A New Architecture for Autonomy: Enabling AI Agents to Query External Knowledge in Intelligent Manufacturing

Introducing OPRA & COPRA: The Observation-Prompt-Response-Action Frameworks

A conceptual design for adaptive decision-making in dynamic, uncertain environments.

Authored by Vagan Terziyan, Oleksandra Vitko, and Oleksandr Terziyan

Modern Manufacturing Demands Unprecedented Intelligence and Adaptability

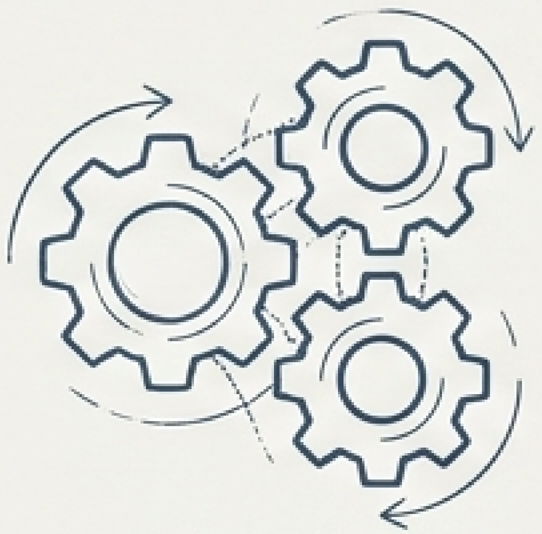

Industry 4.0 & 5.0

- Integration of data-driven intelligence, automation, and networked sensor systems.
- Emphasis on human-centered design, collaboration, and sustainability.

Digital Twins, Materials Informatics, Additive Manufacturing.

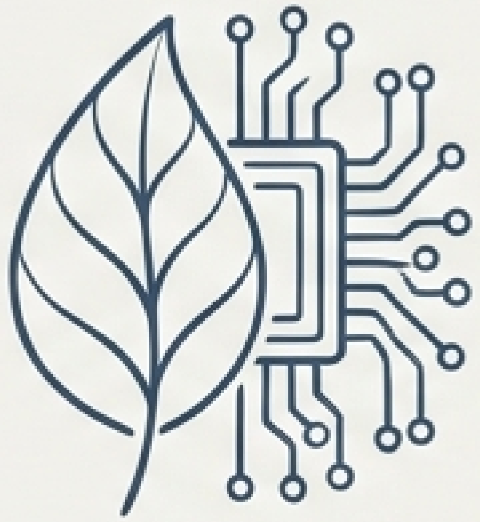

The Rise of Intelligent & Sustainable Manufacturing (ISM)

Core principles: Lean, green, and smart manufacturing.
Goal: Enhance operational excellence and sustainability through real-time decision-making.
Application: Optimizing production, energy use, and predictive maintenance.

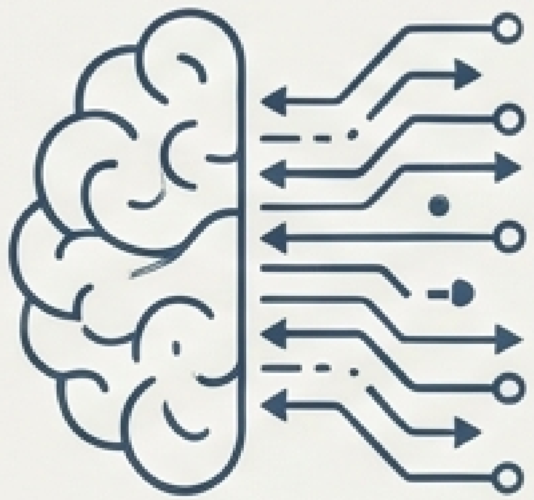

The Role of AI & LLMs

Knowledge-Informed Machine Learning (KIML) integrates domain knowledge for greater accuracy.

Large Language Models (LLMs) provide dynamic insights beyond pre-existing datasets, unlocking novel solutions.

Traditional Autonomous Agents Struggle When Faced with True Novelty and Uncertainty

