Digital Twins That Close the Loop

Agentic AI for Robot Fleets

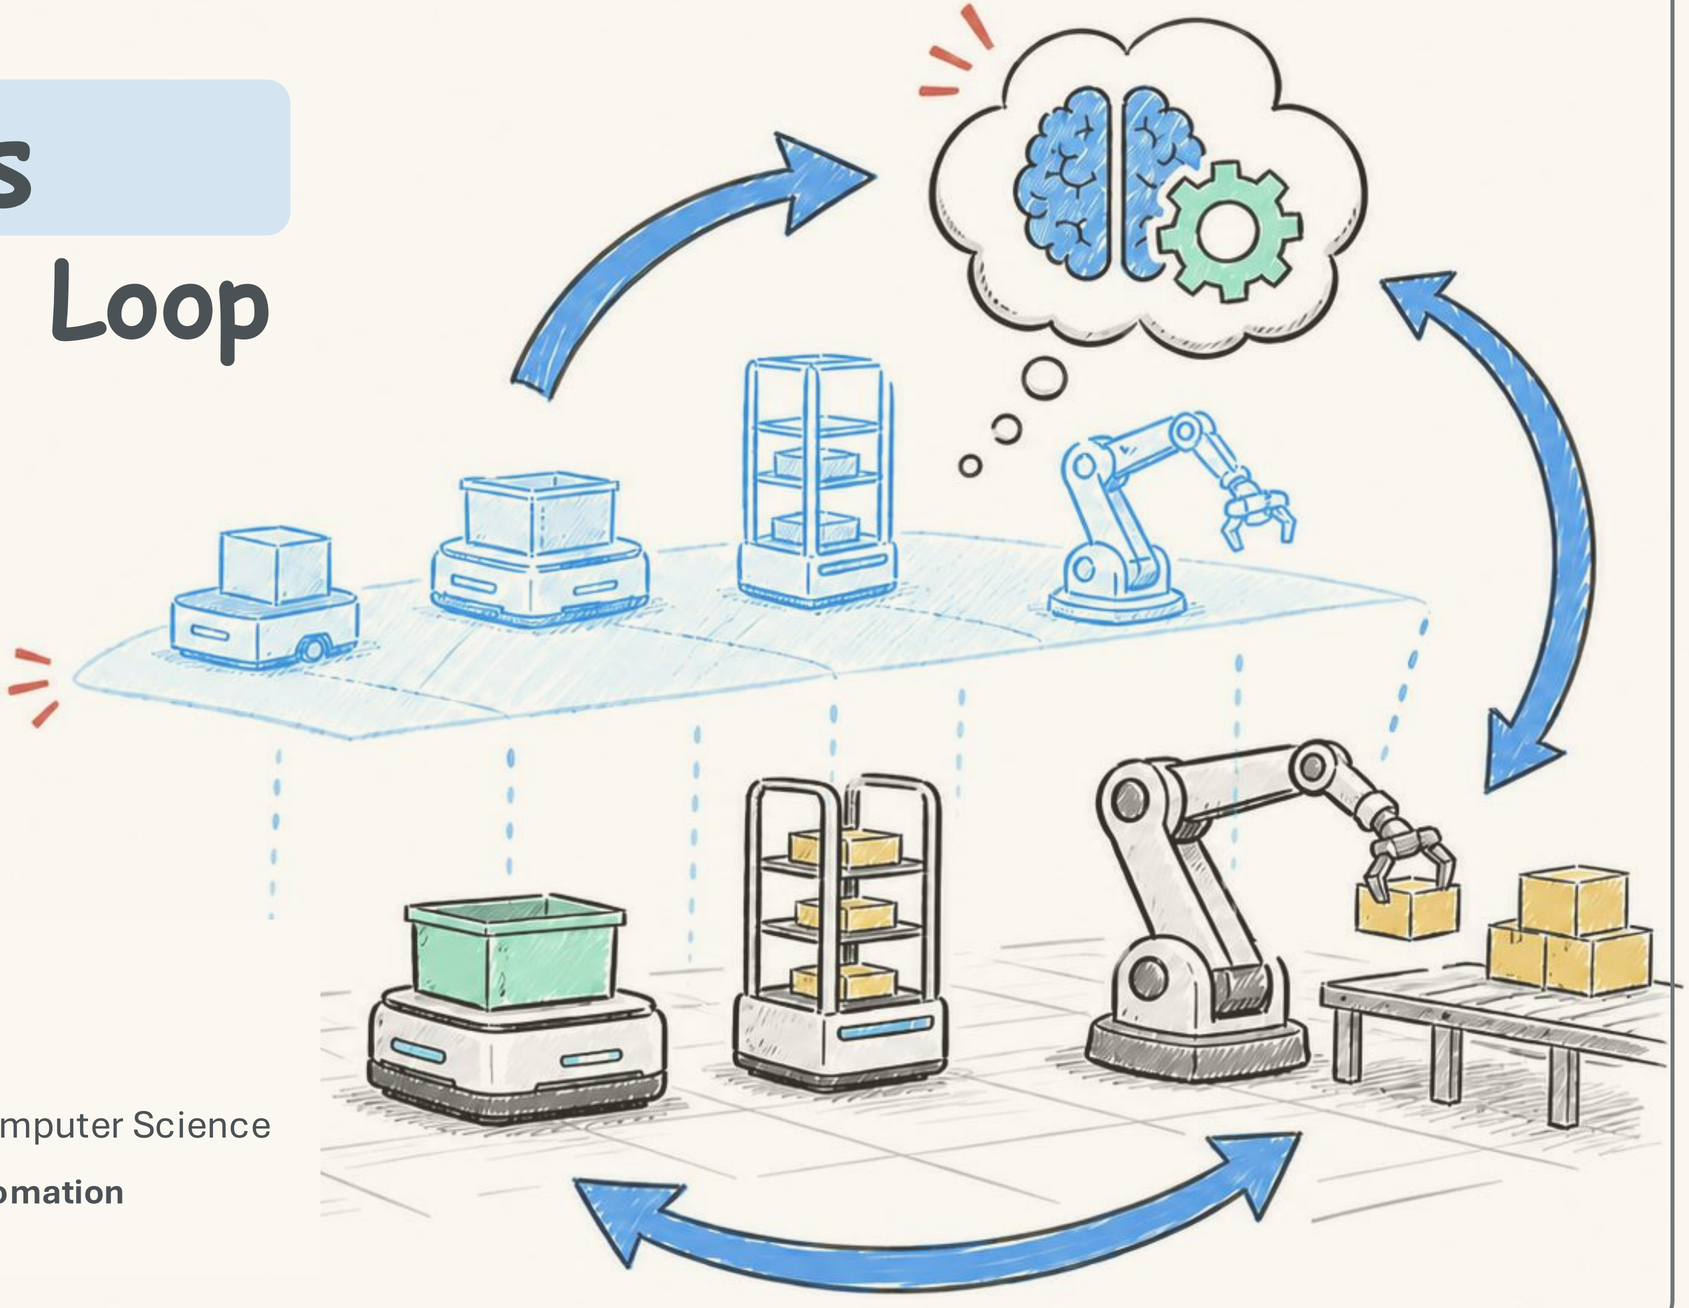

Prakash Aryan
Doctoral Researcher, University of Bern
Software Engineering Group, Institute of Computer Science

Workshop: Agentic AI and Digital Twin for Automation
BITS Pilani, Dubai Campus | 18 August 2026

Robot fleets already act. Their intelligence is fragmented.

LOCAL AUTONOMY

Navigate
Avoid obstacles
Dock

→

FLEET MANAGER

Assign jobs
Resolve traffic
Manage charge

→

HUMAN OPERATIONS

Interpret exceptions
Change priorities
Own consequences

The missing layer is shared operational reasoning across the whole fleet.

Four steps from mirror to bounded action

1 → 2 → 3 → 4

TWIN

Build a synchronized operational world model

FLEET

Find decisions that only exist at system level

AGENT

Reason through tools, simulation, and optimization

ASSURE

Gate, test, observe, and earn autonomy

The goal is not an LLM that drives robots. It is an accountable system that closes a useful loop.

