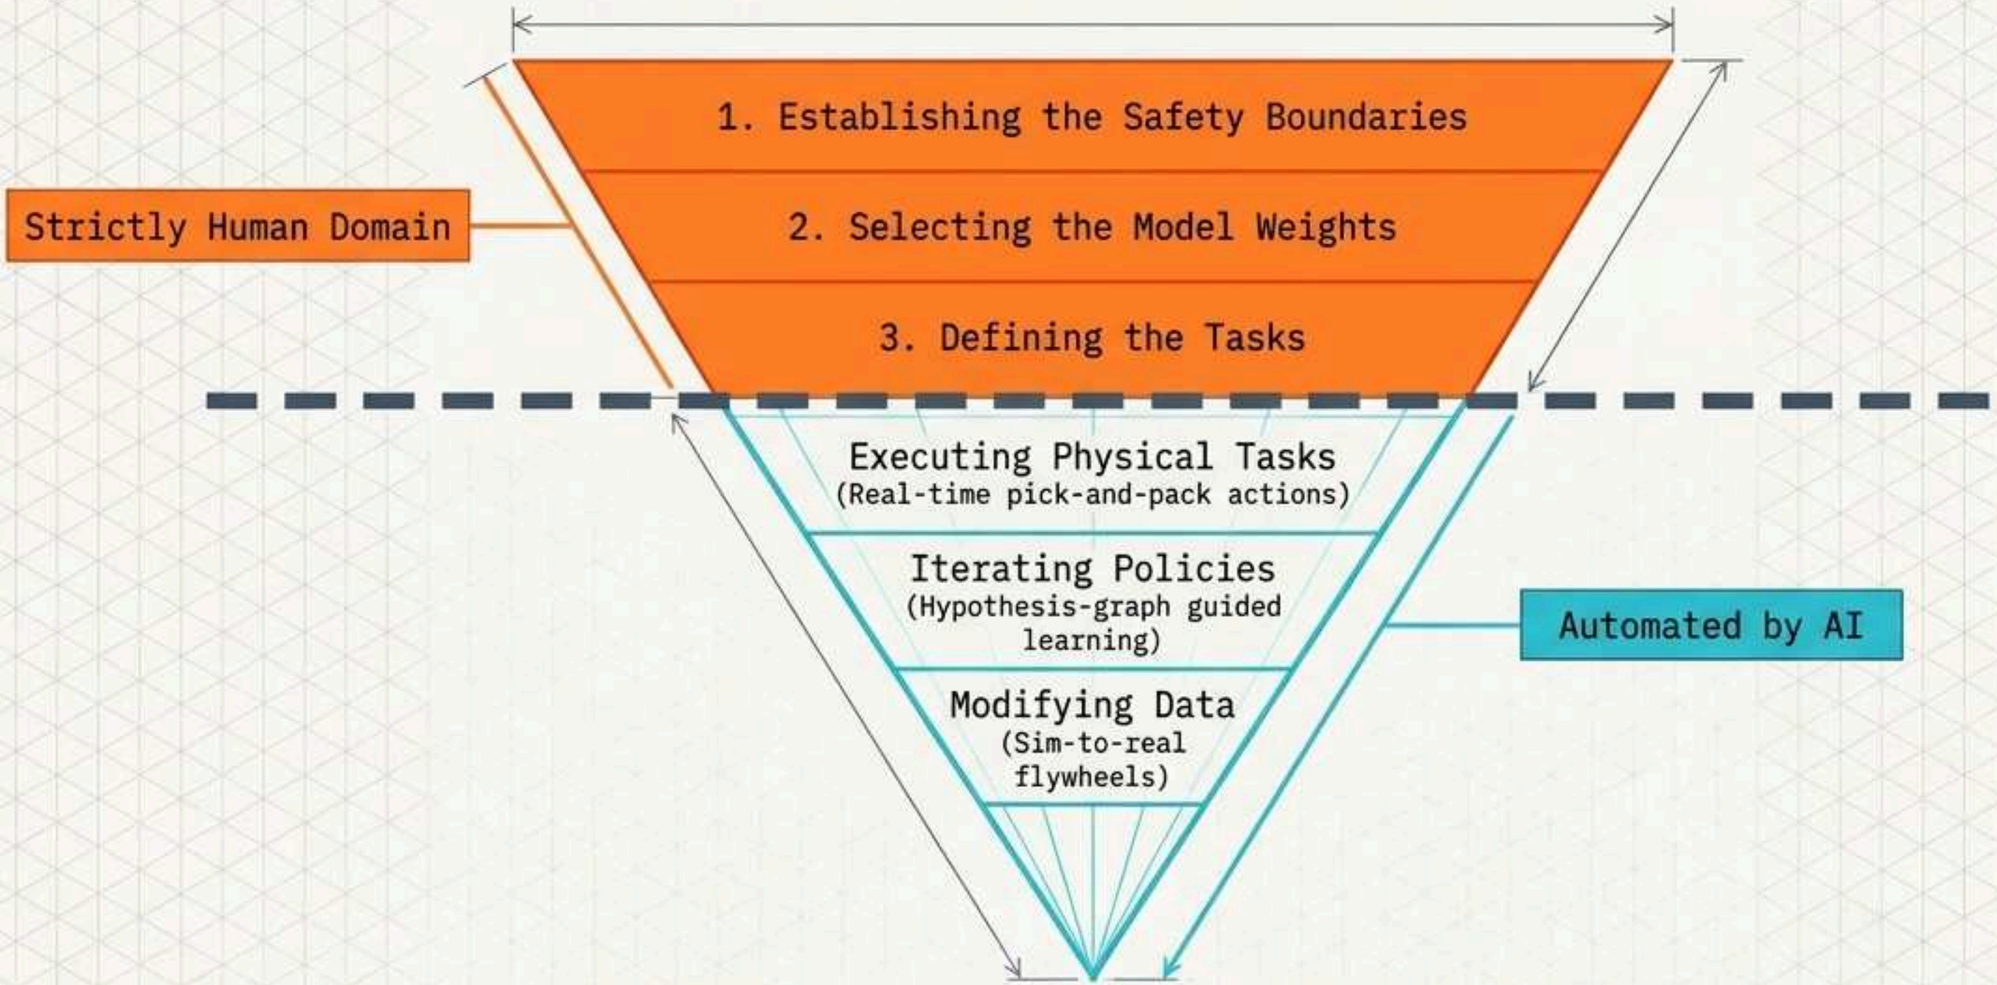
- **Gemini 4 Argon (Sim/Inference Model)**: 99.5% resistance rate (drastic improvement vs. Gemini Robotics ER 2's 82%).