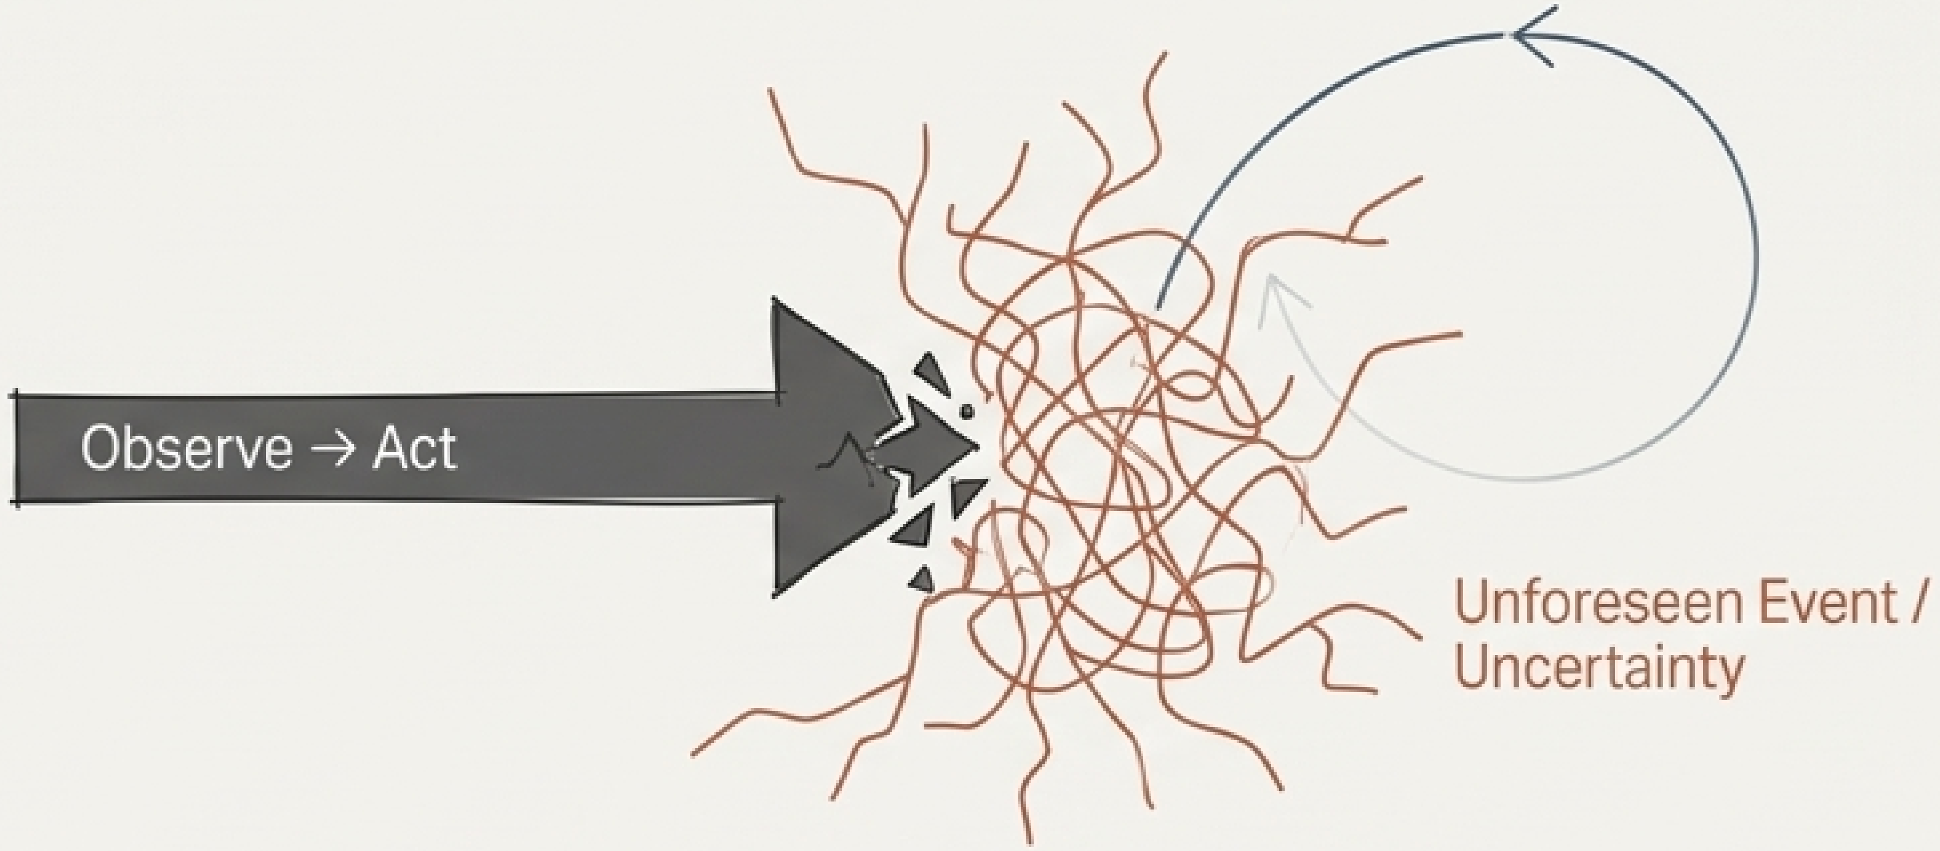

Current agent models (Reactive, BDI, Learning) primarily rely on pre-programmed knowledge, policies, or static rule-bases.

The Critical Limitation: They are effective in well-understood domains but lack the flexibility to handle situations that are:

- Novel (not in training data)
- Complex (multiple interacting factors)
- Uncertain (incomplete or ambiguous data)

This rigidity creates a bottleneck for achieving truly autonomous and resilient manufacturing systems.

The OPRA Framework Introduces an External Knowledge Loop to Bridge an Agent’s Understanding

