

# Human decision-making is our greatest asset—and our biggest bottleneck.

The resilience and sustainability of any organization depend on its people. Yet, high dependency on specific employees, especially in creative, strategic, and emergent decision-making, creates critical points of failure. This creates a fundamental paradox: the very human-centricity we value is also a source of subjectivity, overload, and fragility.



**How do we preserve the unique value of human decision-makers while eliminating them as a single point of failure?**

# Existing approaches force an unacceptable compromise.

Traditional solutions to workplace overload offer a compromise between two extremes, each with significant drawbacks.



## Entirely Human (HR Management)

### Advantage:

- The “gold standard”; humans can explain their reasoning and take responsibility.

### Disadvantage:

- Doesn't scale; creates bottlenecks, high dependency, and subjectivity. Fails in peer-to-peer network workflows.



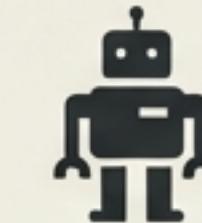
## Human + Support (Decision Support Systems - DSS)

### Advantage:

- Can distribute routine tasks and increase resilience.

### Disadvantage:

- Can *\*increase\** workload through fragmentation, excessive alerts, and interoperability issues. Does not execute decisions.



## Entirely Automated (ADMS)

### Advantage:

- Creates ubiquity (“anywhere, anytime”) and eliminates human bottlenecks.

### Disadvantage:

- Faces challenges in expert knowledge extraction, legal implications, and can still increase stress. Lacks the nuance of human judgment.

A new logic is needed that enables the smart integration of human and machine **intelligence without compromise**.

# Introducing the Pi-Mind Agent: A Responsible, Resilient, Ubiquitous Cognitive Clone.

**Core Concept:** An information technology artefact that acts as an autonomous digital copy of a human decision-maker (the 'donor').

## How it Works

- It learns the decision-making capabilities of its donor, including personal preferences, features, and biases.
- It can act as a cognitive enhancement tool, an autonomous representative, or a replacement when the donor is unavailable.

## The Core Assumption

A decision-maker will be more confident if their autonomous digital representative is as capable and predictable as their exact personal copy, especially when they are sick, overloaded, or need to manage multiple processes simultaneously.



# The Four Design Principles: The Foundational DNA of a Pi-Mind Agent

Every Pi-Mind agent is built on four core principles that ensure it is a true and trustworthy cognitive copy of its donor.



## DPO: The Turing Principle (The Core Code)

*Be an exact copy.* The clone's cognitive behavior must have minimal deviation from its donor's in identical situations.

Precision of imitation (minimal difference = maximal precision)



## DP1: Responsibility (The Ethical Code)

*Inherit the owner's responsibility.* The clone acts as a representative, so the human donor remains morally and legally responsible for its decisions.

Percentage of decision-making duties delegated.



## DP2: Resiliency (The Stress Code)

*Perform under pressure.* The clone must replicate the donor's ability to handle new, challenging, confusing, or stressful situations.

Precision of imitation in an emergency context.



## DP3: Ubiquity (The Replication Code)

*Be everywhere at once.* The clone enables a single person's digital involvement in many processes simultaneously.

Persons/hours saved through parallelization.