A digital twin is a synchronization relationship

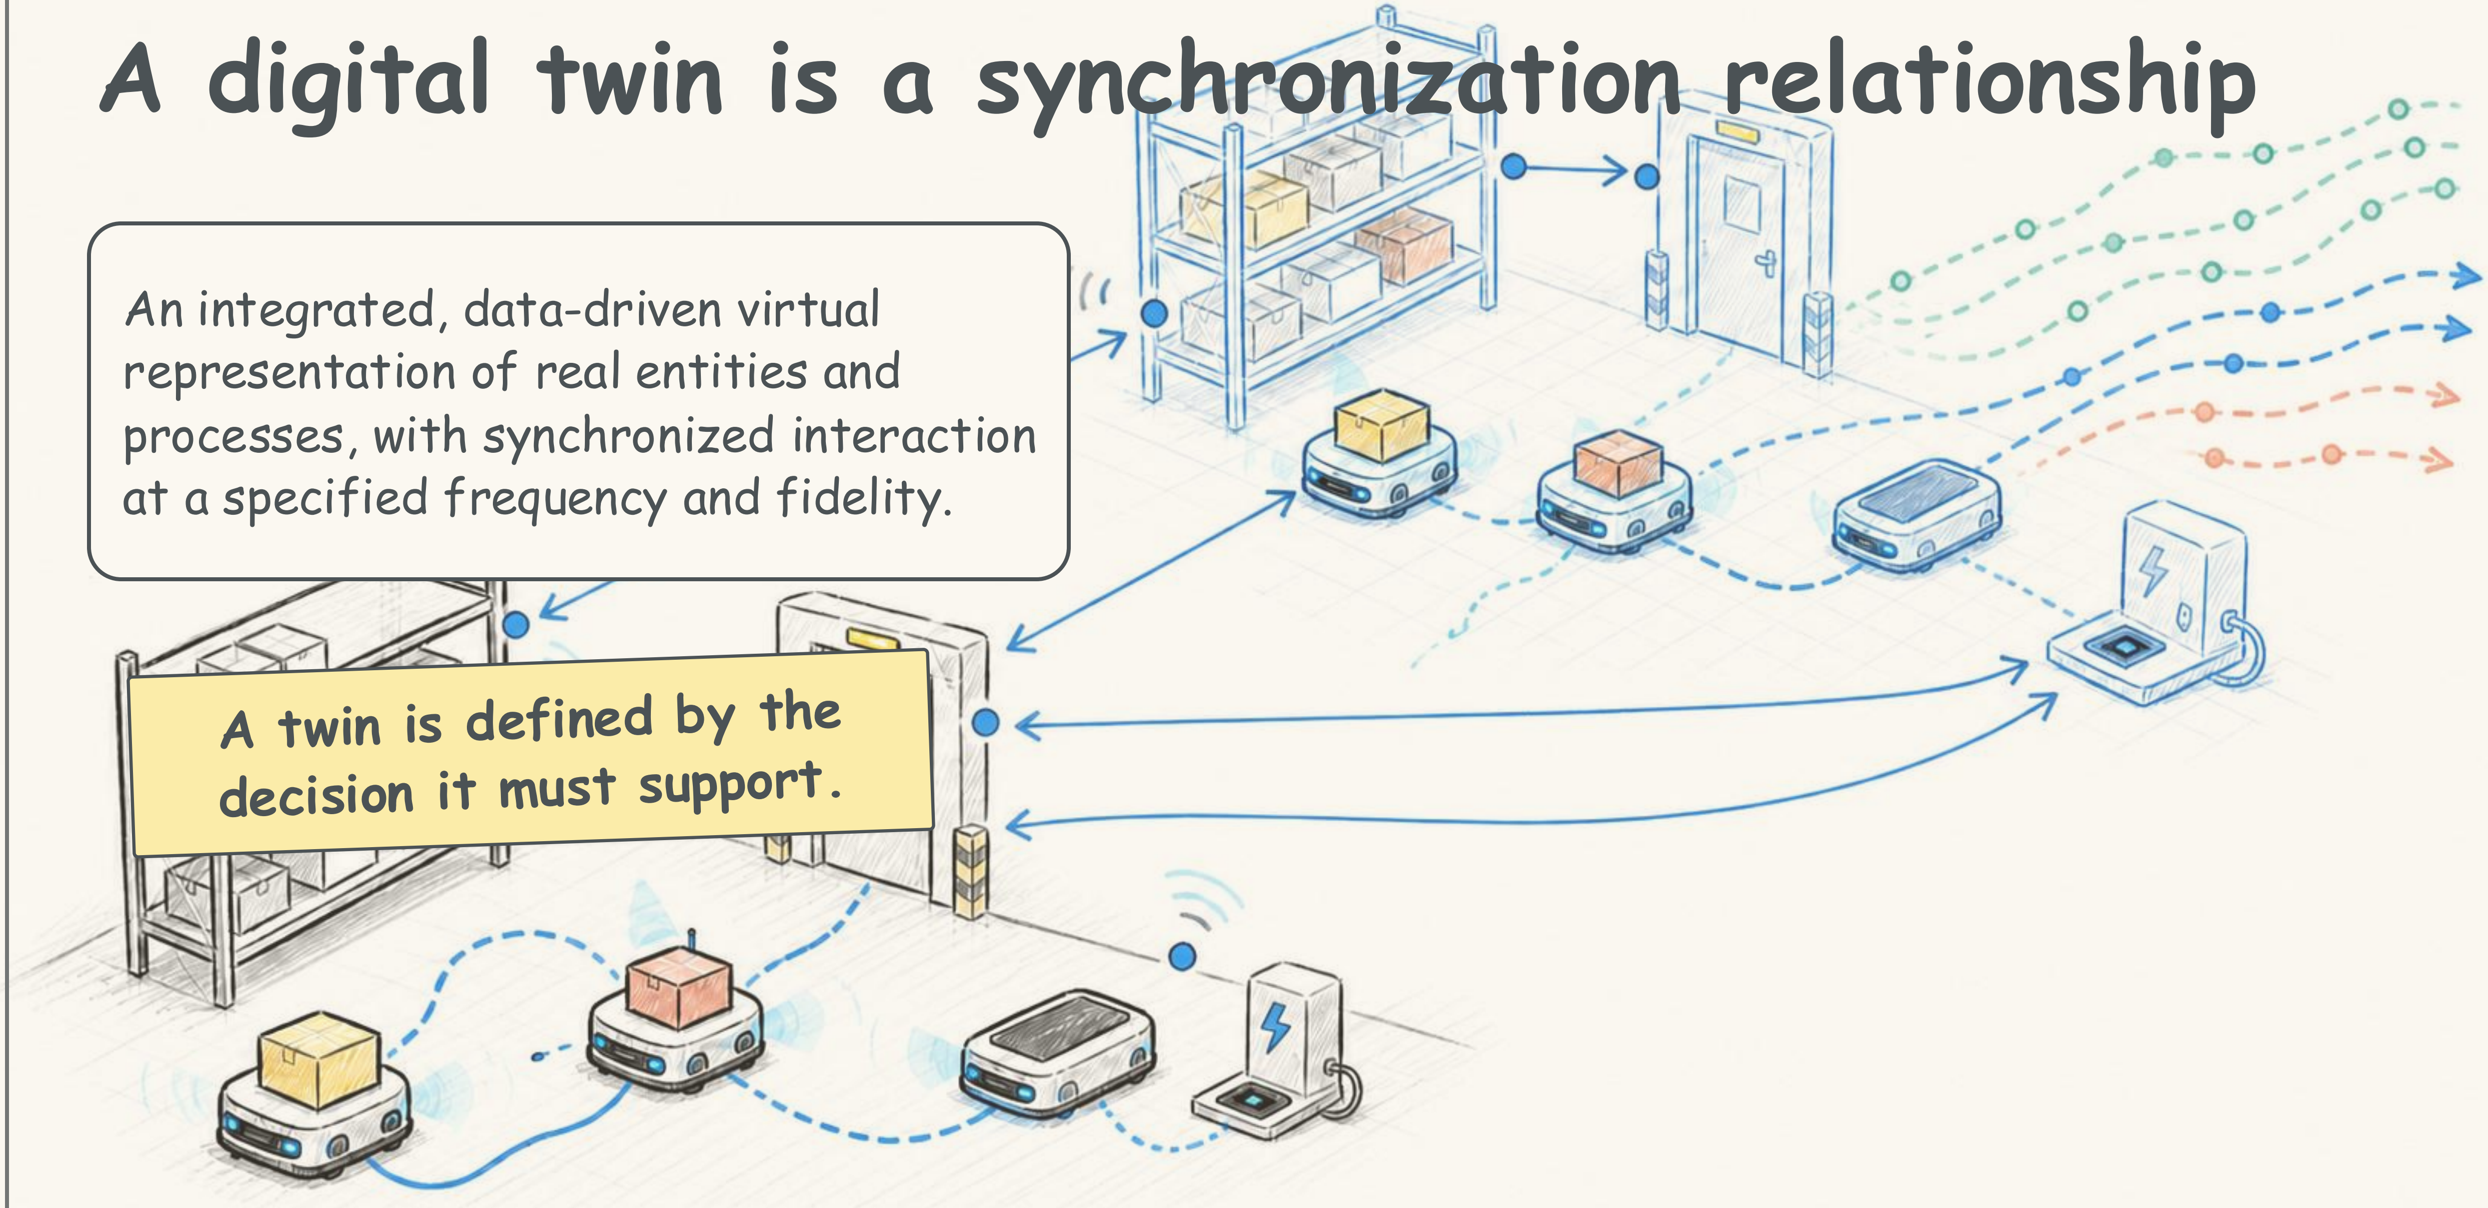

An integrated, data-driven virtual representation of real entities and processes, with synchronized interaction at a specified frequency and fidelity.

A twin is defined by the decision it must support.

Model, shadow, or twin?

MODEL	DIGITAL SHADOW	DIGITAL TWIN
Offline representation No live state No operational action	Physical → virtual Live observation No direct intervention	Physical ↔ virtual Observation and intervention Explicit authority
○	→	↔

If the virtual side cannot be trusted to observe or influence operations, the agent has no stable world to reason about.

Every twin makes a synchronization contract

STATE	FREQUENCY	FIDELITY	AUTHORITY
What is represented? Pose, battery, task, health, map, queue	How fresh must it be? Milliseconds, events, minutes, shifts	How accurate is enough? Routing, energy, wear, human flow	What may flow back? Advice, dispatch, parameters, stop

Staleness is not only a data problem. In a closed loop, it becomes a control risk.

A fleet is a system of systems

Robots compete and cooperate through shared:

- space
- jobs
- chargers
- infrastructure
- people
- time

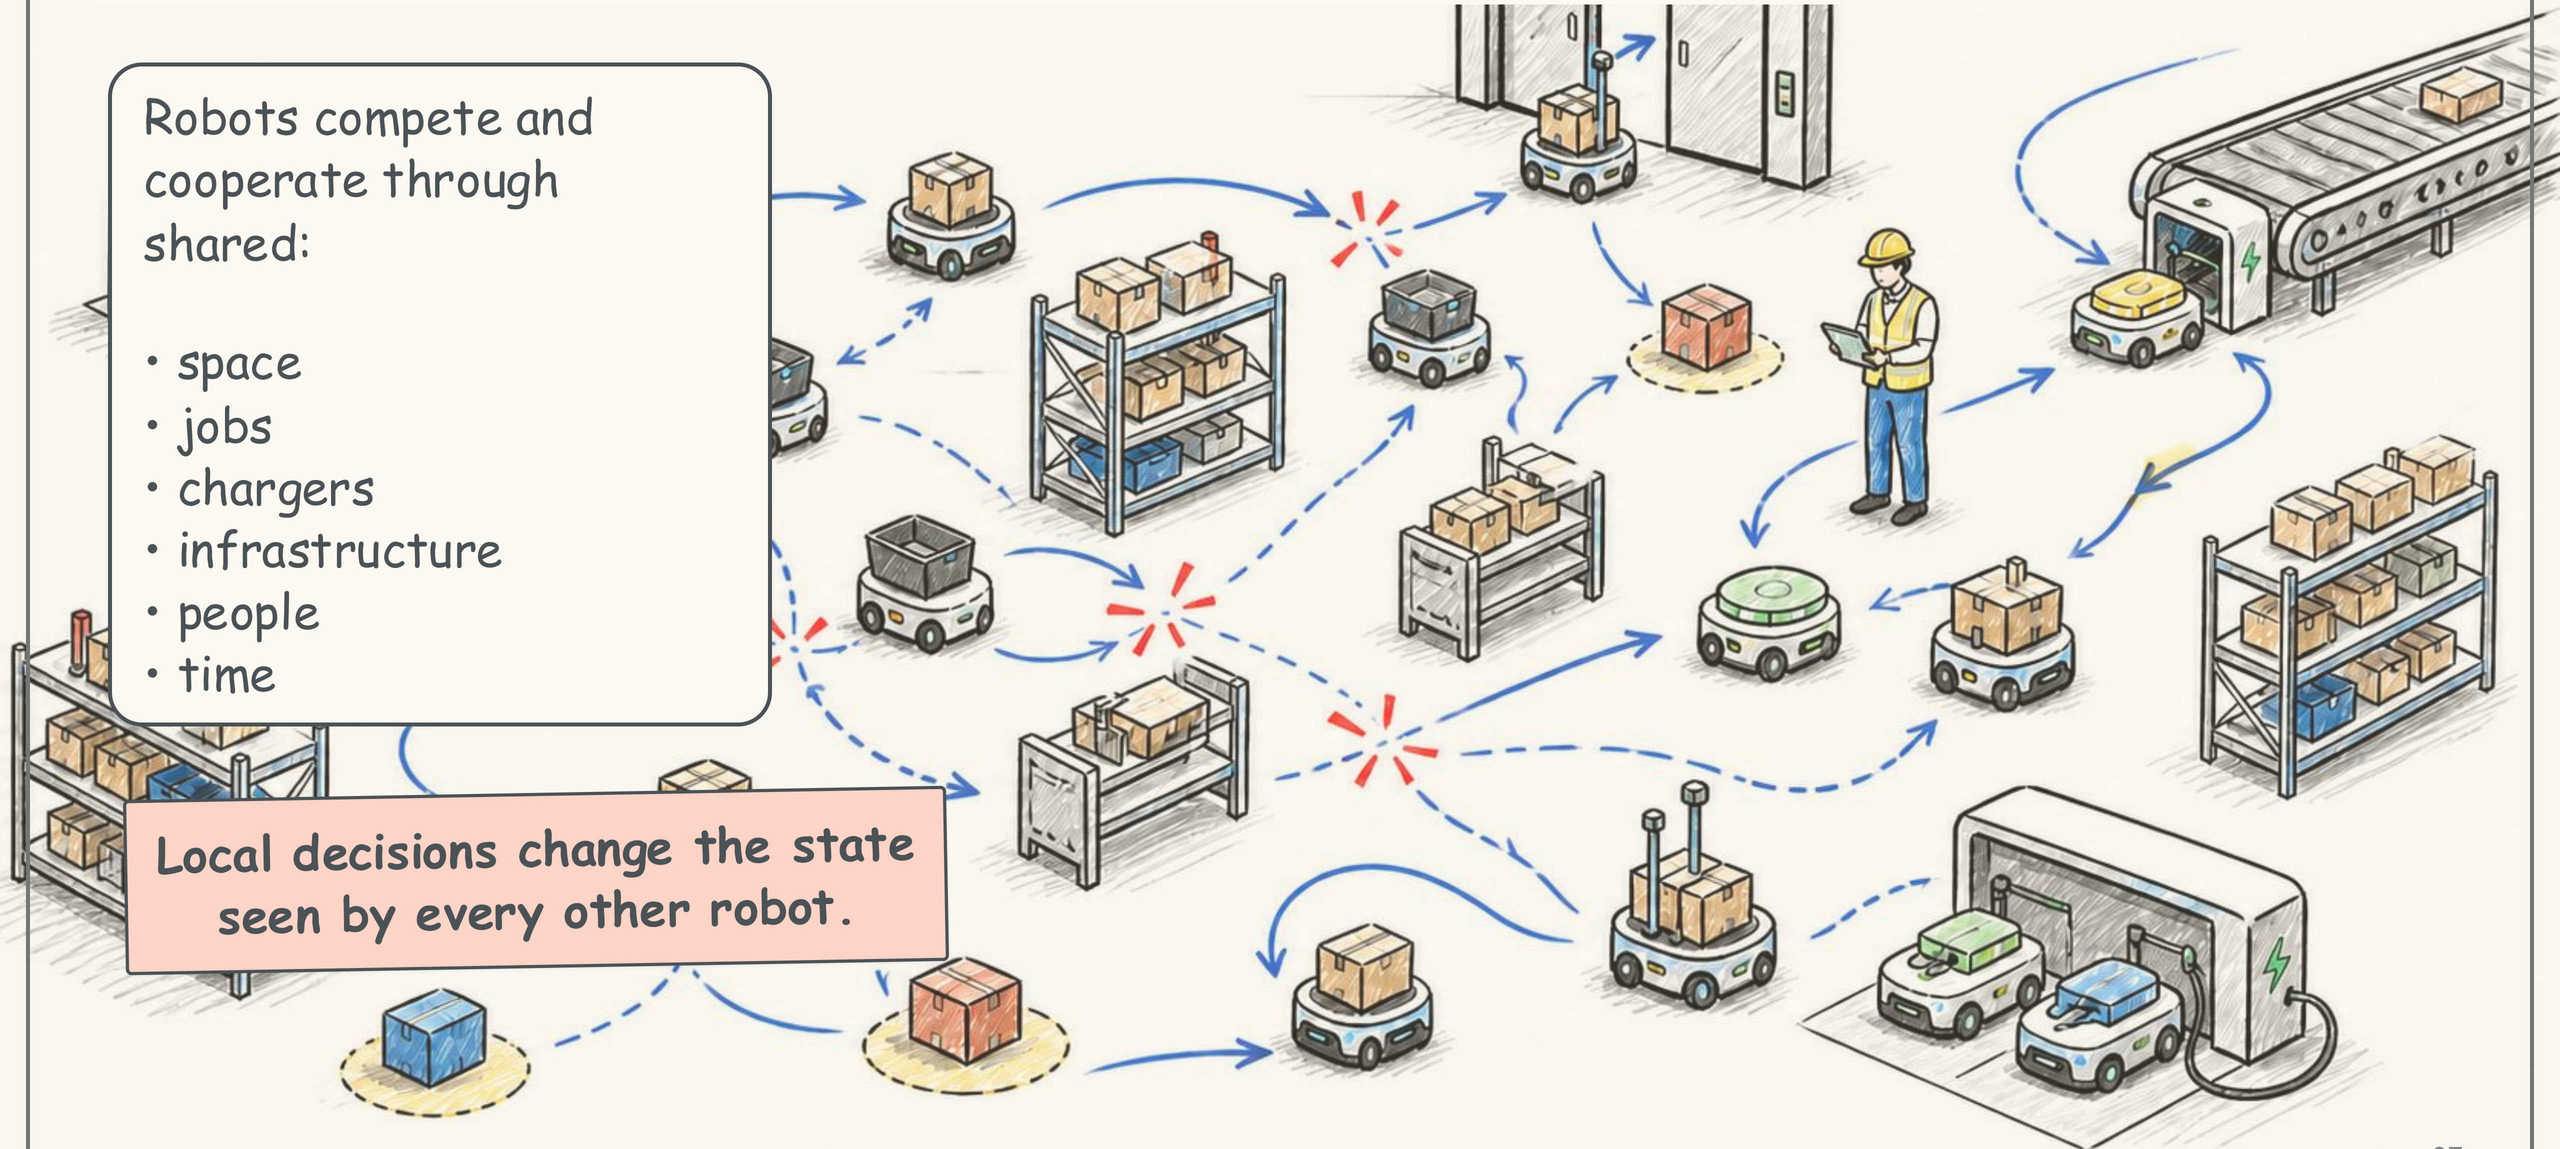

Local decisions change the state seen by every other robot.