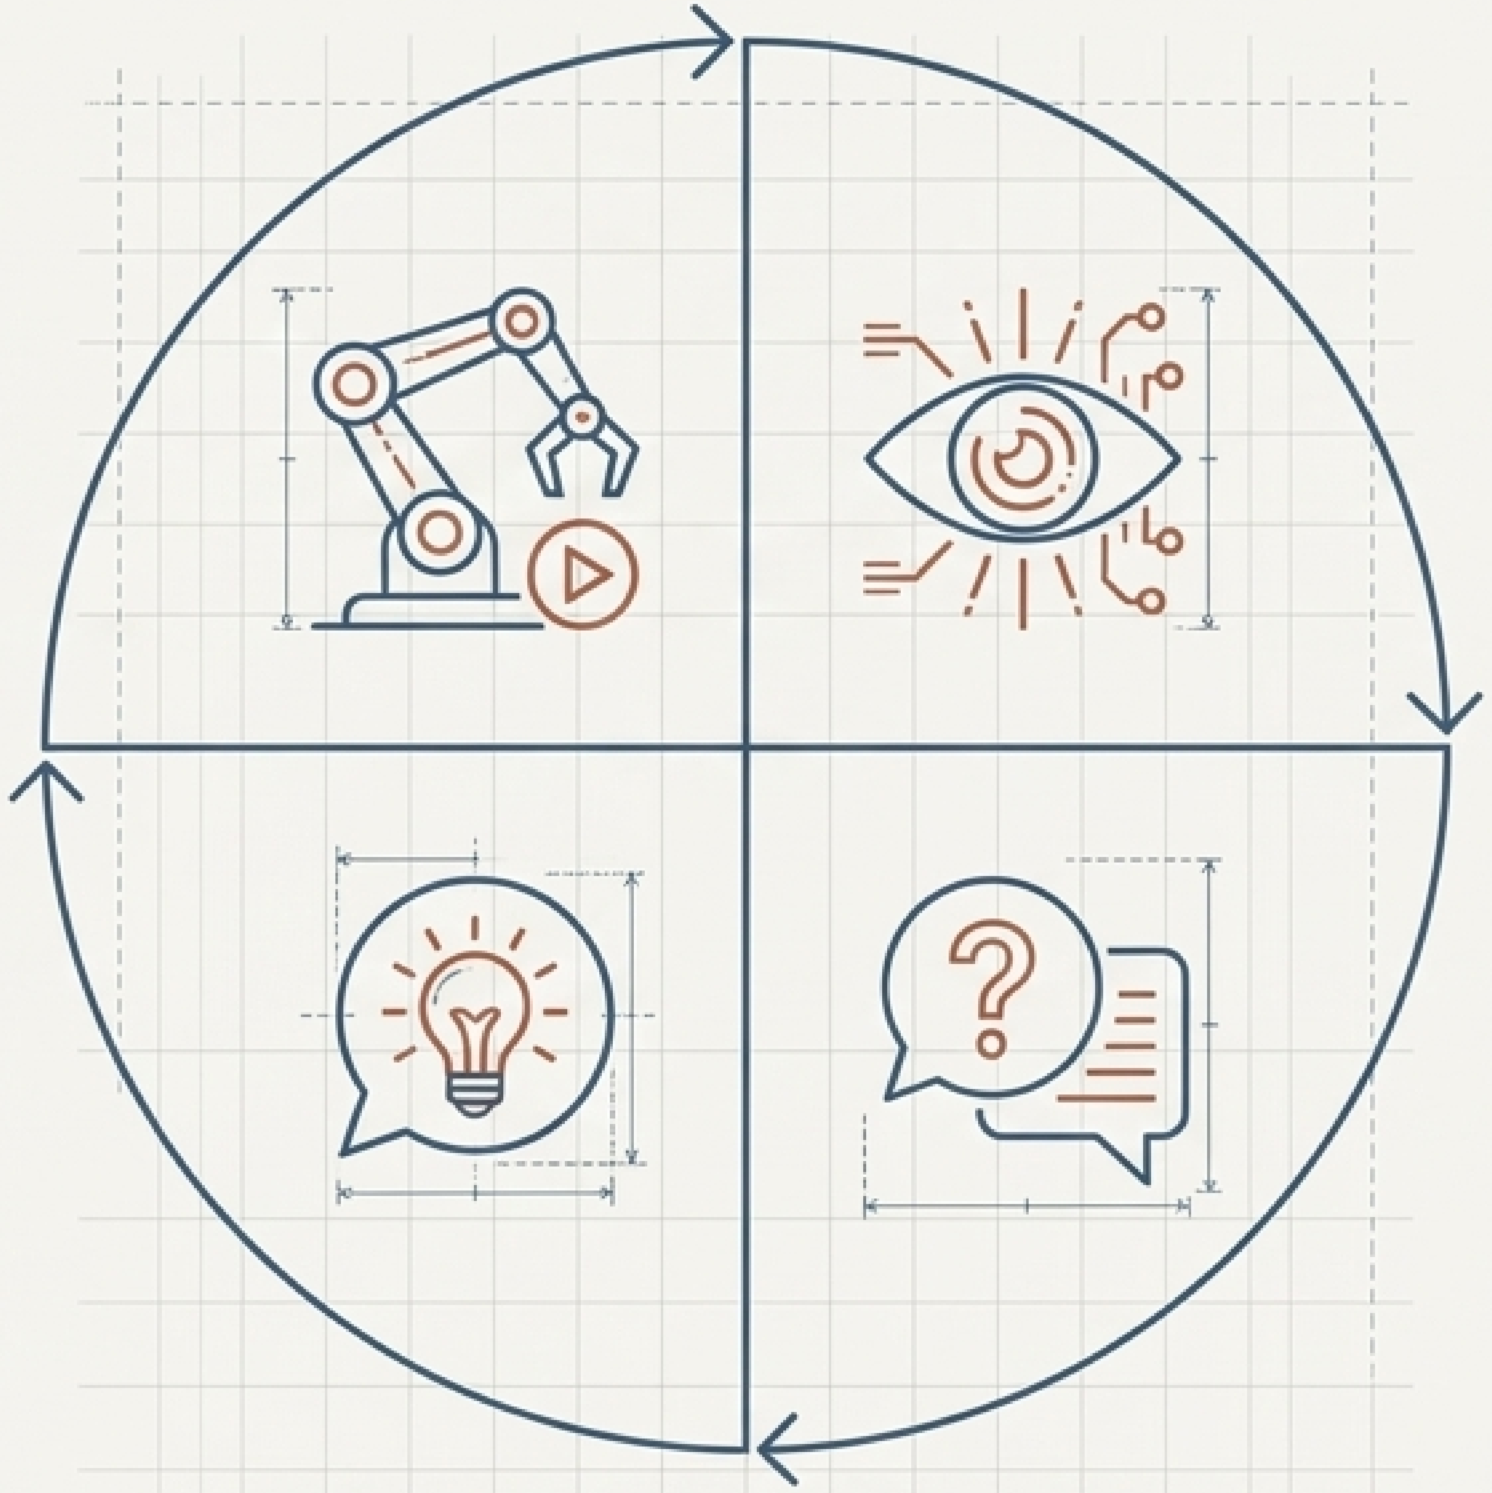

OBSERVATION

The agent collects information about its environment.

Examples: Sensor data from physical environments (camera feeds, vibration levels) or structured inputs from virtual environments (log data).

PROMPT

Based on observations, the agent generates a query for an external model (e.g., ChatGPT).

Purpose: To retrieve additional information, reasoning, or clarification beyond its internal knowledge.

RESPONSE

The external system provides relevant knowledge, insights, or suggestions.

Function: Augments the agent's understanding of the scenario with novel solutions or deeper context.

ACTION

Using its internal state and the new external knowledge, the agent selects and executes an action.

Outcome: Adjusting a strategy, refining internal models, or executing a physical task.

Translators Mediate Between an Agent's Formal Language and the LLM's Natural Language

