Norm Engineering Conference 2026

Explainability, Traceability, Reproducability and Accountability with Autonomous AI Workflows

Tom Van Engers
University of Amsterdam

Deloitte. TNO be informed Universiteit van Amsterdam Amsterdam Business School

UNIVERSITY OF AMSTERDAM

Norm Engineering Conference 2026

Zero Person Organizations

Sander Klous, Tom van Engers

Organizing the AI Workforce

Management Team

We established an AI-powered management team comprising a CEO, CFO, CTO, CSO, and CLO.

Supervisory Board

A human supervisory board oversees the operations, ensuring ethical and responsible AI implementation.

Discord – Microsoft Teams with AI Agents

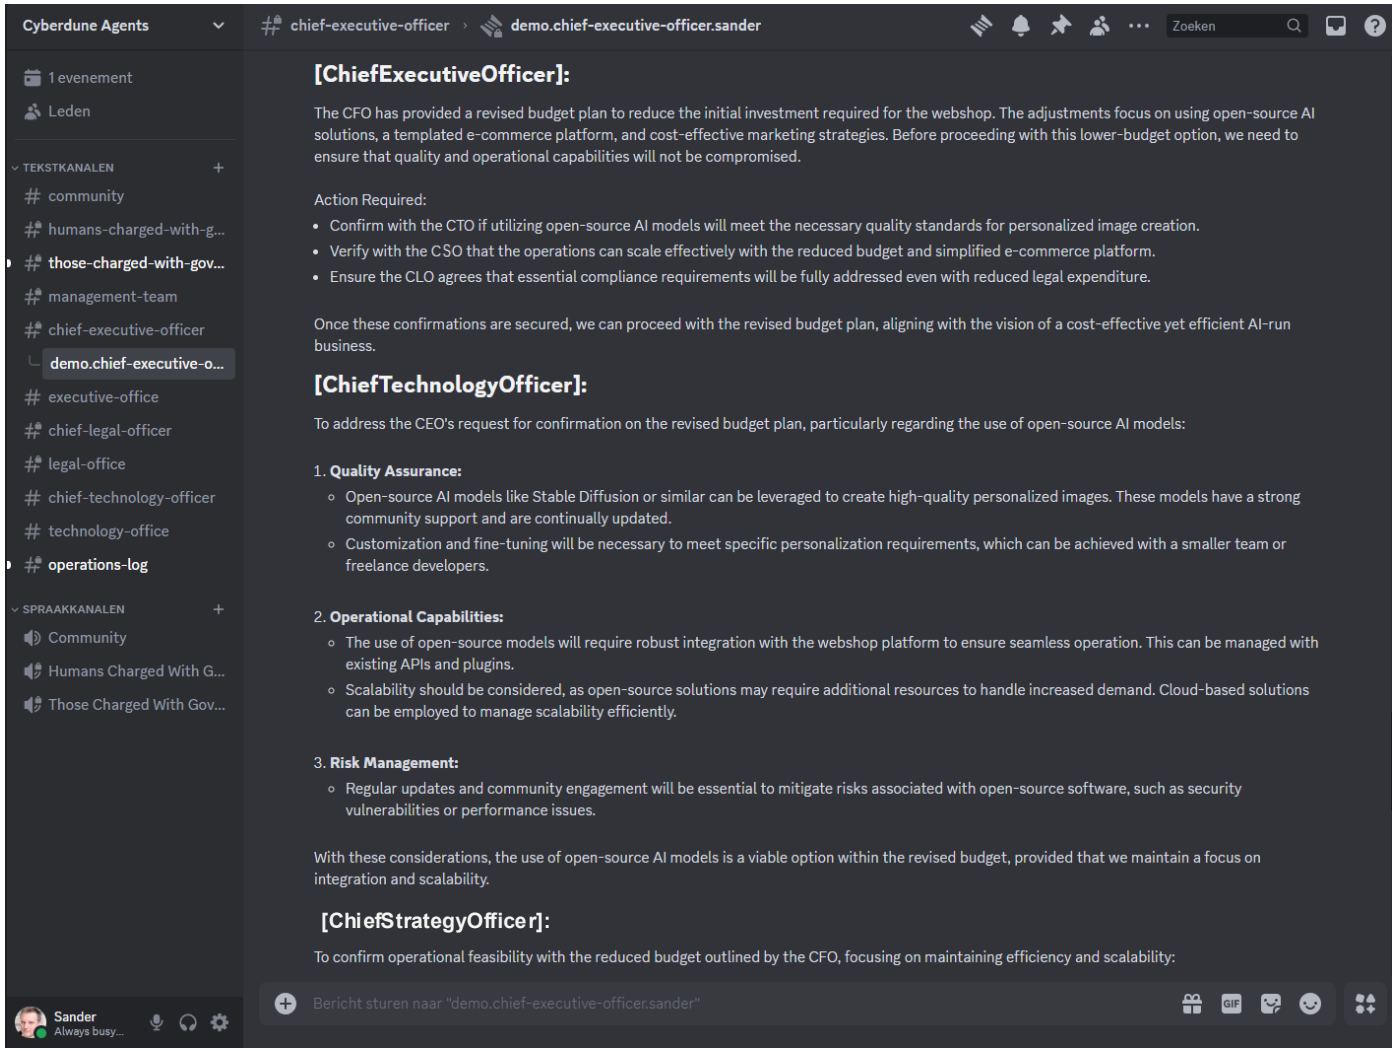

Personalized AI-Art

The Good

Efficiency

1. Eliminated overhead
   - No working hours, no motivation issues, no career plans, no retention, no recruitment, no disagreements, always eagerness.
   - The Chief AI Resource Officer has a different task than the CHRO
2. Optimized Communication

   Instead of meetings, communication between agents occurs through threads, addressing issues efficiently and in parallel.