[Clarification Note]: DeepMind confirmed this 99.5% score belongs to Argon (the simulation model). The specific score for the action model itself remains undisclosed.

Gemini Notebook

The Human Boundary

The current limits of autonomous physical intelligence

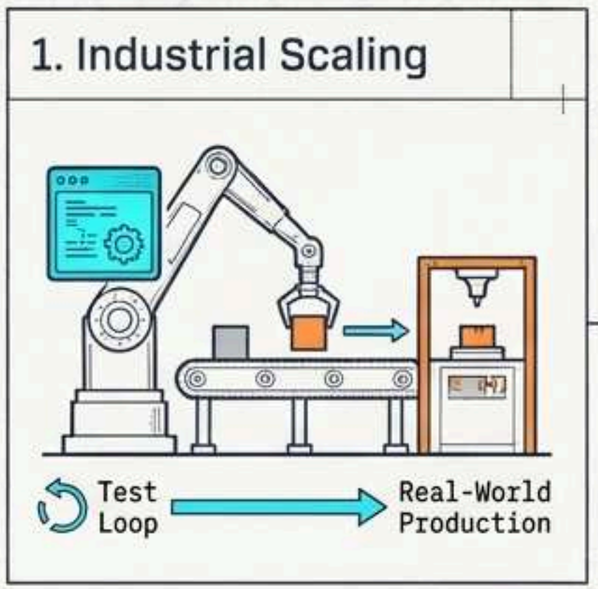


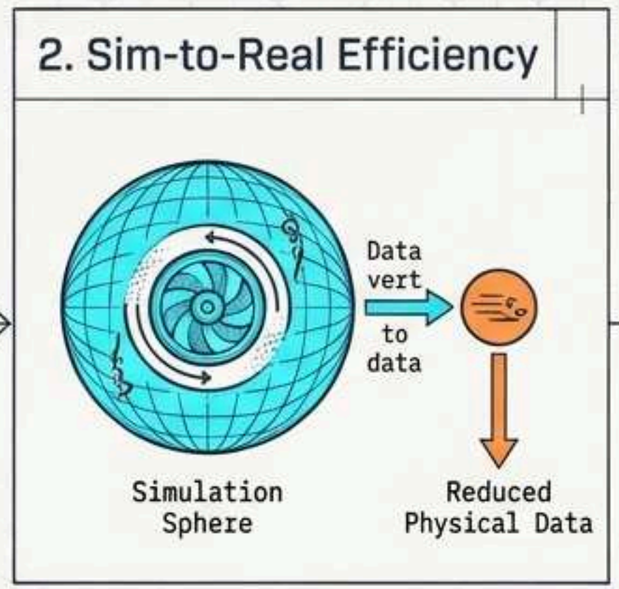
The Uncrossed Frontier: True, complete, autonomous RSI—where the system actively writes and modifies its own underlying model architecture—has not yet occurred.