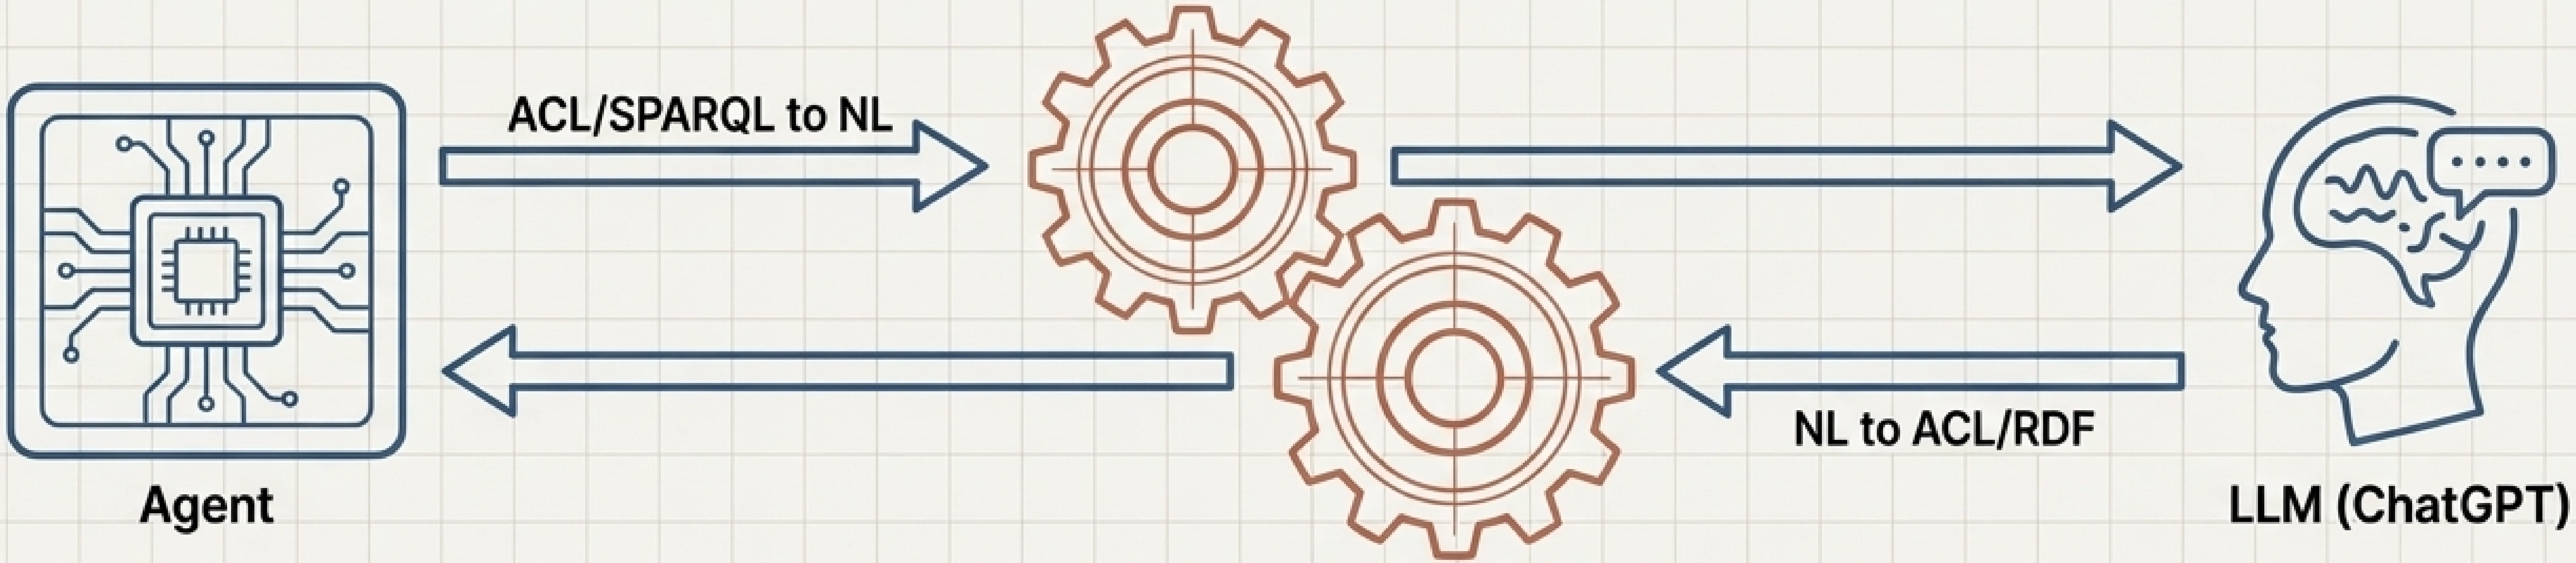

Agents communicate using formalized languages like Agent Communication Language (ACL) or SPARQL for querying knowledge graphs. LLMs operate on Natural Language (NL).

OPRA requires translators to bridge this gap:**

ACL to NL
Converts structured agent messages into NL queries.

SPARQL to NL
Transforms database queries into NL questions.

```
(request :content (temperature :location "factory_section_A"))
```
Translated NL Query
"Could you provide the current temperature in factory section A?"

This ensures agents can leverage NL-based reasoning while maintaining interoperability with standardized industrial communication protocols

Scenario: An OPRA Agent Monitors a Power Plant Turbine for Predictive Maintenance