Fleet objectives are coupled

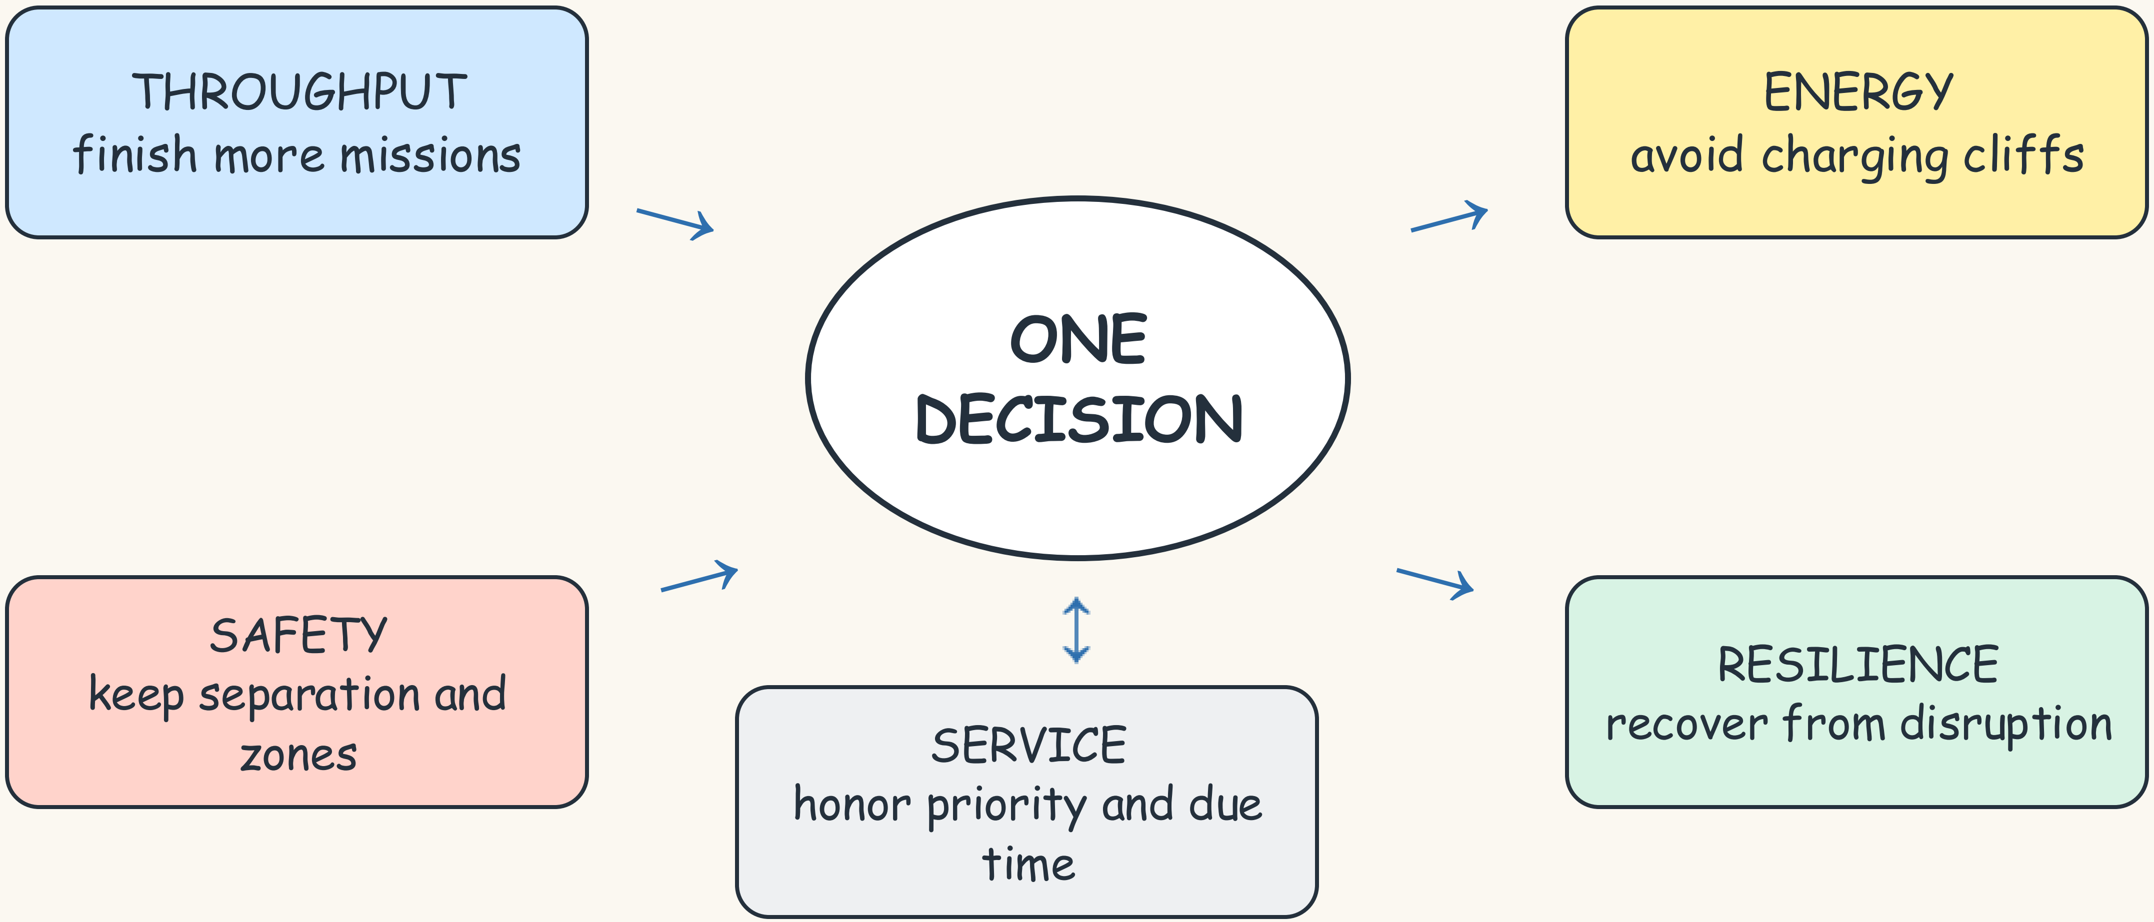

Agentic AI adds goal-directed tool use

TRADITIONAL PIPELINE

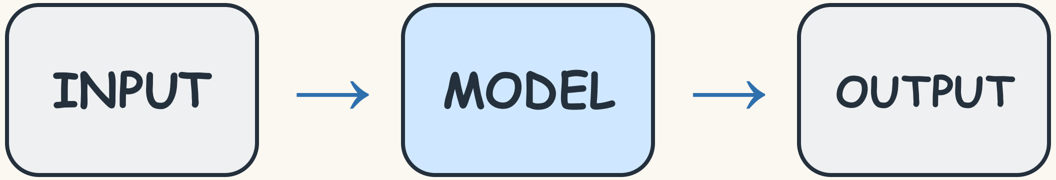

one pass

AGENTIC LOOP

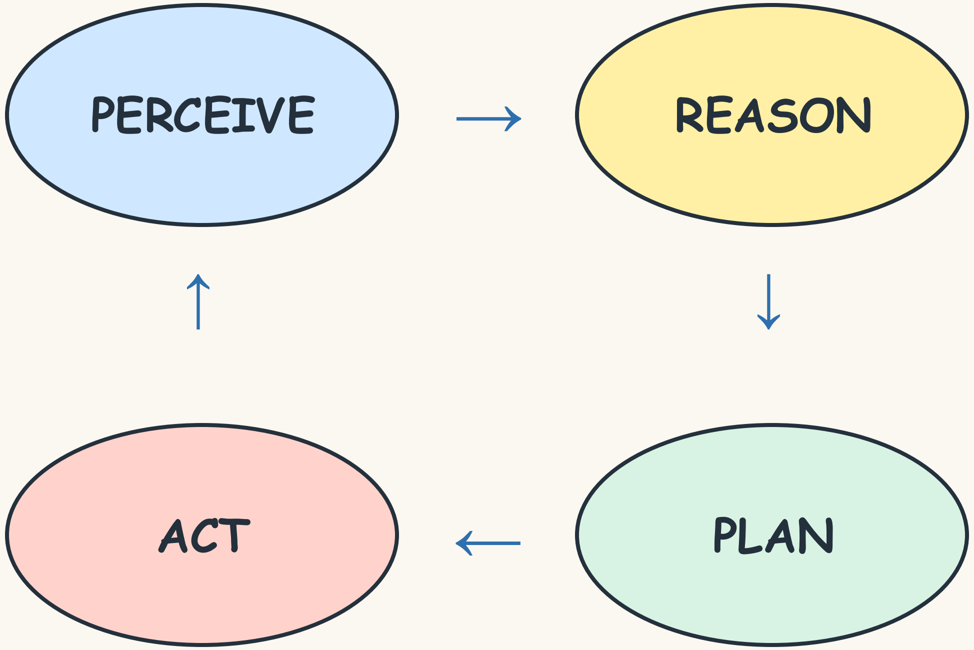

Agency appears at the system level: memory + tools + feedback + policy.

Keep the language model at the semantic layer

LANGUAGE AGENT	Interpret goals, policies, exceptions, and unstructured reports
↓	
OPTIMIZERS + SIMULATORS	Generate feasible schedules, routes, allocations, and predicted outcomes
↓	
FLEET + ROBOT CONTROL	Dispatch tasks, negotiate traffic, execute trajectories, enforce safety

The LLM should not close the millisecond control loop.