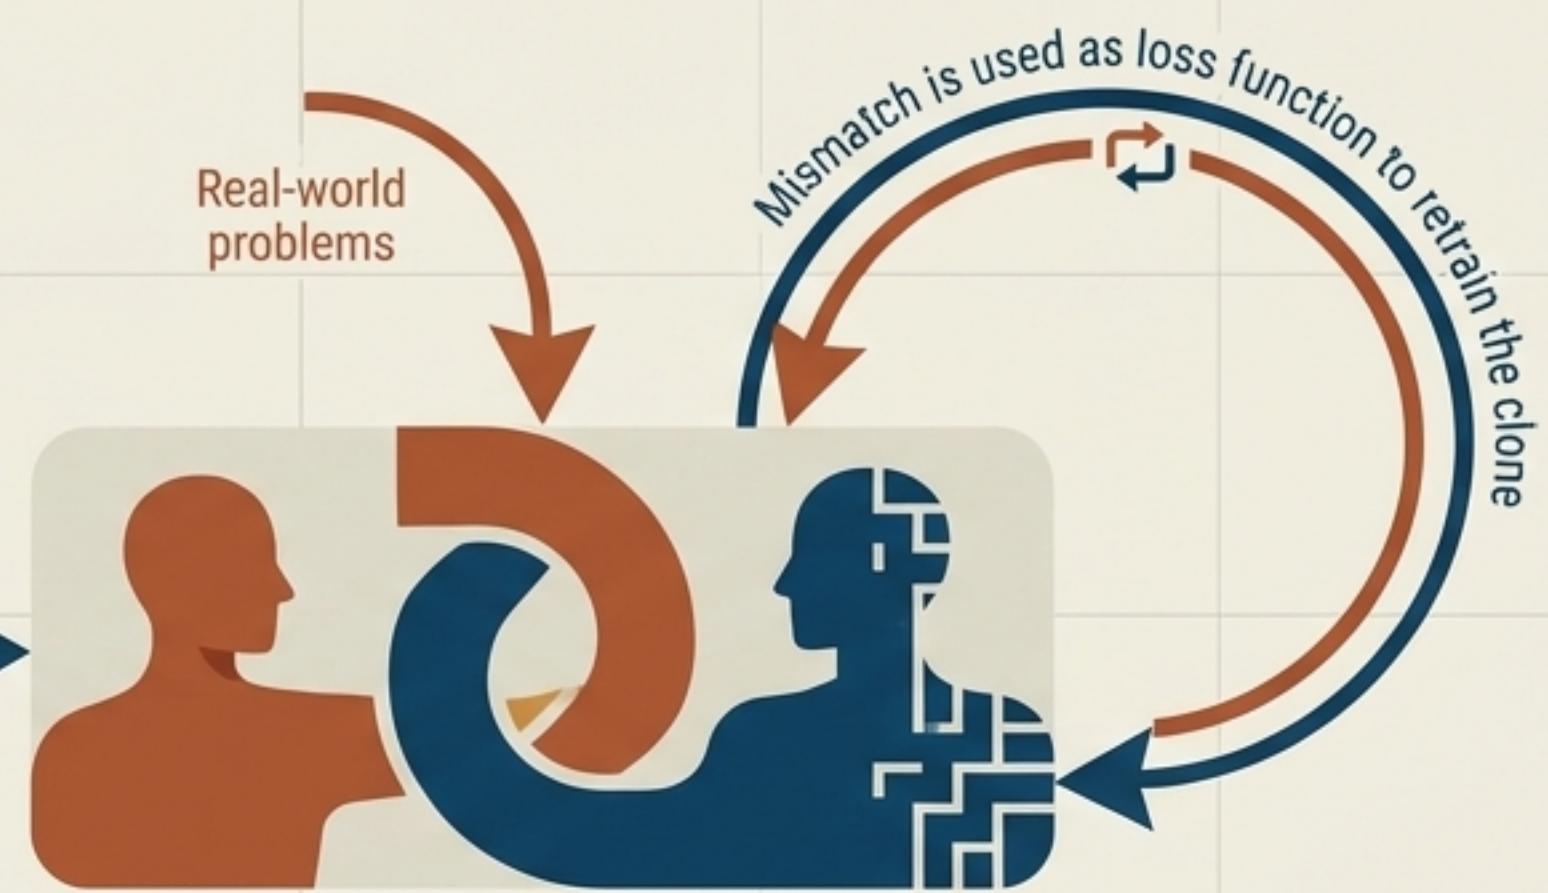
# How a Clone Learns: Training Through Adversarial Collaboration

The clone's cognitive capabilities are developed through adversarial learning in an environment we call **T|C-SGAN**. This isn't just about feeding it data; it's about teaching it to think under pressure.



**The Generator**  
("The Sparring Partner")

An AI component (G) whose job is to create challenging, novel, or "fake" decision scenarios to try and confuse the team. It constantly improves its ability to create difficult problems.