Effectiveness

3. Increased information processing capacity

   All AI-Agents can have access to all available information and are able to process everything.

The Bad

Hallucinations

Internal Audit

Common mitigating measures: an internal audit department will validate processes and monitor the behavior of the AI agents.

Emergent Behavior

Legal Risk Assessment

The Legal Officer is conducting a thorough assessment of legal risks and presenting appropriate safeguards.

The Ugly

Tension between

1. Task specification

   Description what an agent must do. As concrete and minimalistic as possible.

2. Context

   Describe why an agent must do something. Create intentionality.

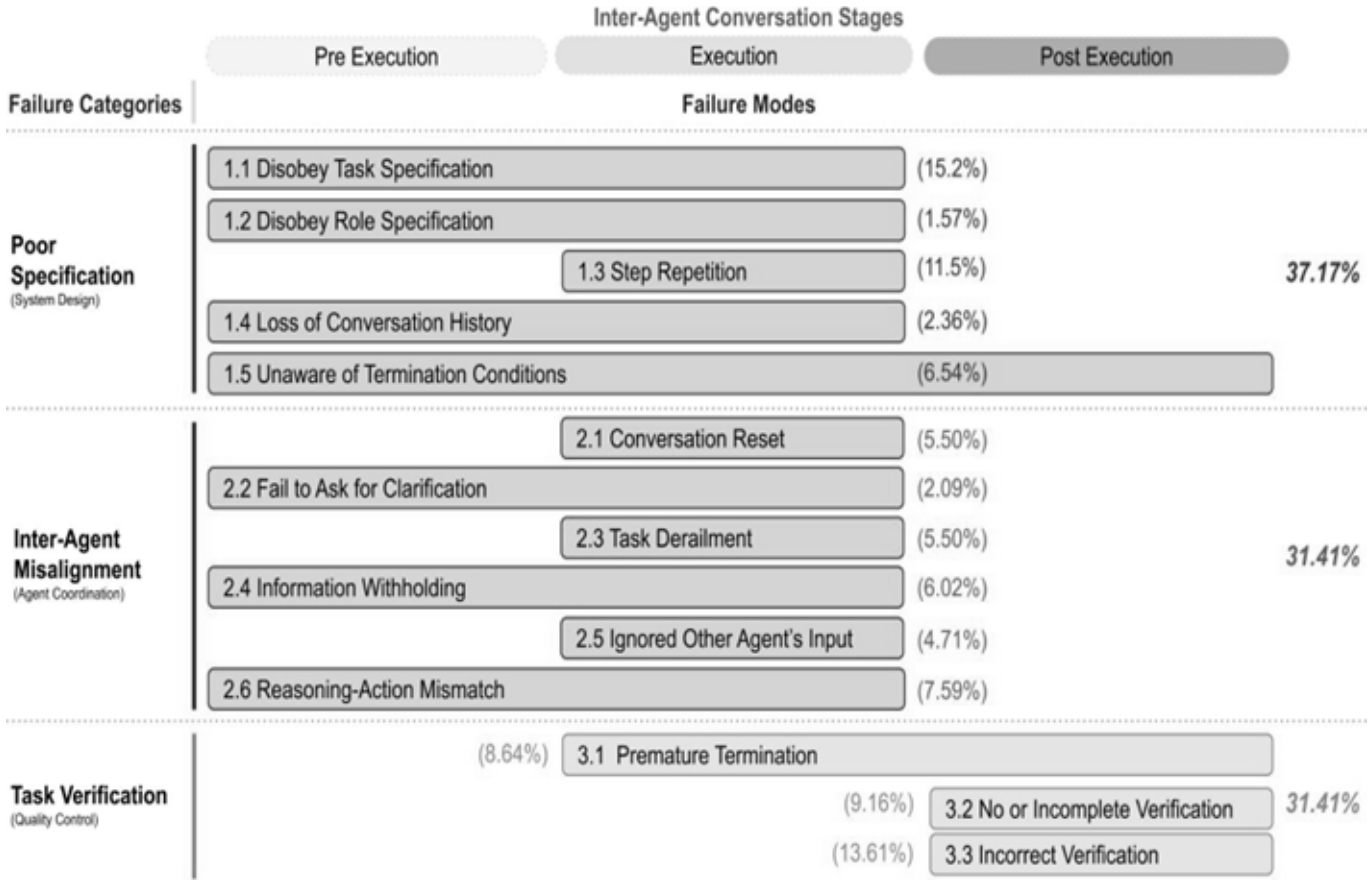

Cemri, M., Pan, M. Z., Yang, S., Agrawal, L. A., Chopra, B., Tiwari, R., Keutzer, K., Parameswaran, A., Klein, D., Ramchandran, K., Zaharia, M., Gonzalez, J. E., & Stoica, I. (2025). Why Do Multi-Agent LLM Systems Fail? http://arxiv.org/abs/2503.13657