Twin = operational world model

OBSERVE
Twin state

DIAGNOSE
Evidence + context

SIMULATE
Candidate futures

ACT + LEARN
Bounded intervention

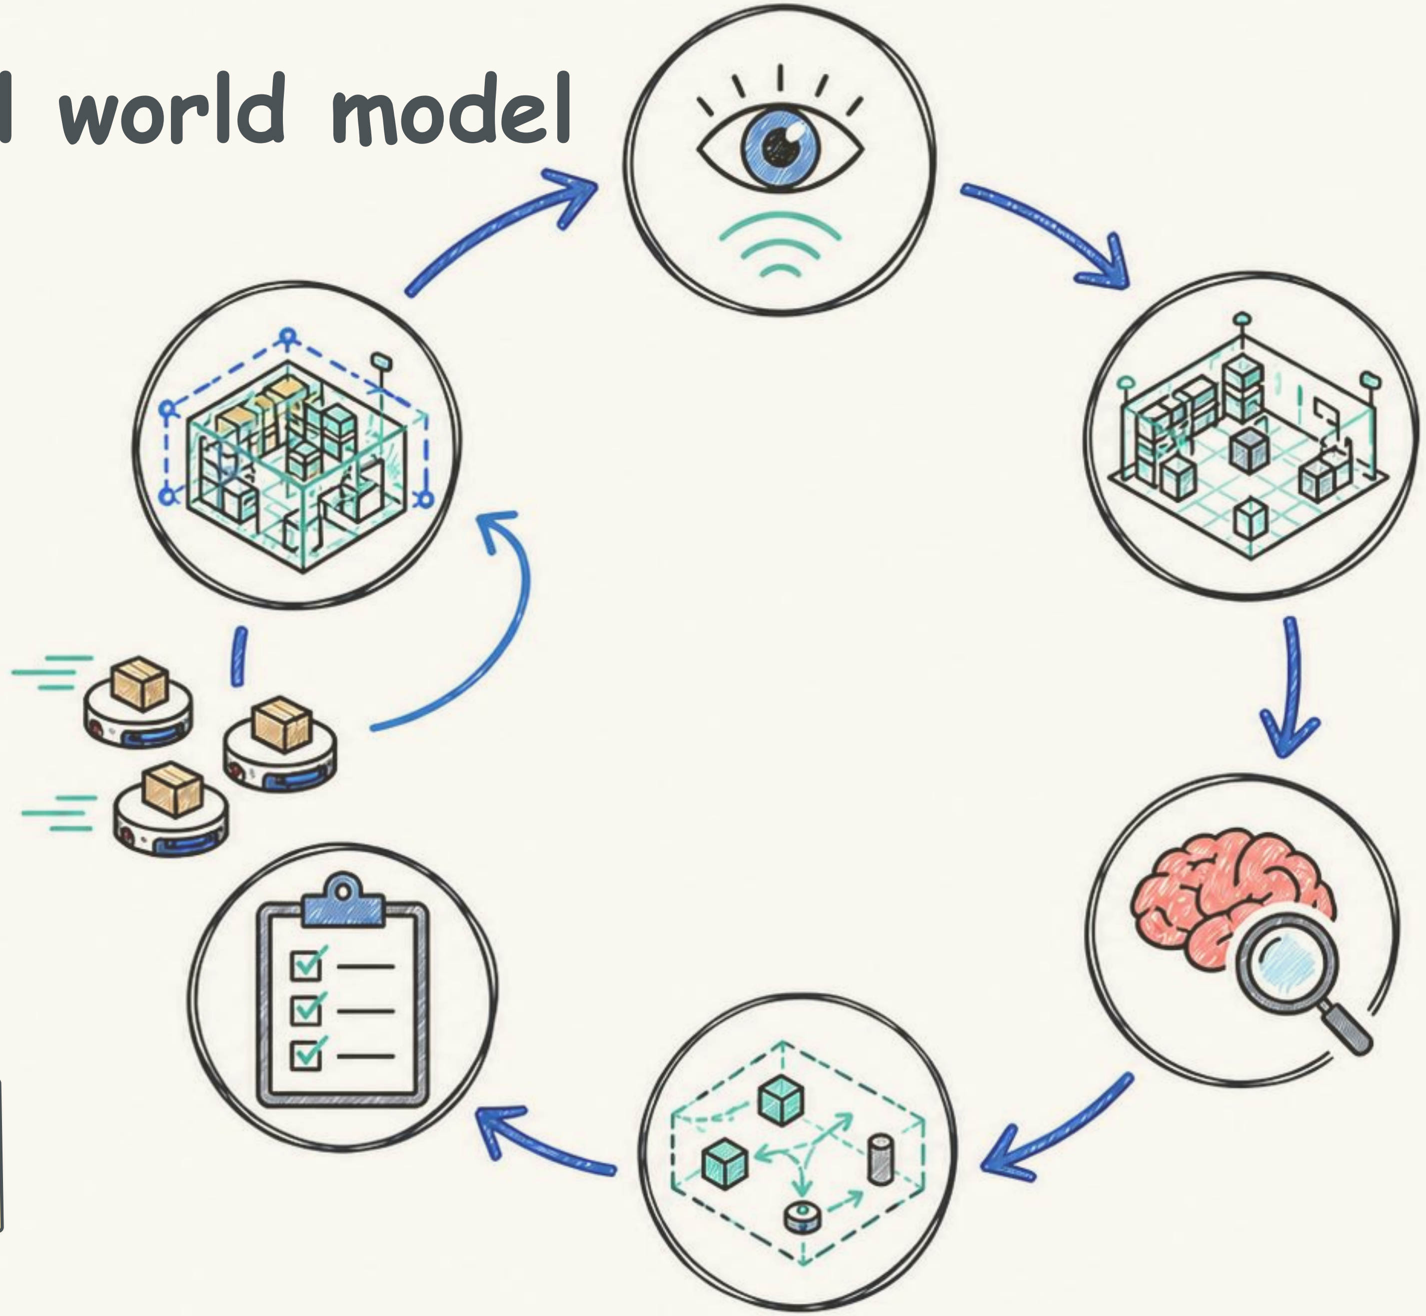

The twin is both memory and rehearsal space.

Reference architecture: agent + twin + fleet

HUMAN + POLICY	goals | approval | escalation | accountability
AGENT	memory | reasoning | tool selection | explanation
TWIN + TOOLS	state graph | history | simulation | optimizer | rules
INTEROPERABILITY	events | schemas | identity | fleet adapters | audit
FLEET MANAGEMENT	task allocation | traffic | charging | recovery
PHYSICAL SYSTEM	robots | chargers | doors | conveyors | people

Give the agent tools, not raw authority

READ	ANALYZE	ACT
query_fleet_state	simulate_candidate	propose_dispatch
get_incident_context	solve_allocation	request_zone_change
retrieve_policy	check_constraints	start_canary

Typed · scoped · authorized · reversible · logged

Illustrative scenario: blocked aisle

40 AMRs

Blocked aisle
Urgent order at risk
Two robots queued
One charger occupied
One battery near reserve

Goal: recover service without trading away safety or charge reserve.