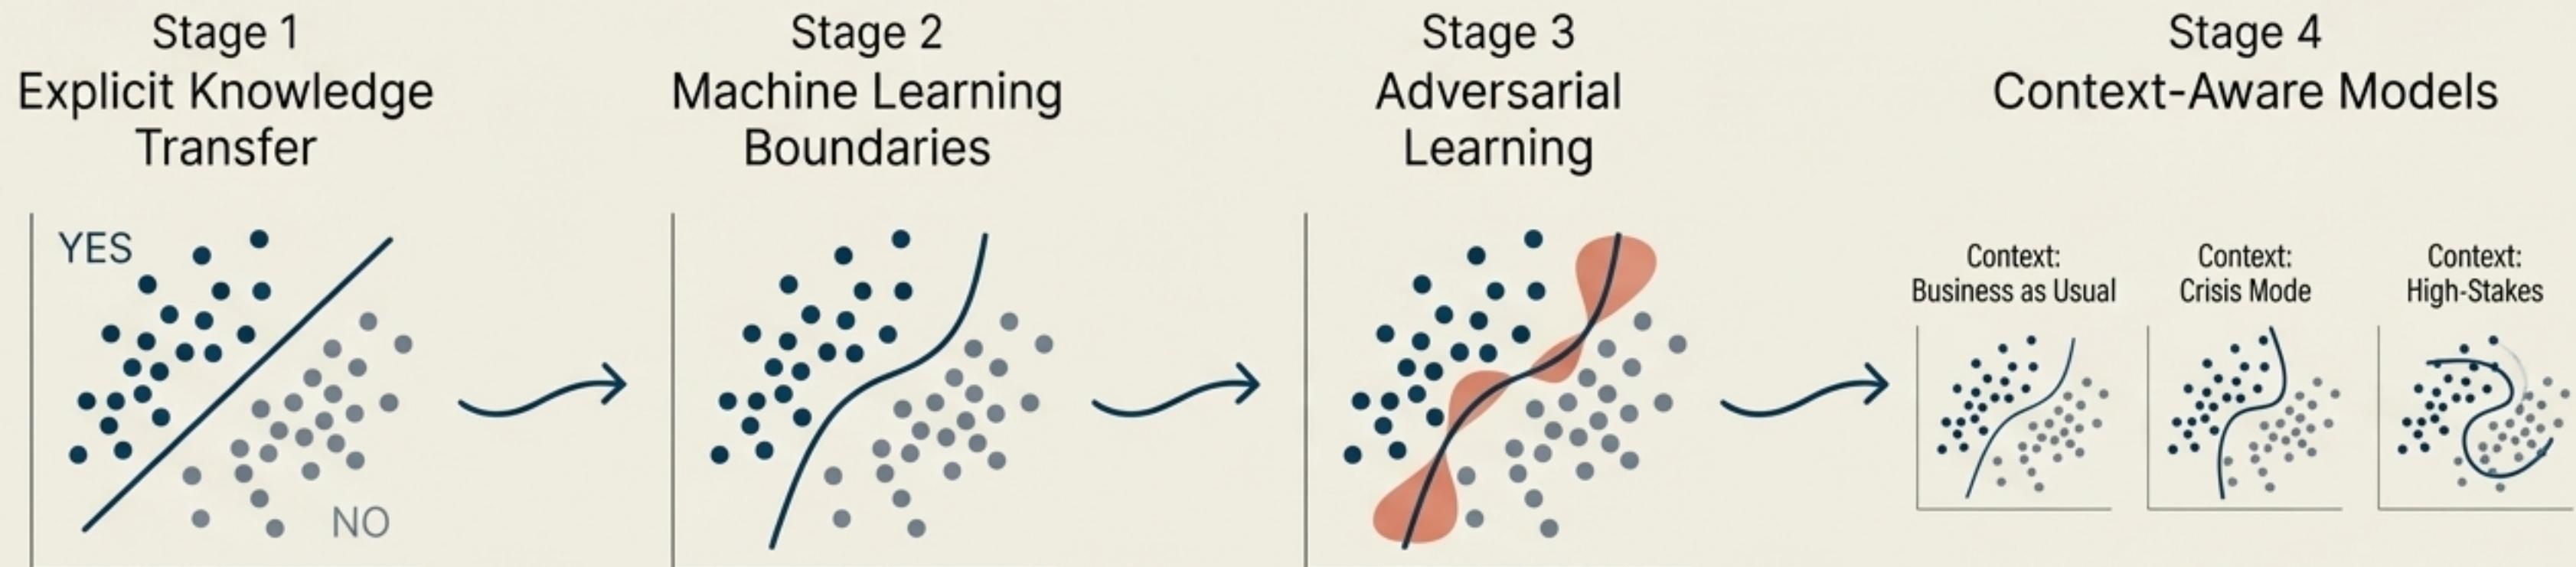
**The Turing Discriminator**  
("The Mentor-Apprentice Team")

A "Human Donor" (H) and their digital "Clone" (C) work together. They are presented with both real-world problems and the "fake" problems from the Generator. The Clone learns by trying to match the Human Donor's decisions as precisely as possible.