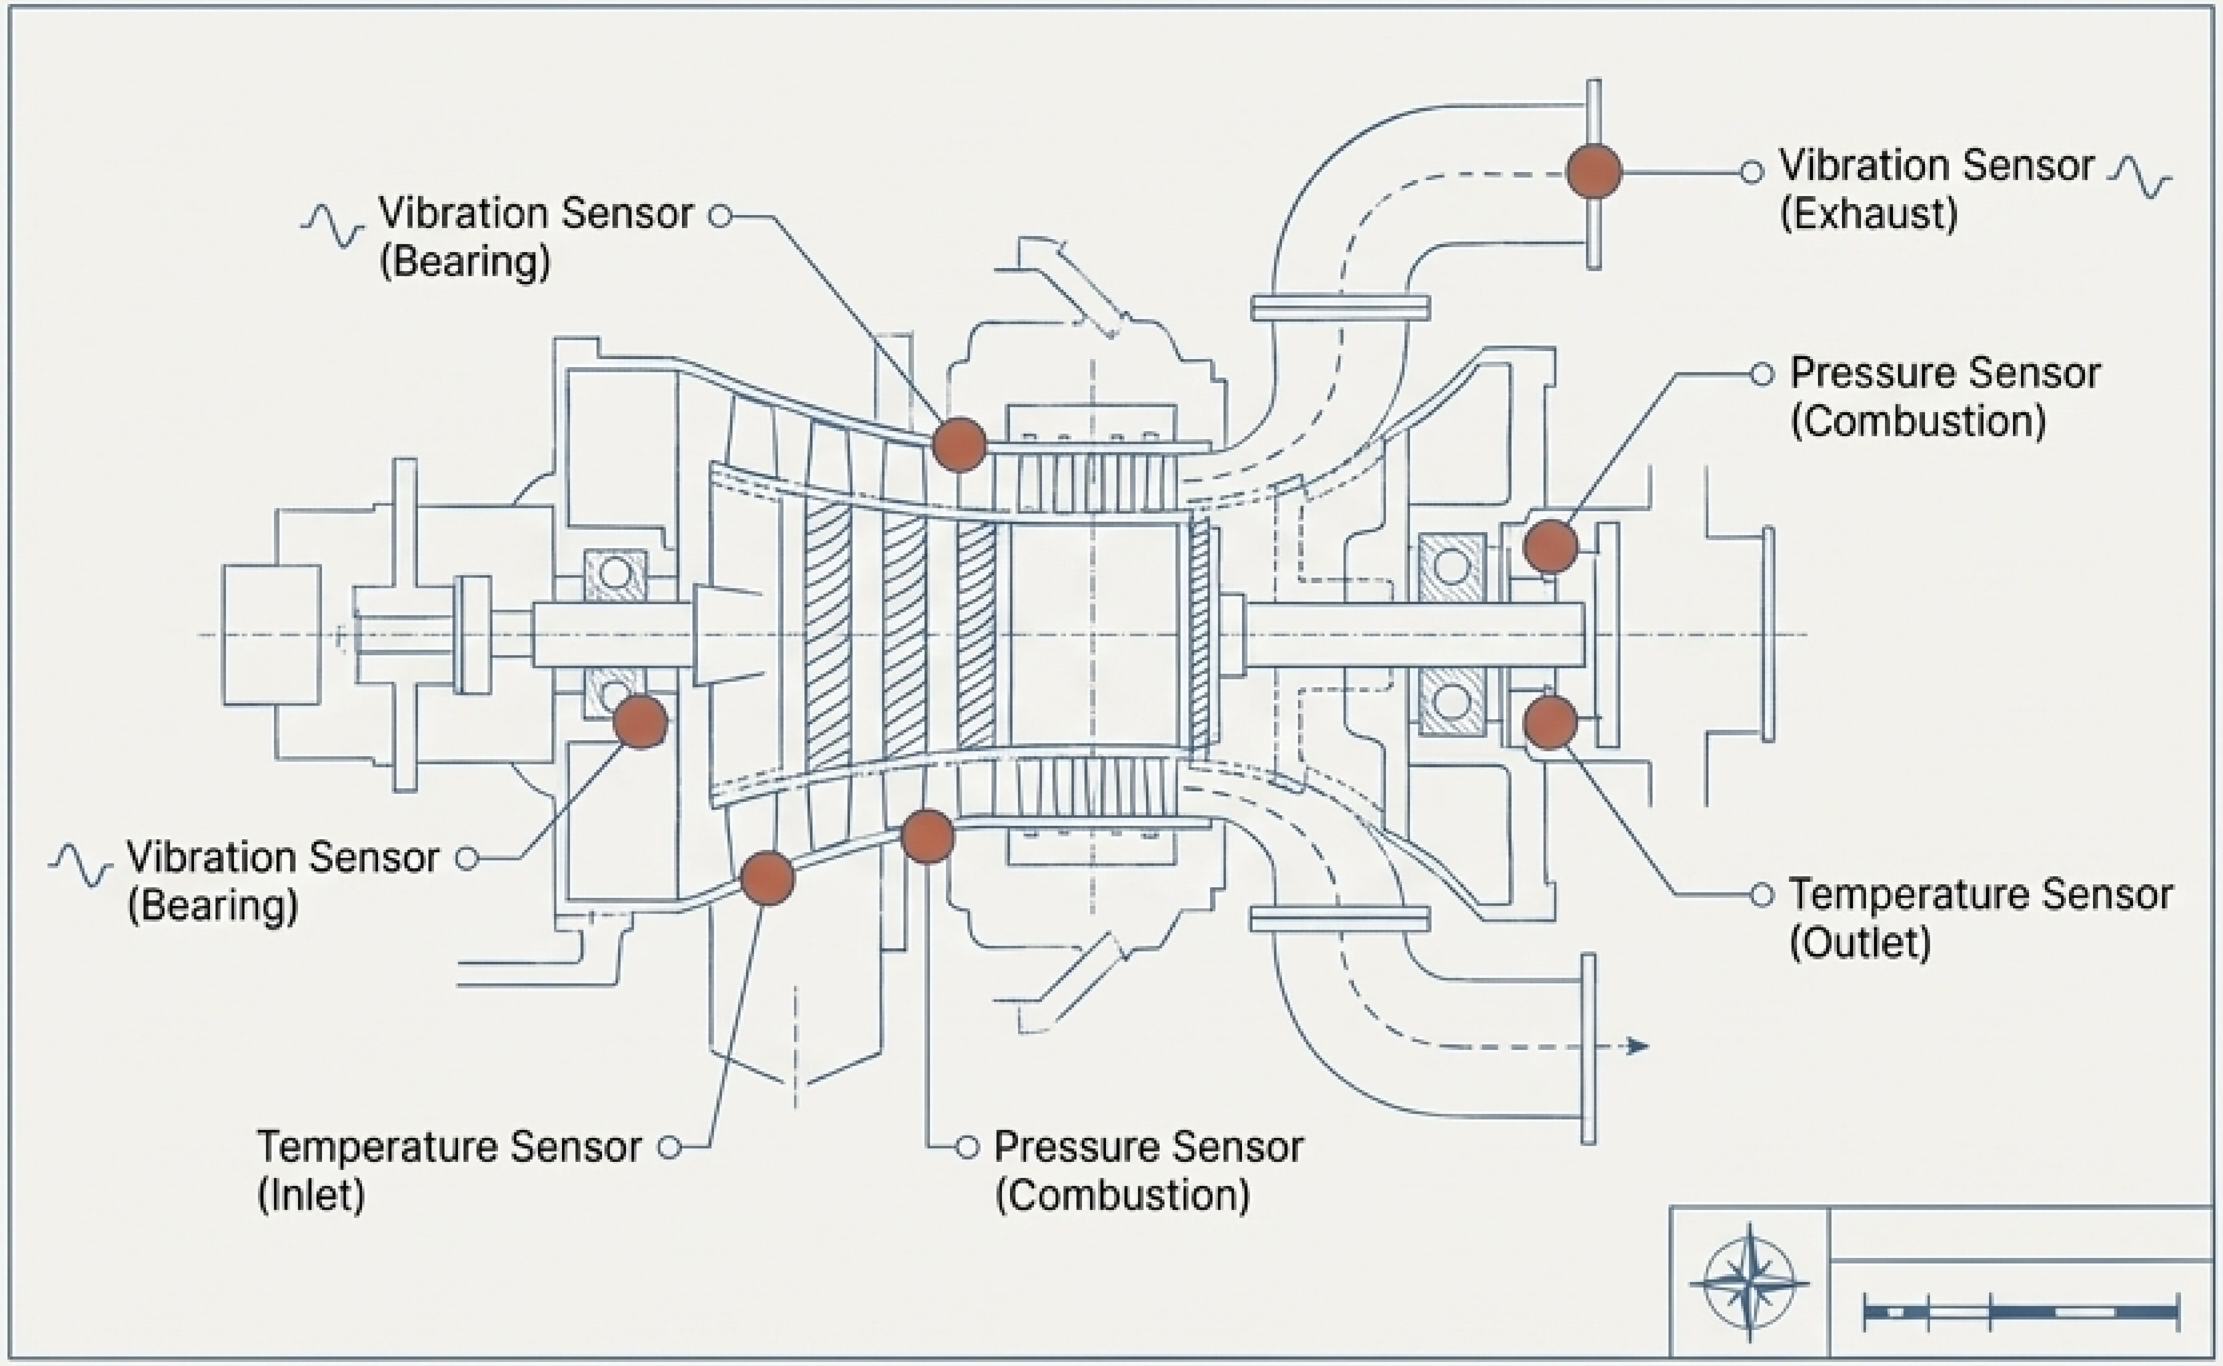

Asset: Industrial Turbine

Agent's Role: Monitor 'health' via real-time sensor data to avoid unexpected downtime.