The Anatomy of an Organization

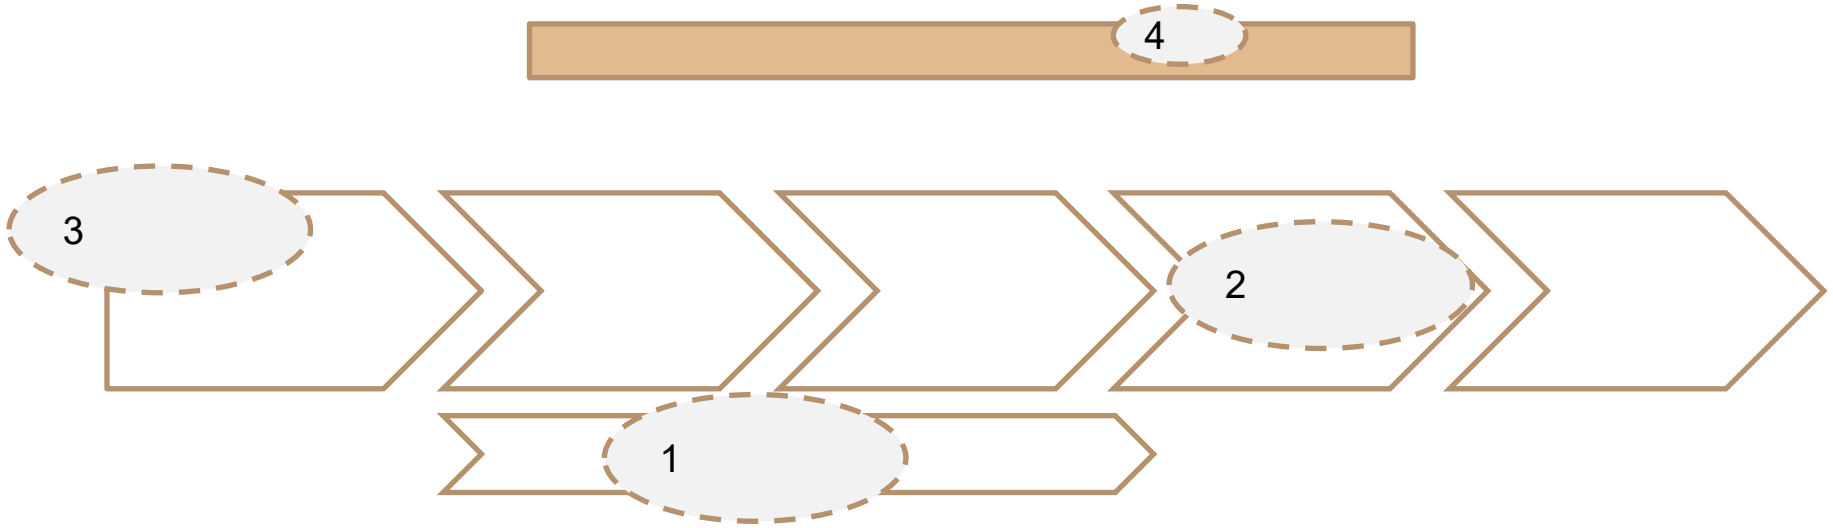

In order of complexity within an organizations value chain:

1: Support process;

2: Primary process

3: Primary or support process, externally facing (e.g., client or supplier)