## The Outcome

This "game" drives the coevolution of the team and the challenger. The Clone learns not just the rules, but the **"nuance"** of the Donor's decision-making, even in situations it has never seen before, ensuring resiliency (DP2).

# Capturing More Than Rules: From Explicit Knowledge to Hidden Intuition



Starts with explicit rules provided by the donor.

The clone learns hidden decision boundaries from observing the donor's choices.

Adversarial training discovers these "corner cases" to make the clone's boundaries more precise.

Finally, the clone learns that decision logic changes with context, creating a rich, personal decision ontology.

The Pi-Mind agent integrates all these layers to create a truly representative clone capable of nuanced, context-dependent decisions.

# The Evidence: A Three-Part Validation of the Pi-Mind Agent

We evaluated the Pi-Mind agent using a multi-stage Design Science Research Methodology across real-world environments, including a Ukrainian higher education institution (NURE) and a military logistics laboratory (NATO).



## Proof of Concept

Does the artefact work as intended? We measure its effectiveness and decision accuracy.



## Proof of Value

Is the artefact useful? We measure its efficiency and the tangible benefits it provides.

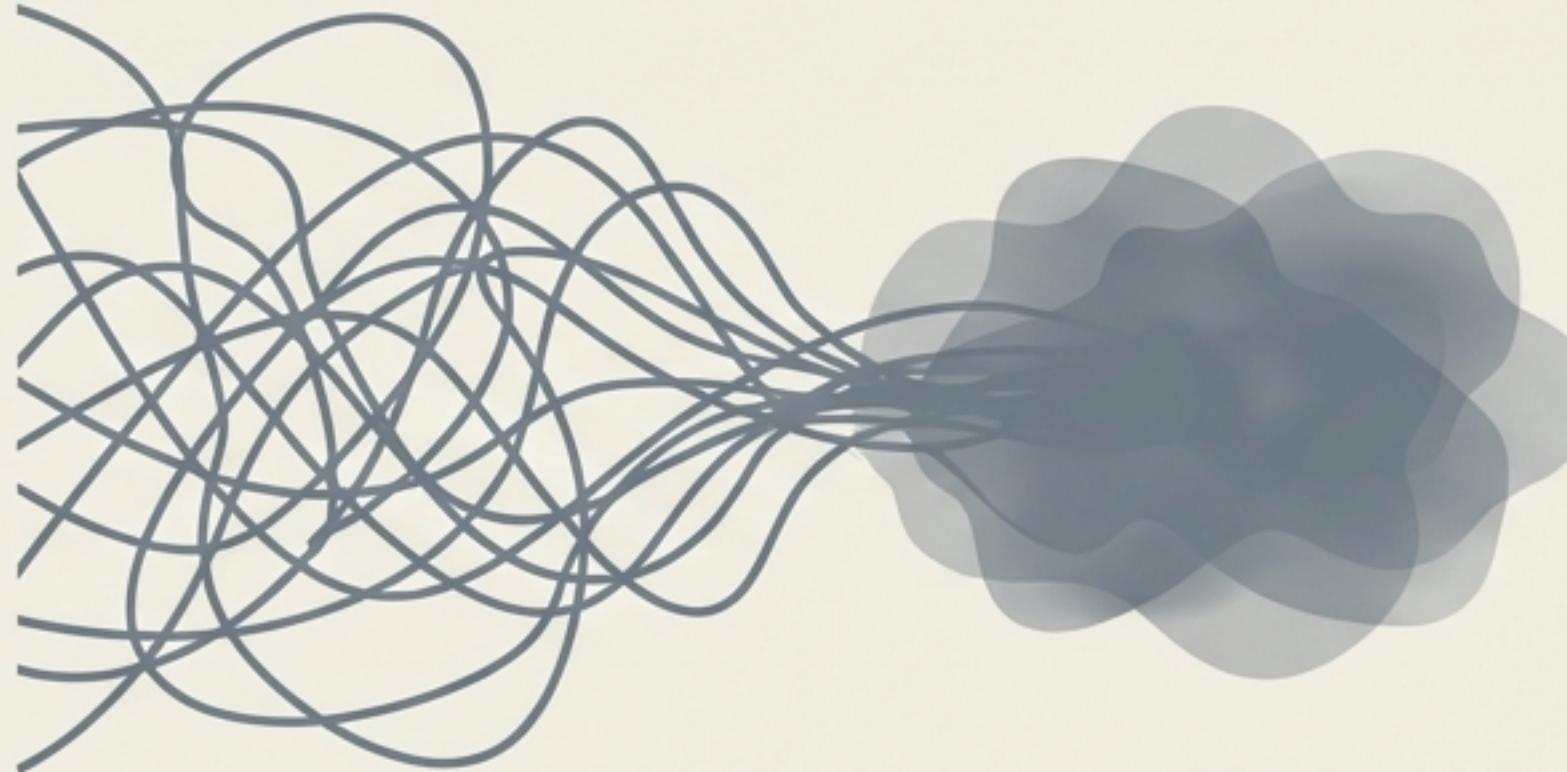


## Proof of Use

Does the artefact have a broader impact? We measure its influence on organizational processes and society.

# Proof of Concept: Increasing Decision Accuracy from 49% to 96%.

NURE Employee Motivation - The annual process of distributing monetary bonuses to academic staff.

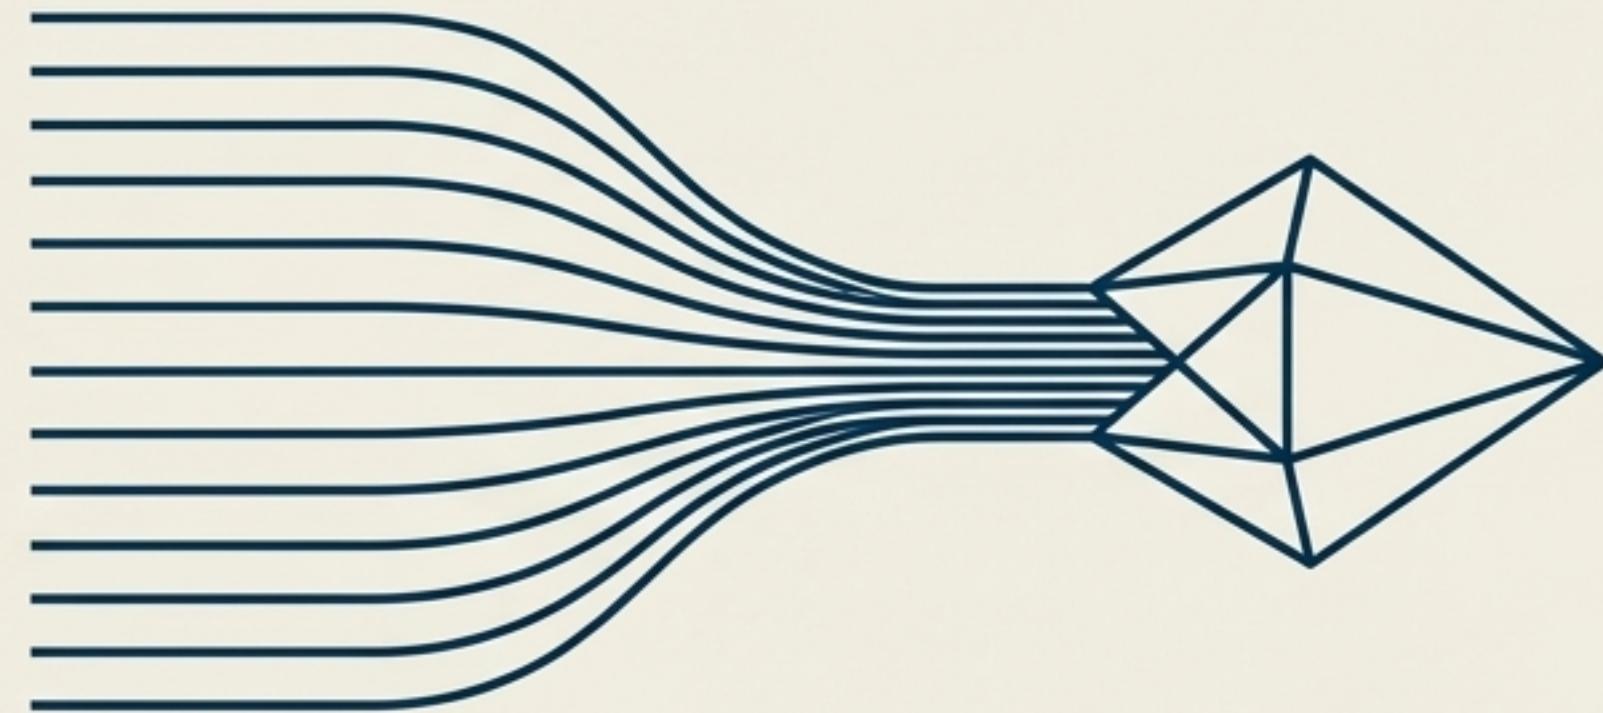


**Process:** Decisions made by an Award Committee.

**Problem:** Highly subjective, prone to bias and manipulation, leading to poor alignment with actual performance.

# F1-Score: 0.49

Barely better than a coin flip.