Key Monitored Parameters & Normal Ranges

Vibration Level
Normal Range: 1–10 mm/s

1 mm/s | 10 mm/s
Normal | Warning | Critical

Temperature
Normal Range: 80–100 °C

Normal | Warning | Critical
0 °C | 80 °C | 100 °C

Pressure
Normal Range: 25–28 bar

Normal | Warning | Critical
0 bar | 25 bar | 28 bar

The Situation: The agent observes multiple sensor readings that are outside their normal range, indicating a potential issue that requires more than a simple rule-based response.

The Agent Observes Critical Anomalies and Formulates a Detailed Prompt for External Advice

Observation

Vibration:
15 mm/s ▲ HIGH

Temperature:
110 °C ▲ HIGH

Pressure:
30 bar ▲ HIGH

Prompt

>Consider me as an agent responsible for monitoring and maintaining the health of an industrial asset (a turbine in a power plant)... I have observed the following situation:

> • Vibration level of 15 mm/s (normal range: 1–10 mm/s),
> • Temperature of 110 °C (normal range: 80–100 °C),
> • Pressure of 30 bar (normal range: 25–28 bar).

Based on the provided details:
>1. Please provide your opinion on what could be causing these observations.
>2. Assess the urgency or emergency level of the observed situation.
>3. Recommend what my next action should be...
>4. Suggest what kind of external assistance or expertise is needed...

If you need further information... please ask follow-up questions...

An Iterative Dialogue Refines the Diagnosis by Incorporating New Contextual Data

CHATGPT'S INITIAL RESPONSE

**Assessment*: "High Urgency... could contribute to mechanical failure."

**Recommendation*: "Immediate controlled shutdown is recommended."

**Request for Info*: "I would ask... about the turbine's operational history, recent maintenance activities, or sensor calibration..."

AGENT'S RESPONSE (PROVIDING MORE DETAILS)

>The turbine was last maintained two months ago. No unusual sounds have been reported, but there has been a slight increase in load over the past week.

CHATGPT'S REFINED RECOMMENDATION

**Updated Analysis*: "The increased load over the last week likely contributed to the observed anomalies..."

**Revised Action: "I recommend scheduling an immediate maintenance inspection... It is advisable to reduce the load on the turbine but not shut it down entirely unless the vibration levels continue to rise."

The Agent Takes Calibrated Action to Mitigate Risk Without Unnecessary Downtime

Final Recommendation from ChatGPT

Reduce the load immediately while waiting for the maintenance team.

****Reasoning:** Reduces mechanical stress, thermal overload, and safety hazards. Continuing at full load could escalate the issue.

Agent's Final Actions

1. **Reduce the load** on the turbine to prevent further stress.
2. **Schedule maintenance** to check for misalignment or imbalance.
3. **Continue monitoring** vibration, temperature, and pressure levels.

Embedded Safety Thresholds (Informed by Dialogue)

The agent will automatically shut down the turbine if:

- **Vibration** exceeds **20 mm/s** OR
- **Temperature** exceeds **120 °C**

COPRA Extends the Framework for Collaborative, System-Wide Intelligence

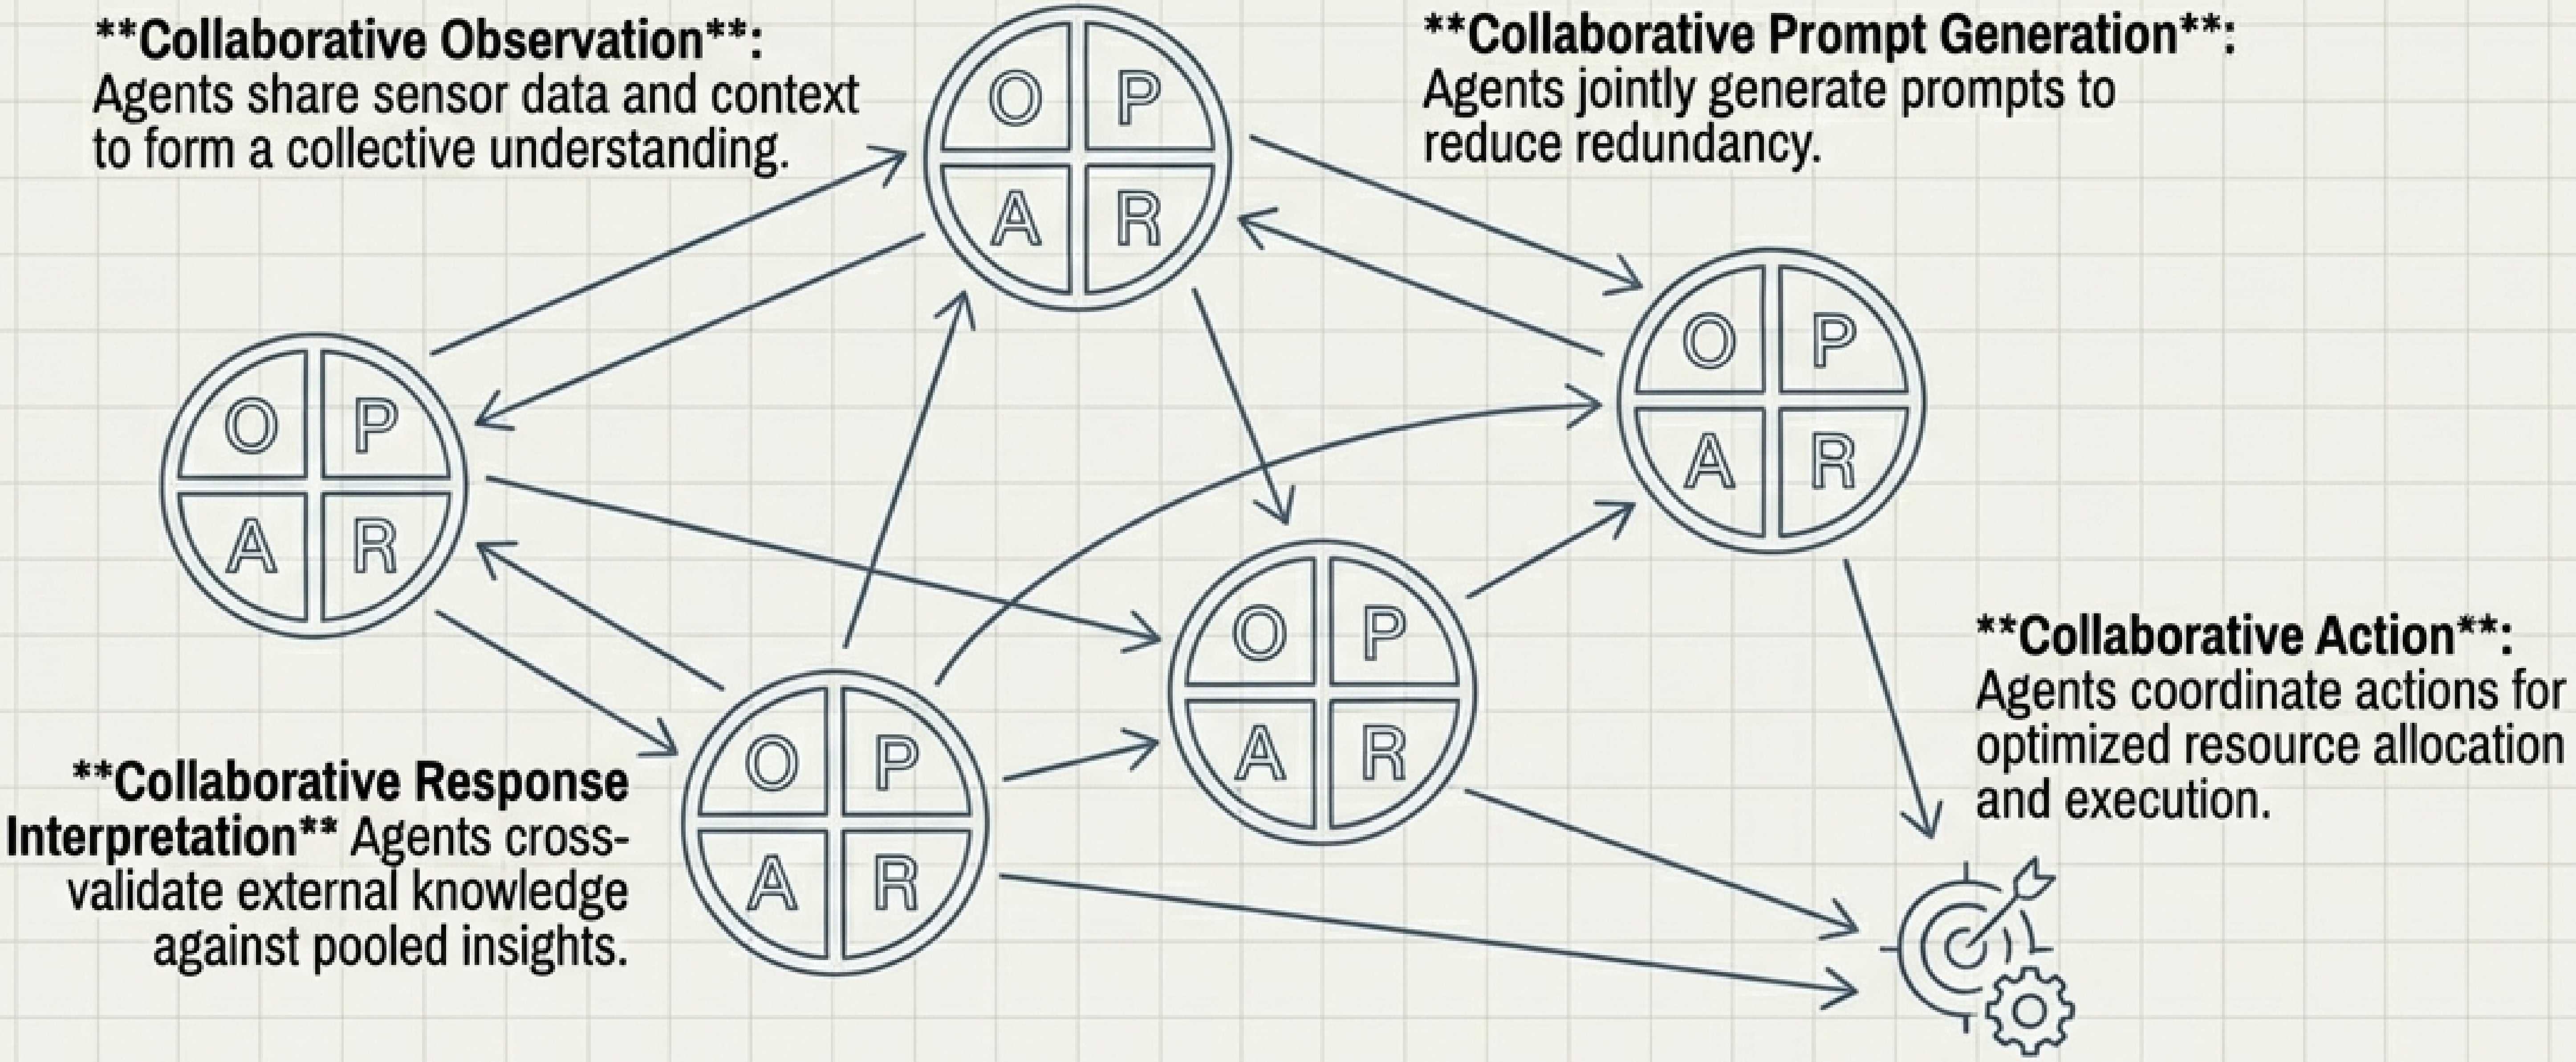

Key Concept: In COPRA, knowledge generation is a shared responsibility. This promotes distributed intelligence, reducing bottlenecks and ensuring decisions benefit from pooled insights across multiple assets or factory locations.

OPRA's Hybrid Approach Complements Traditional Architectures by Offloading Deliberation

Architecture	Adaptability to Novelty	Reliance on Pre-Programmed Knowledge	Handling of Uncertainty	Core Mechanism
OPRA	**High** Can query external knowledge for unforeseen events.	**Low** Relies on external, up-to-date knowledge source on-demand	**High** Iterative dialogue clarifies ambiguity.	External Deliberation via Prompt-Response Loop
Reactive	**Low** Bound by predefined stimulus-response rules.	**High** Actions are pre-programmed.	**Low** Cannot reason about novel situations.	Direct Observation-Action Mapping
Deliberative (BDI)	**Moderate** Can plan, but within limits of its internal belief model.	**High** Relies on internal beliefs, desires, and intentions.	**Moderate** Can struggle if internal knowledge is incomplete.	Internal Symbolic Reasoning and Planning
Learning (RL)	**High (over time)** Adapts through experience and iterative updates.	**Moderate** Starts with a policy, learns from interaction.	**Moderate** Requires sufficient experience to handle uncertainty.	Policy Optimization via Trial and Error
Knowledge-Based	**Low** Limited by the static, pre-defined knowledge base.	**Very High** Entirely dependent on its structured knowledge base.	**Low** Fails when rules for a situation are absent.	Inference Engine on a Static Knowledge Base

The OPRA Framework Is Designed for Scenarios Where Predefined Rules Are Absent or Insufficient

OPRA is most effective when...