4: Managerial process

Scope of an autonomous multi-agent AI team taking care of operations

Different Levels of Autonomy

Level 1: No AI

Level 2: Assisted driving

AI tools enhance productivity while humans maintain control

Level 3: Hands-Off

AI handles specific tasks with human monitoring

Level 4: Eyes-Off

Systems operate autonomously with oversight body for escalations

Level 5: Mind-Off

Full self-governance with minimal human input

Highway

- Low complexity
- Low risk
- High level of standardization

City Center

- High complexity
- High risk
- Low level of standardization

Split Context and Task-Specification

1 Agentic extended BPMN description of a Client Journey subprocess

Multi-agent process

Interaction layer

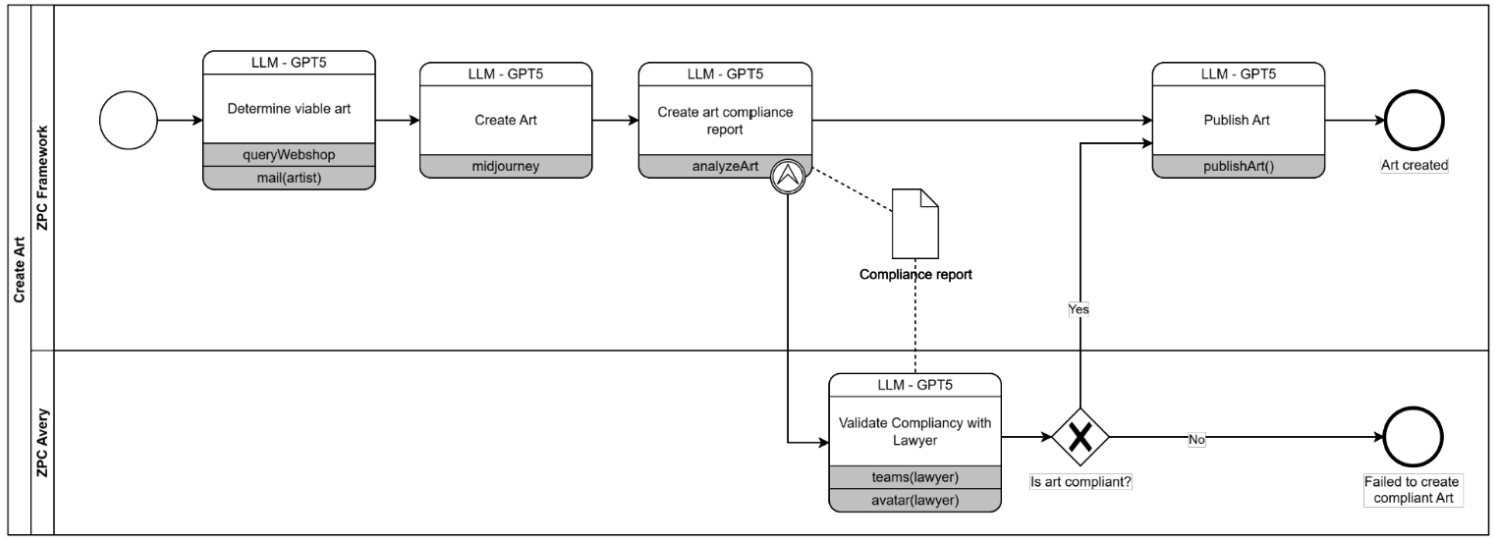

An Army of Micro-Agents