**Process:** The Rector's Pi-Mind agent generated rankings based on actual, verified achievements.

**Outcome:** Transparent, objective, and defensible decisions. Accidental and deliberate errors were excluded.

# F1-Score: 0.96

A dramatic increase in accuracy.

# Proof of Value: Saving Over 1,500 Human Hours Annually

NURE Recruitment - Selecting job candidates by the academic council.

NURE Recruitment - Selecting job candidates by the academic council.

## The Old Process

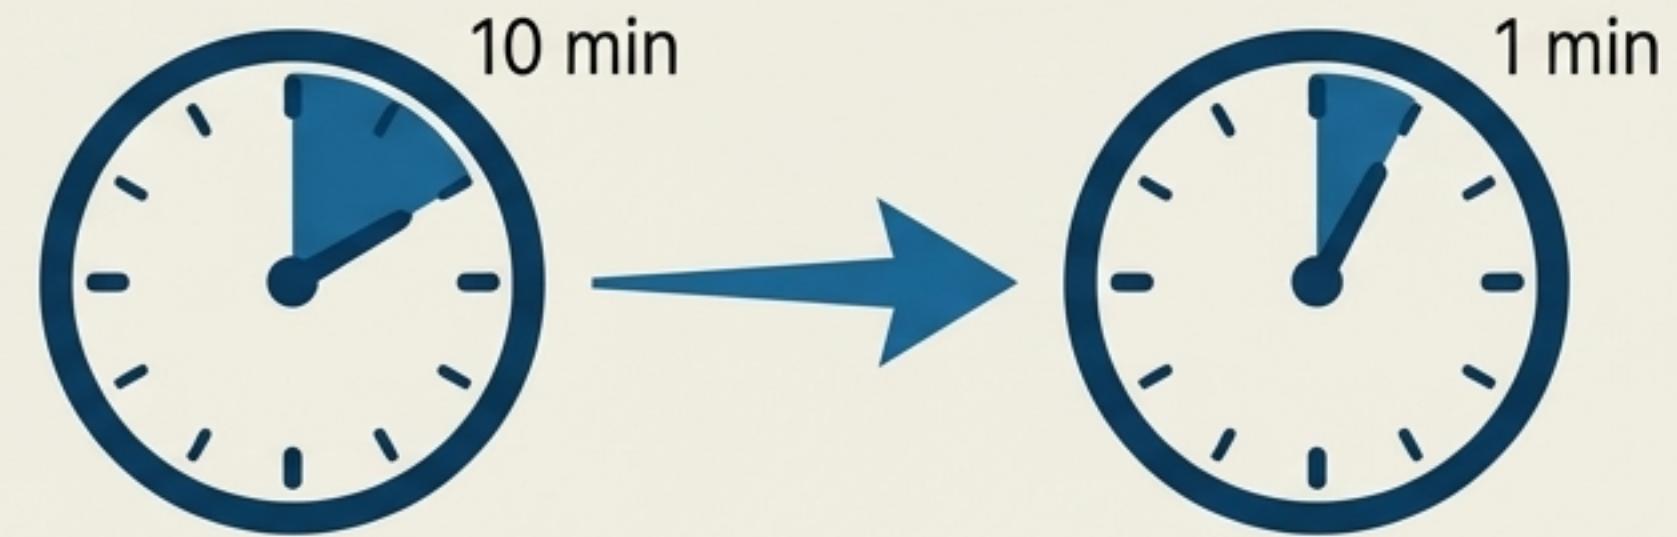
- Required lengthy in-person meetings for the hiring committee (23-48 members).
- Each candidate review took at least 10 minutes of collective time for document review and voting.

## The New Process

- Pi-Mind agents for each council member generated proactive assessments based on candidate achievements.
- This automated the “vote,” replacing the meeting with an aggregation of ranking lists.
- The approval time for each candidate was reduced to just **one minute**.