Fast-Evolving Fields: Staying updated with best practices in cybersecurity or medical tech.

High-Risk Scenarios: Consequences of an incorrect decision are severe.

Uncertainty in Data: Sensor readings are incomplete or contradictory.

Lack of Clear Guidelines: Monitoring new equipment with no documented fault history.

Complex Diagnostics: Sensor data has multiple possible interpretations.

Novel Situations: An unexpected drop in production efficiency with no alarms.

Multi-Factor Decisions: Balancing performance, safety, and cost.

The Future is Explainable: The Roadmap Extends to OPRA+ and COPRA+ for Enhanced Transparency

OPRA+

Introduces an 'Explanation' step to provide context-aware justifications for chosen actions.

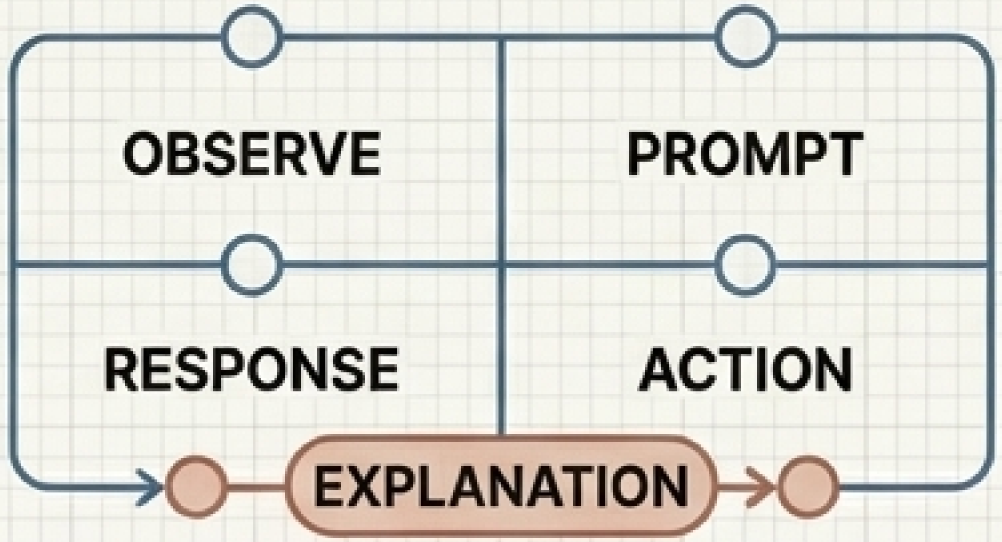

OPRA+: Introduces an 'Explanation' step to provide context-aware justifications for chosen actions, fostering user trust and interpretability.

COPRA+

Integrates the explanation phase withir multi-agent, allowing agents to rationale behind actions for improved coordination.

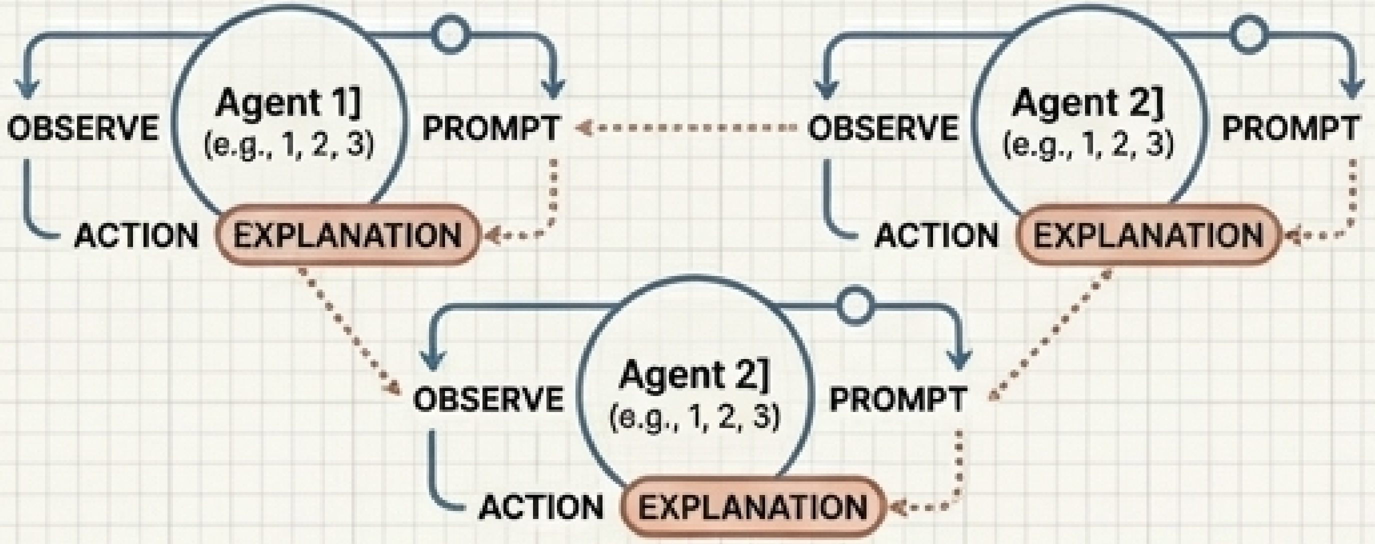

COPRA+: Integrates the explanation phase within the multi-agent context, allowing agents to share the rationale behind actions for improved coordination.

Core Value: By leveraging LLMs to provide explanations alongside recommendations, OPRA+ and COPRA+ strengthen the system's ability to make context-aware, goal-aligned, and trustworthy decisions.

OPRA Re-imagines Agent Autonomy by Integrating Dynamic, Collaborative, and Explainable Knowledge

A New Paradigm for Adaptability

OPRA and COPRA shift from rigid, internal models to dynamic, external knowledge querying. This enables agents to handle novelty and uncertainty in complex ISM environments where predefined rules fail.

Proven Value in High-Stakes Applications

The predictive maintenance scenario demonstrates a tangible reduction in risk and operational downtime through calibrated, context-aware actions.

A Foundation for Trustworthy AI

The roadmap toward OPRA+ and COPRA+ directly addresses the critical need for transparency and explainability, paving the way for more resilient and collaborative human-AI systems in manufacturing.