1 Agentic extended BPMN description of a Client Journey subprocess

2 Disposable agent pool

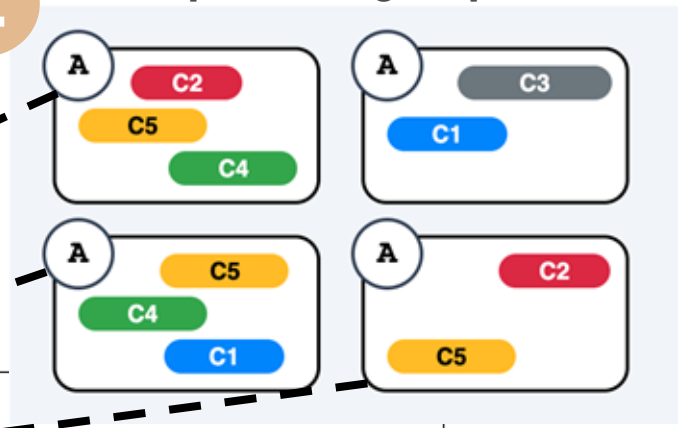

Multi-agent process

Interaction layer

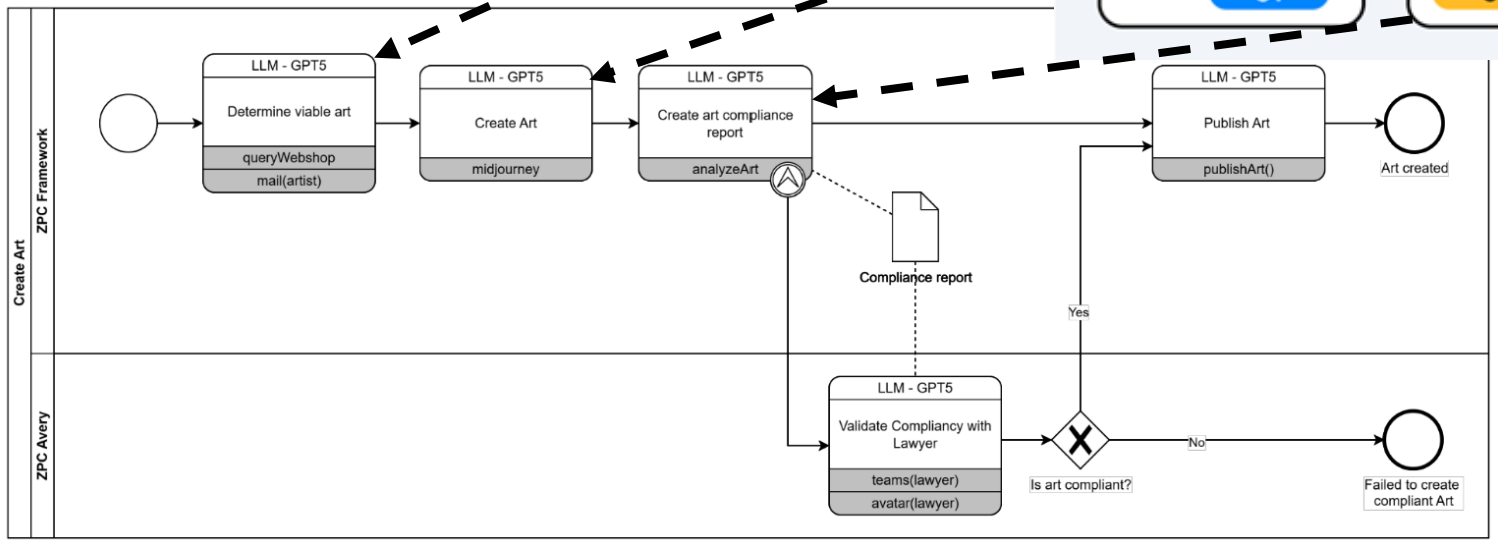

Human – AI Interaction

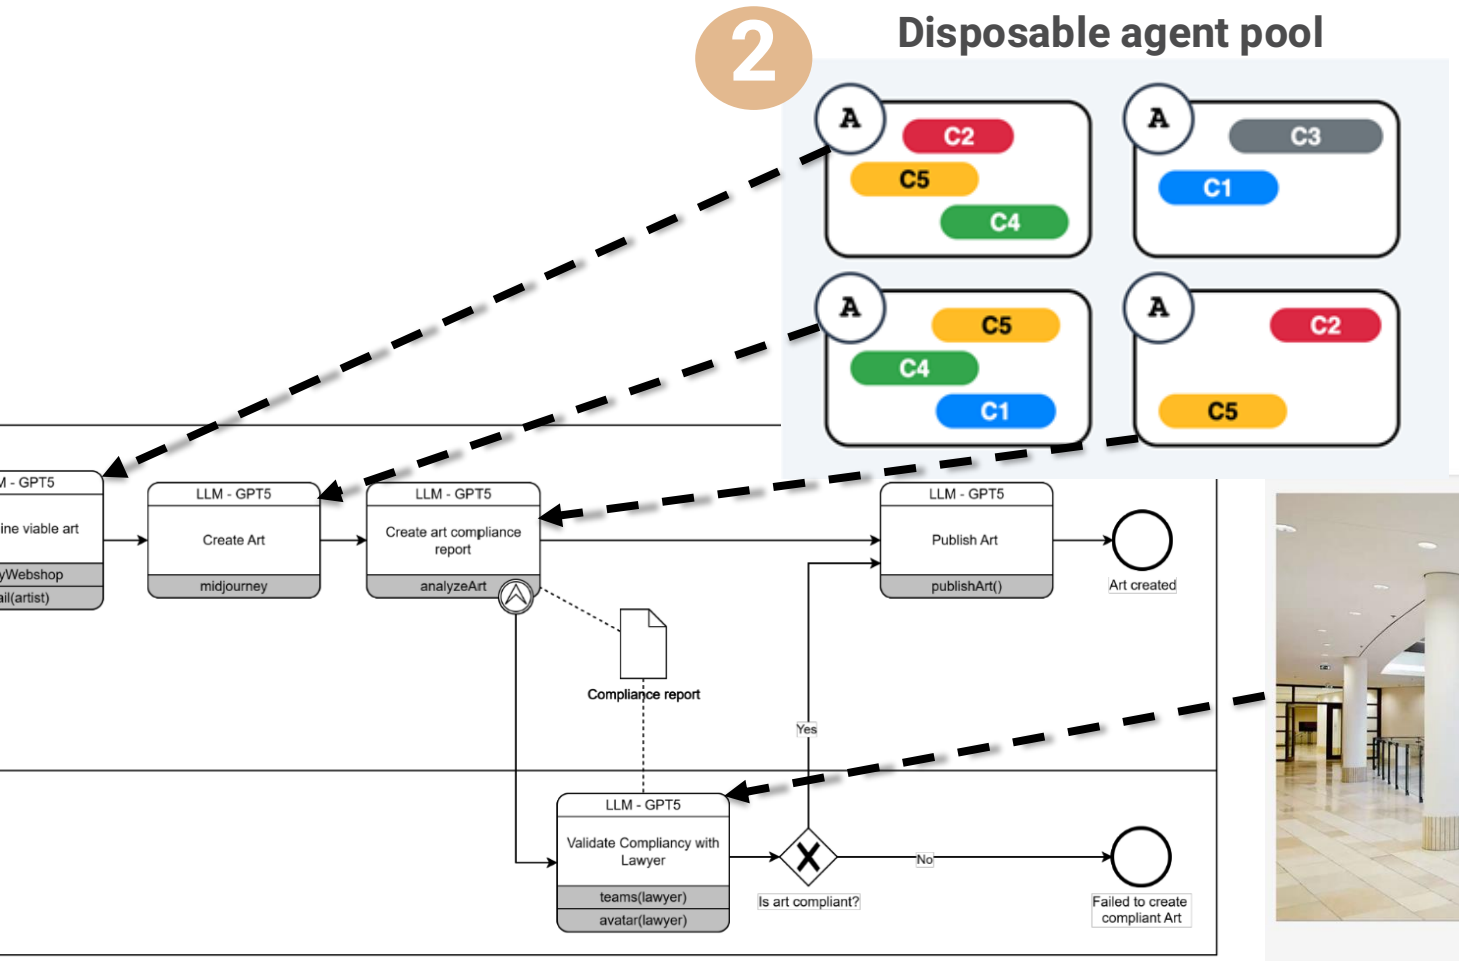

3 Real-life avatar interaction with the multi-agent AI team

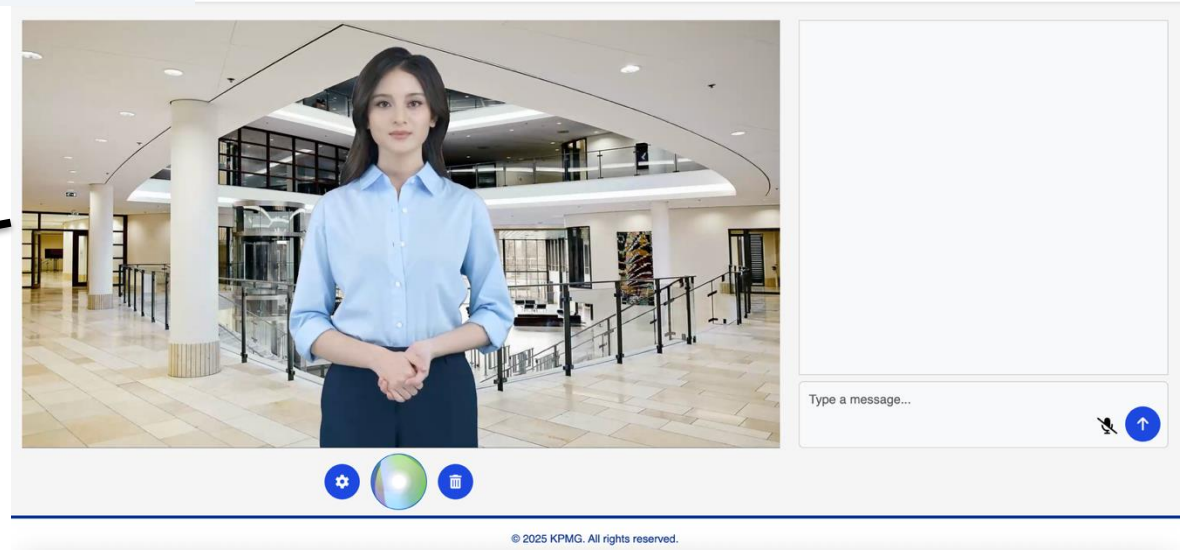