## The Bottom Line

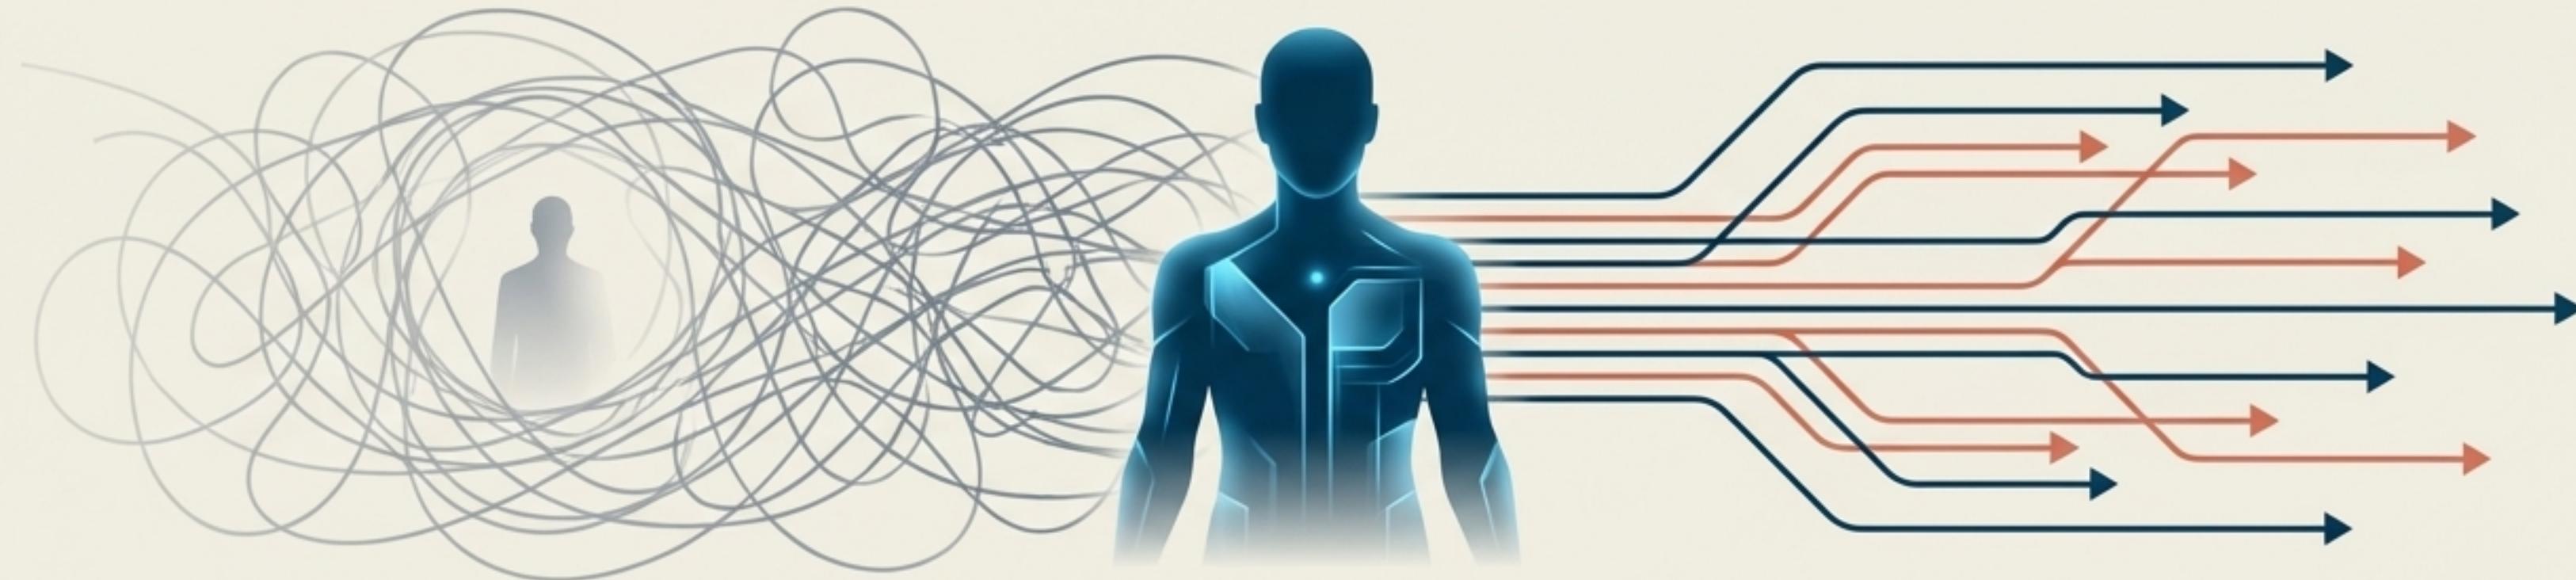
- **9 minutes saved per transaction.**
- This resulted in a total saving of **1,591 human hours per year** for the university.