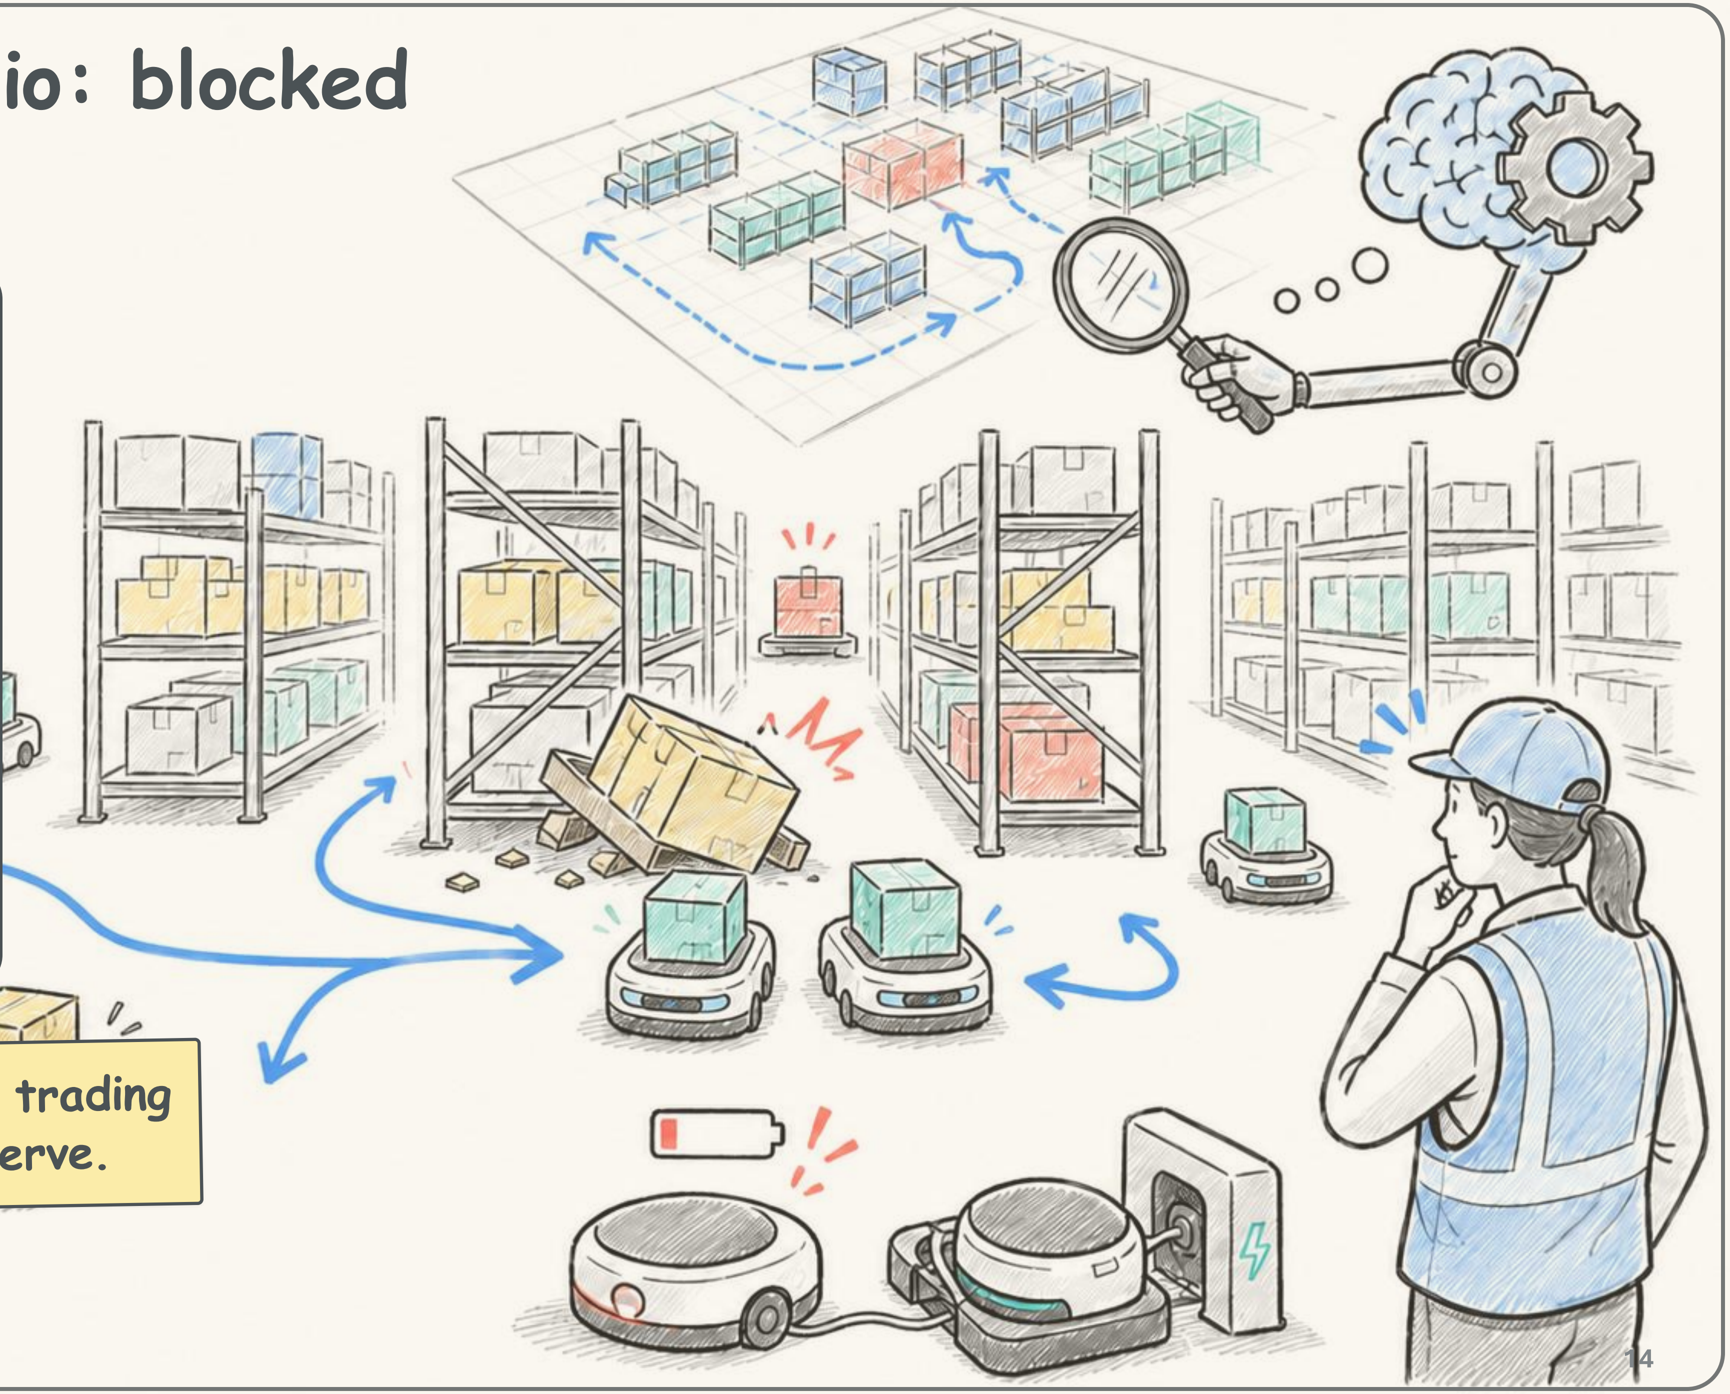

Step 1: observe and diagnose from evidence

TWIN OBSERVATIONS

Aisle B-7 blocked

Order O-184 due in 11 min

Robot R-12 at 14%

Charger C-2 free in 6 min

→

DIAGNOSIS

The delay is not only a path problem. It is a coupled traffic, task, and charging problem.

Evidence confidence: high
Twin freshness: 1.2 s

Ground every conclusion in observable state.

Step 2: compare candidate futures

Candidate	Predicted service	Energy	Risk	
A Reroute all	Order on time	High detour	Congestion	
B Reassign + reroute	Order on time	Reserve held	Low	CHOOSE
C Wait for clearance	Order late	Lowest	Low	

Language proposes. Simulation and optimization provide the evidence.

Illustrative predicted outcomes

Step 3: act, verify, and roll back

1 → 2 → 3 → 4

APPROVE

Operator accepts the bounded recovery plan

CANARY

Apply to two tasks and one zone first

MONITOR

Compare delay, battery, and conflicts to prediction

ROLL BACK

Return to the safe schedule if a limit trips

Exit conditions: safety event · KPI deviation · stale twin · tool failure · human override

A decision needs a record, not a rationale

Trigger	Blocked aisle B-7 at 18:22:14
Evidence	State snapshot + freshness + confidence
Candidates	A / B / C with predicted KPIs
Constraints	Safety zones, battery floor, order SLA
Authorization	Approved by operator 23; scope: zone B
Outcome	Delay -7 min; no safety event; twin error 4%

Store state and consequences, not hidden chain-of-thought.

Autonomy is a policy, not a model feature

0
OBSERVE

Twin only

1
EXPLAIN

Diagnose and summarize

2
RECOMMEND

Rank plans for a human

3
APPROVAL

Prepare action, wait for consent

4
BOUNDED AUTO

Act inside policy and stop conditions

More authority requires more evidence, narrower scope, and stronger recovery.

Safety gates belong below the agent

The agent proposes intent.

Hard constraints decide feasibility.

Safety controllers retain final authority.

A valid tool call can still be an unsafe physical action.