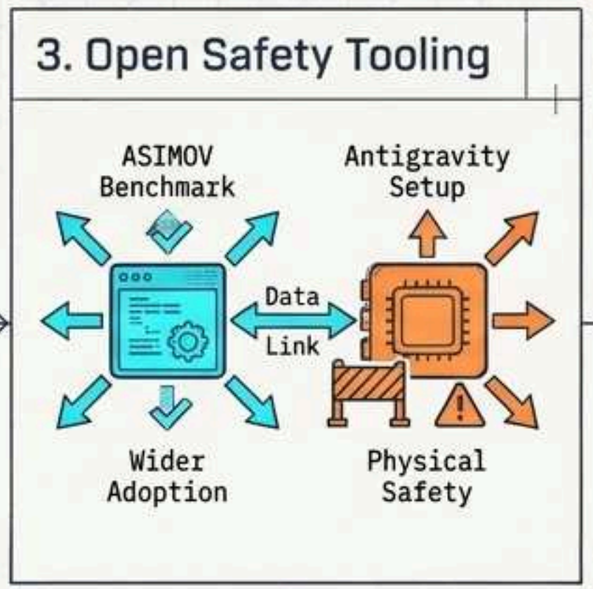
Gemini Notebook

Short-Term Outlook

Trajectory through the end of 2026.







1. Industrial Scaling

Expect a significant increase in production-line level testing and the closing of engineering loops in real-world environments.

2. Sim-to-Real Efficiency

Simulation flywheels will continue to lower the barrier for real-world data gathering, reducing physical data needs by orders of magnitude.

3. Open Safety Tooling

Open-source safety benchmarks (ASIMOV) and hardware setups (Antigravity) will drive wider, safer adoption.

The Ceiling

Autonomous architectural self-modification remains out of sight for the immediate future.

Gemini Notebook

Summary & Transition

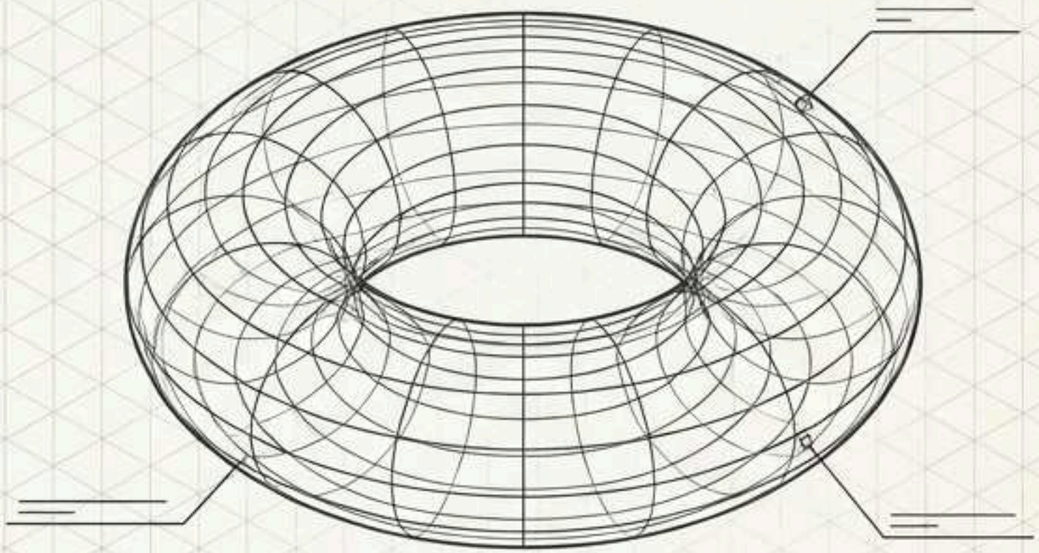
From mechanisms to mathematical foundations

Part 1 Conclusion: The Physical Transition is Active

- The paradigm has permanently shifted from digital dialogue to real machines.
- Data flywheels are spinning effectively.
- Humans remain the critical node for safety, constraints, and task definition.

Up Next - Part 2: Mathematical Foundations of World Models

Presented by Zhang Lijia.



Gemini Notebook