**>1,500  
Hours Saved**

# Proof of Resilience: Decisive Action During the COVID-19 Crisis.

Case study: "NURE Extreme Resource Reconfiguration."



## The Challenge

The sudden announcement of the COVID-19 quarantine required an urgent transition to crisis management and full remote operations. The primary decision-maker, the university rector, was unavailable due to a business trip.

## The Solution

The administration used the rector's existing Pi-Mind agent. Though created for other business processes, it accurately reflected his administrative preferences for resource allocation.

## The Result

**3 working days**

to completely restructure and adapt all university processes to remote mode.

**F1-Score of 0.95**

decision accuracy, comparing clone-suggested orders to those ultimately effective.

**Key Insight:** The Pi-Mind agent ensures cognitive continuity and protects critical workflows even when key personnel are unavailable, demonstrating true organizational resilience.

# Proof of Use: Accelerating National Higher Education Reform

The Pi-Mind agent was implemented as part of the TRUST Portal, a digital infrastructure designed to increase transparency and decrease corruption in Ukrainian Higher Education (HE).

During the 2015 election for Ukraine's first national independent QA agency, the Pi-Mind agent was used to evaluate candidates based on merit. The agent's transparent rankings revealed that officials were attempting to nominate less-qualified, politically chosen candidates. Under from the solid, transparent facts provided by the agent, the Ministry of Education and Science was forced to annul the nominations and select the more qualified candidates.



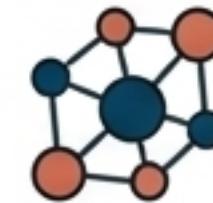
**>5,000**  
registered users



**>500**  
Pi-Mind agents created



**>1,700**  
procedures launched  
by agents

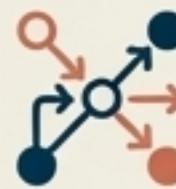


**23 Million**  
knowledge triples in the  
ontological database

This demonstrates "last-mile DSR," where the research artefact has a direct, measurable impact on society and governance.

# Verified Performance: A Quantitative Summary

## Turing Principle (DPO) - Effectiveness



F1-Score (Decision Accuracy)

Increased from **0.49** to **0.96** in the NURE business-as-usual case.

Reached **0.995** correlation in the NATO security case.

## Responsibility (DP1) - Effectiveness



% of Responsibilities Delegated

**100%** in NURE case due to stored personal value systems (PSVs).

**>85%** in NATO business-as-usual processes.

## Resiliency (DP2) - Effectiveness



F1-Score (Crisis Decision Accuracy)

**0.95** in the NURE crisis management case.

Increased to **0.7** even under adversarial evasion attacks in the NATO case.



## Ubiquity (DP3) - Efficiency

Human Hours Saved / Processes Run

**>1,500 hours saved** annually in one NURE procedure.

**>1,700 processes** run in the real-world TRUST environment.

# The Vision: A Future of Human-AI Symbiosis

The Pi-Mind agent is more than a tool; it's a step towards a new paradigm of collaborative intelligence that keeps humans in the loop.



## Implications for...

### Responsible AI



Provides a framework for transparently distributing responsibility between a donor, developer, and user, balancing AI autonomy with human control.

## Implications for...

### The Future of Work (Industry 4.0)



Preserves the human-centric nature of cybernized manufacturing and enables highly personalized product creation by cloning the expertise of workers, operators, and even customers.

## Implications for...

### Ethical Autonomy



Offers a path to resolving ethical dilemmas in areas like self-driving vehicles. If a digital driver makes the same choices a thoughtful human owner would, trust is significantly enhanced.

Society doesn't need faceless analytics; it needs AI with a 'human face'—a sustainable and trusted partner to help overcome complex challenges.

René Descartes, 1637: "Cogito, ergo sum." (I think, therefore I am.)

My clone thinks  
when I cannot,  
therefore, I **(still)** am.

The Pi-Mind agent ensures the sustainable cognitive involvement of human experts,  
preserving their essential role in an increasingly complex and fragile world.