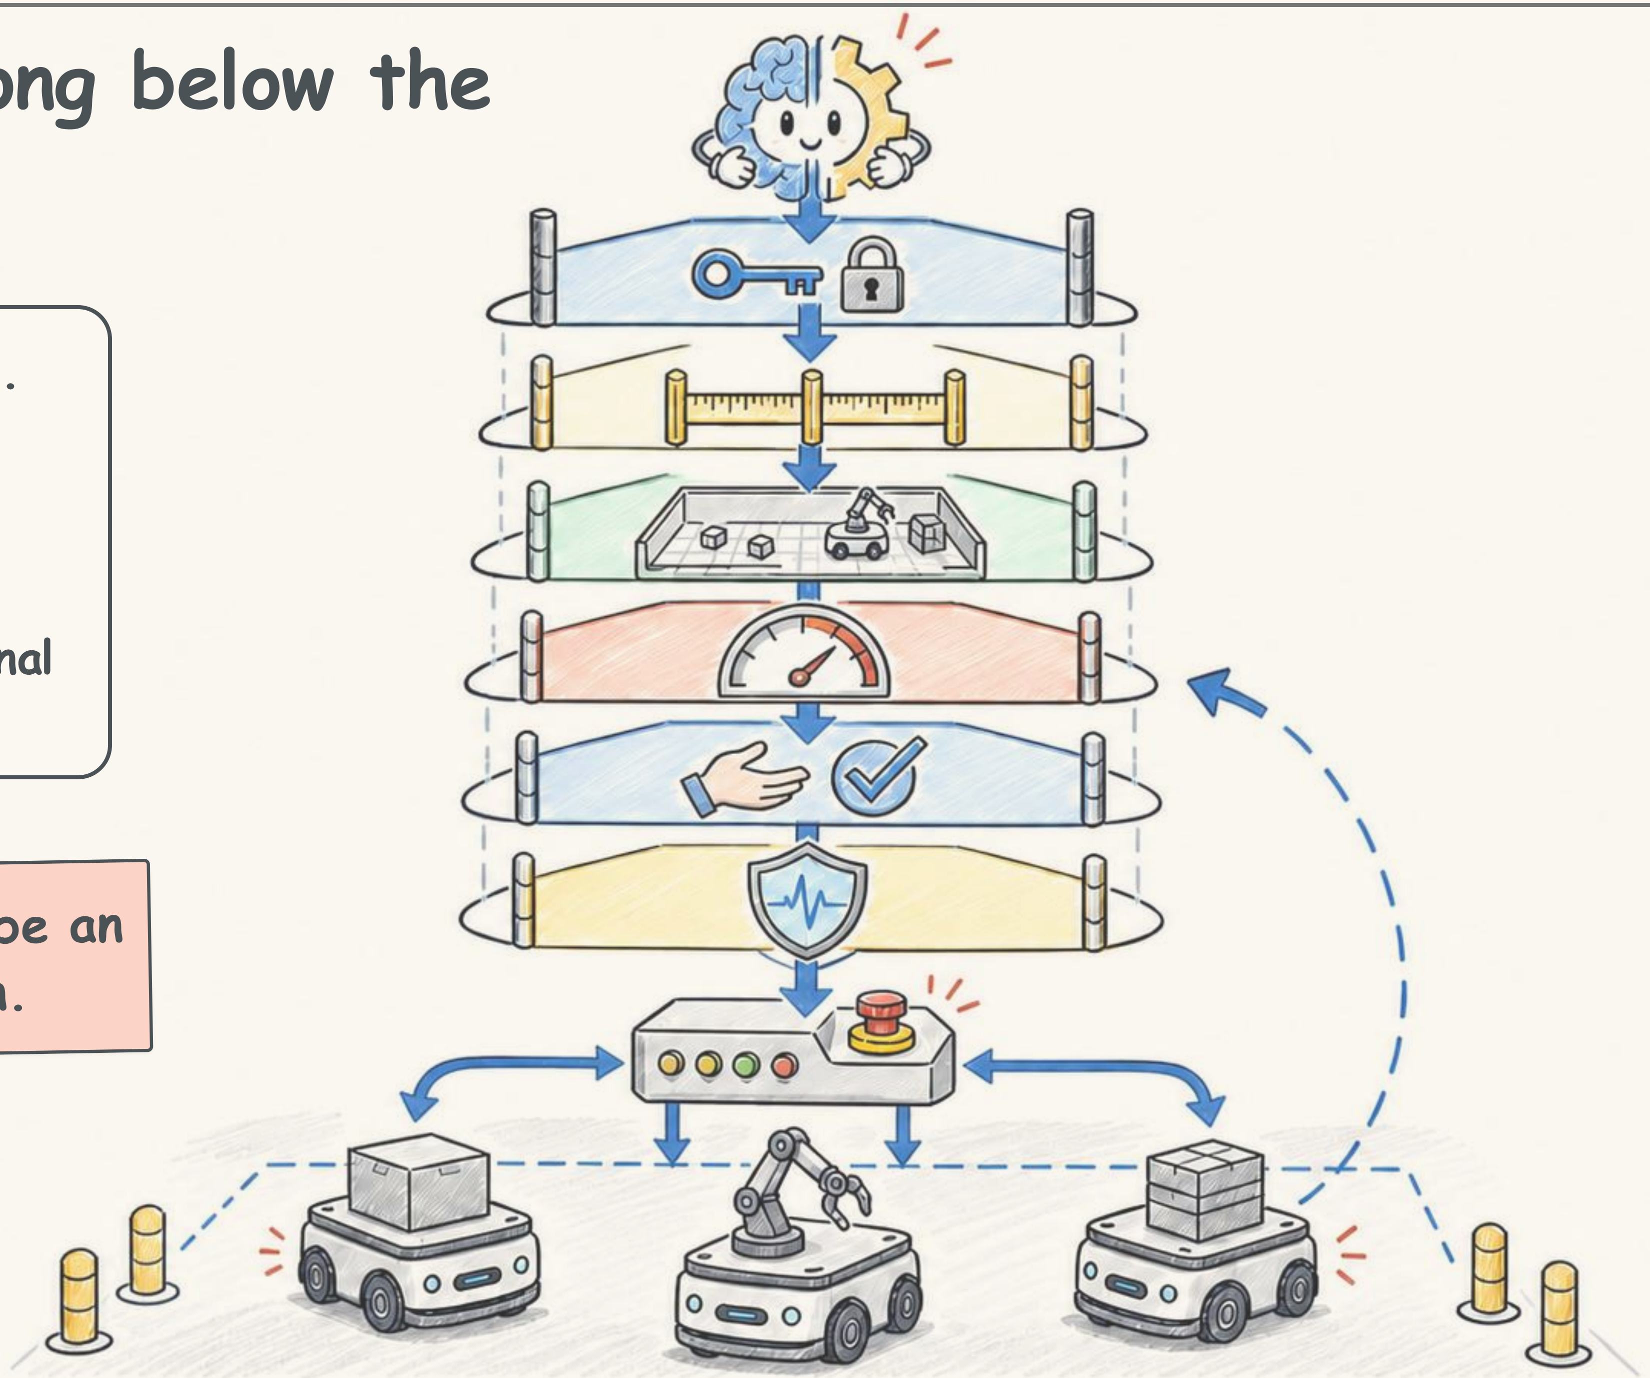

Two coupled loops can amplify errors

STALE TWIN
agent reasons about the past

OBJECTIVE DRIFT
proxy wins, mission loses

DATA POISONING
false state enters the twin

TWIN
↔
AGENT

SIM GAP
predicted future is not credible

TOOL ERROR
valid syntax, wrong intent

CONFLICT CASCADE
one fix creates another queue

Test the policy before the plant floor

1 REPLAY Historical incidents and edge cases

2 SHADOW Live recommendations with no authority

3 COUNTERFACTUAL Twin simulations against alternative actions

4 STRESS Latency, packet loss, stale data, adversarial events

5 CANARY Small scope, tight limits, immediate rollback

6 RUNTIME Continuous monitors and human override

Interoperability creates safe handles

ISO 23247	VDA 5050 v3	Open-RMF	ISO 21423
Manufacturing twin architecture and information exchange	Command and status interface for mixed mobile-robot fleets	Task allocation, traffic negotiation, and shared infrastructure	Emerging AMR communication and interoperability standard

Interoperability defines what can be exchanged. Safety still needs its own standard, control logic, and verification.

For driverless industrial trucks and AMRs, see ISO 3691-4:2023.

Earn autonomy in stages

INSTRUMENT
1
trusted telemetry

TWIN
2
synchronized state

SHADOW AGENT
3
diagnose only

RECOMMEND
4
rank options

APPROVAL
5
prepare + wait

BOUNDED AUTO
6
act + recover

Do not automate a decision you cannot observe, simulate, constrain, and reverse.

A practical first 90 days

DAYS 0-30

Choose one decision

Map state, authority, and failure cost

Instrument freshness and data quality

DAYS 31-60

Build the twin slice

Add replay and scenario simulation

Expose read-only tools

DAYS 61-90

Run a shadow agent

Compare recommendations to operators

Define gates for a canary

Best first target: frequent, costly, bounded, and currently manual.

Research questions that matter now

1. How should twin fidelity adapt to the decision and risk?

2. How do we measure agent value beyond task success?

3. How can fleet memory remain useful without becoming stale?

4. How should humans supervise many interacting agents and robots?

5. What evidence is sufficient to earn more action authority?

Three ideas to carry into your next system

1

THE TWIN

is a synchronized decision model, not a dashboard

2

THE AGENT

reasons through tools, models, and constraints

3

AUTONOMY

is earned through policy, evidence, and recovery

Discussion: What decision in your operation is frequent, costly, and still manual?

Prakash Aryan | prakash.aryan@unibe.ch | github.com/prakash-aryan

References I: digital twins and assurance

Digital Twin Consortium. Definition of a Digital Twin.
https://www.digitaltwinconsortium.org/initiatives/the-definition-of-a-digital-twin/

NIST. Digital Twins for Advanced Manufacturing.
https://www.nist.gov/programs-projects/digital-twins-advanced-manufacturing

ISO 23247-2:2021. Digital twin framework for manufacturing: reference architecture.
https://www.iso.org/standard/78743.html

Shao, Hightower, and Schindel. Credibility Consideration for Digital Twins in Manufacturing.
https://www.nist.gov/publications/credibility-consideration-digital-twins-manufacturing

Voas et al. NIST IR 8356: Security and Trust Considerations for Digital Twin Technology.
https://doi.org/10.6028/NIST.IR.8356

NIST. AI Risk Management Framework and Playbook.
https://airc.nist.gov/airmf-resources/airmf/

ISO 3691-4:2023. Safety requirements for driverless industrial trucks and systems.
https://www.iso.org/standard/83545.html

References II: agents, fleets, and interoperability

Yao et al. ReAct: Synergizing Reasoning and Acting in Language Models.
https://arxiv.org/abs/2210.03629

Li et al. Large Language Models for Multi-Robot Systems: A Survey.
https://arxiv.org/abs/2502.03814

Deng et al. LLMs and Digital Twins for Adaptive Multi-Robot Task Allocation.
https://arxiv.org/abs/2506.18178

Burr et al. Agentic Digital Twins: A Taxonomy of Capabilities.
https://arxiv.org/abs/2601.18799

Open-RMF demos and integration concepts.
https://github.com/open-rmf/rmf_demos

VDA 5050 version 3.0 release for mixed mobile-robot fleets.
https://www.vda.de/en/press/press-releases/2026/260421_PM_VDA_5050_EN

ISO/FDIS 21423. Industrial mobile robots: communications and interoperability.
https://www.iso.org/standard/86749.html