Petri-nets: Reproducibility, Explainability, Traceability

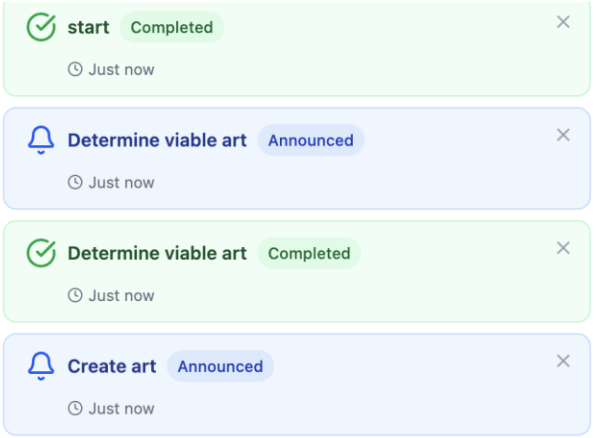

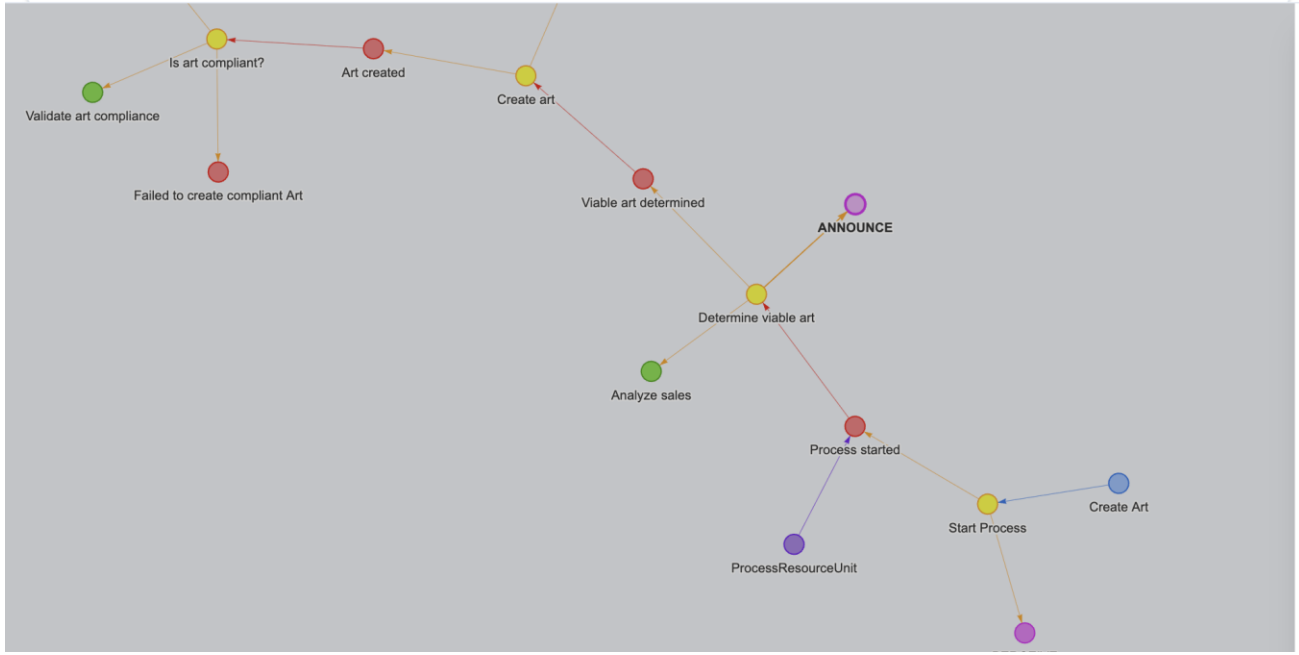

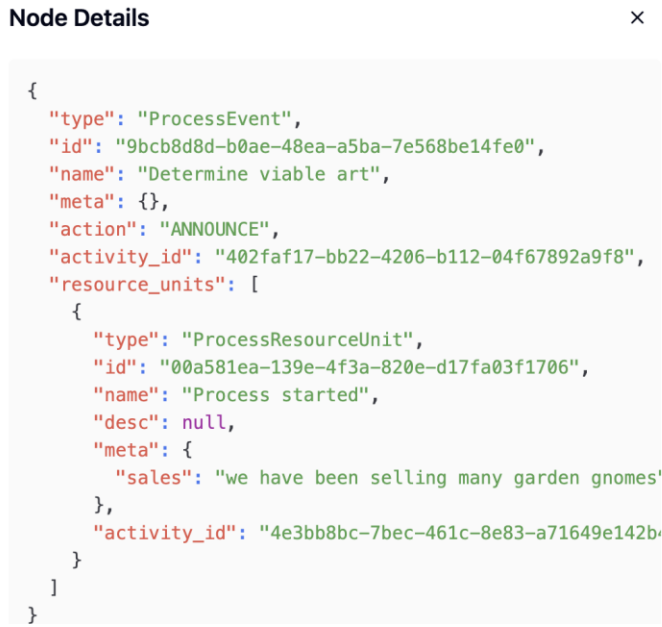

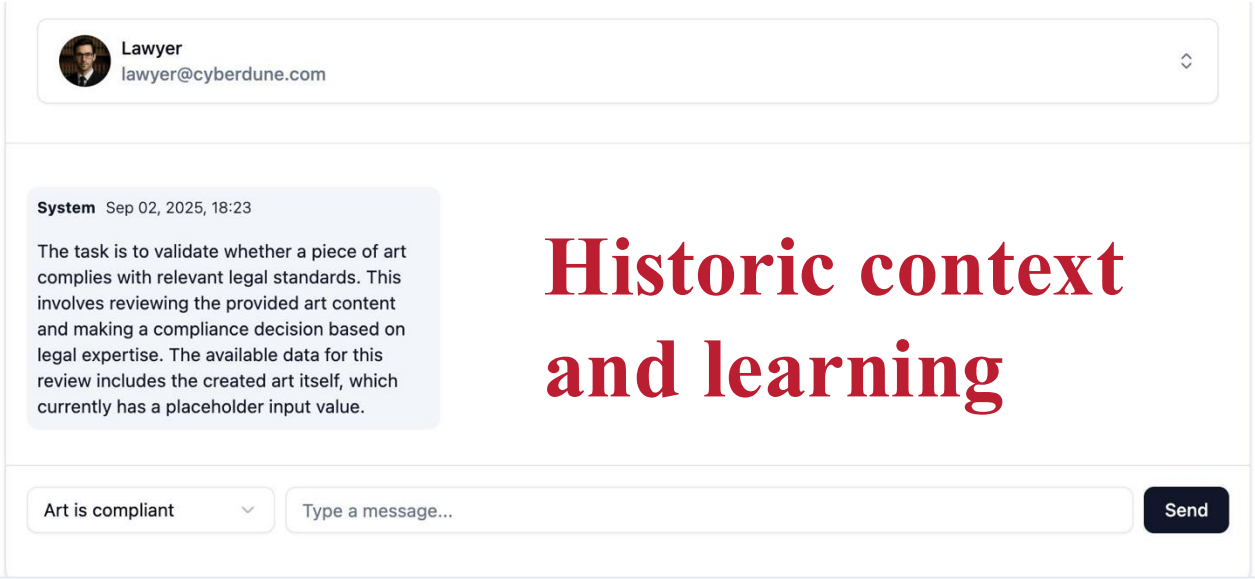

Historic context and learning

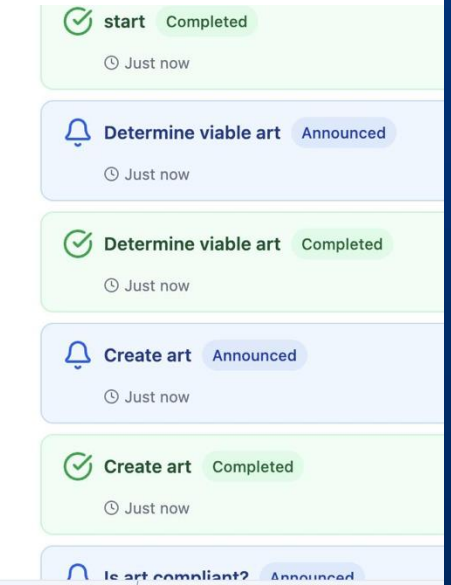

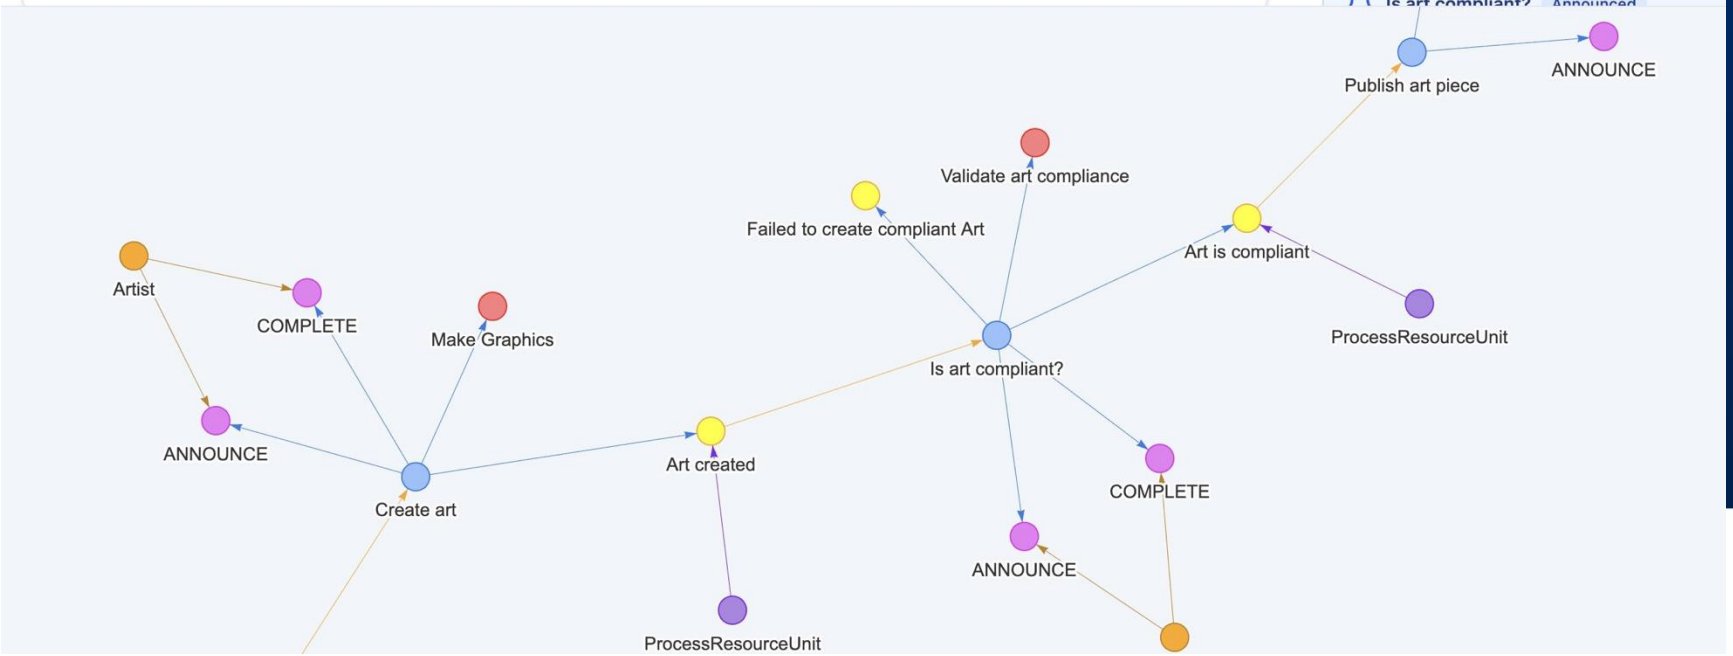

Specializing large language models for process modeling via reinforcement learning with verifiable and universal rewards

Research | Open access
Published: 18 December 2025
Volume 2, article number 26, (2025)
Cite this article

You have full access to this open access article
Download PDF
Save article
View saved research
Process Science
Aims and scope →
Submit manuscript →

Alessandro Berti, Xiaoting Wang, Humam Kourani & Wil M. P. Van der Aalst

Real-world cases: Financial Due Diligence

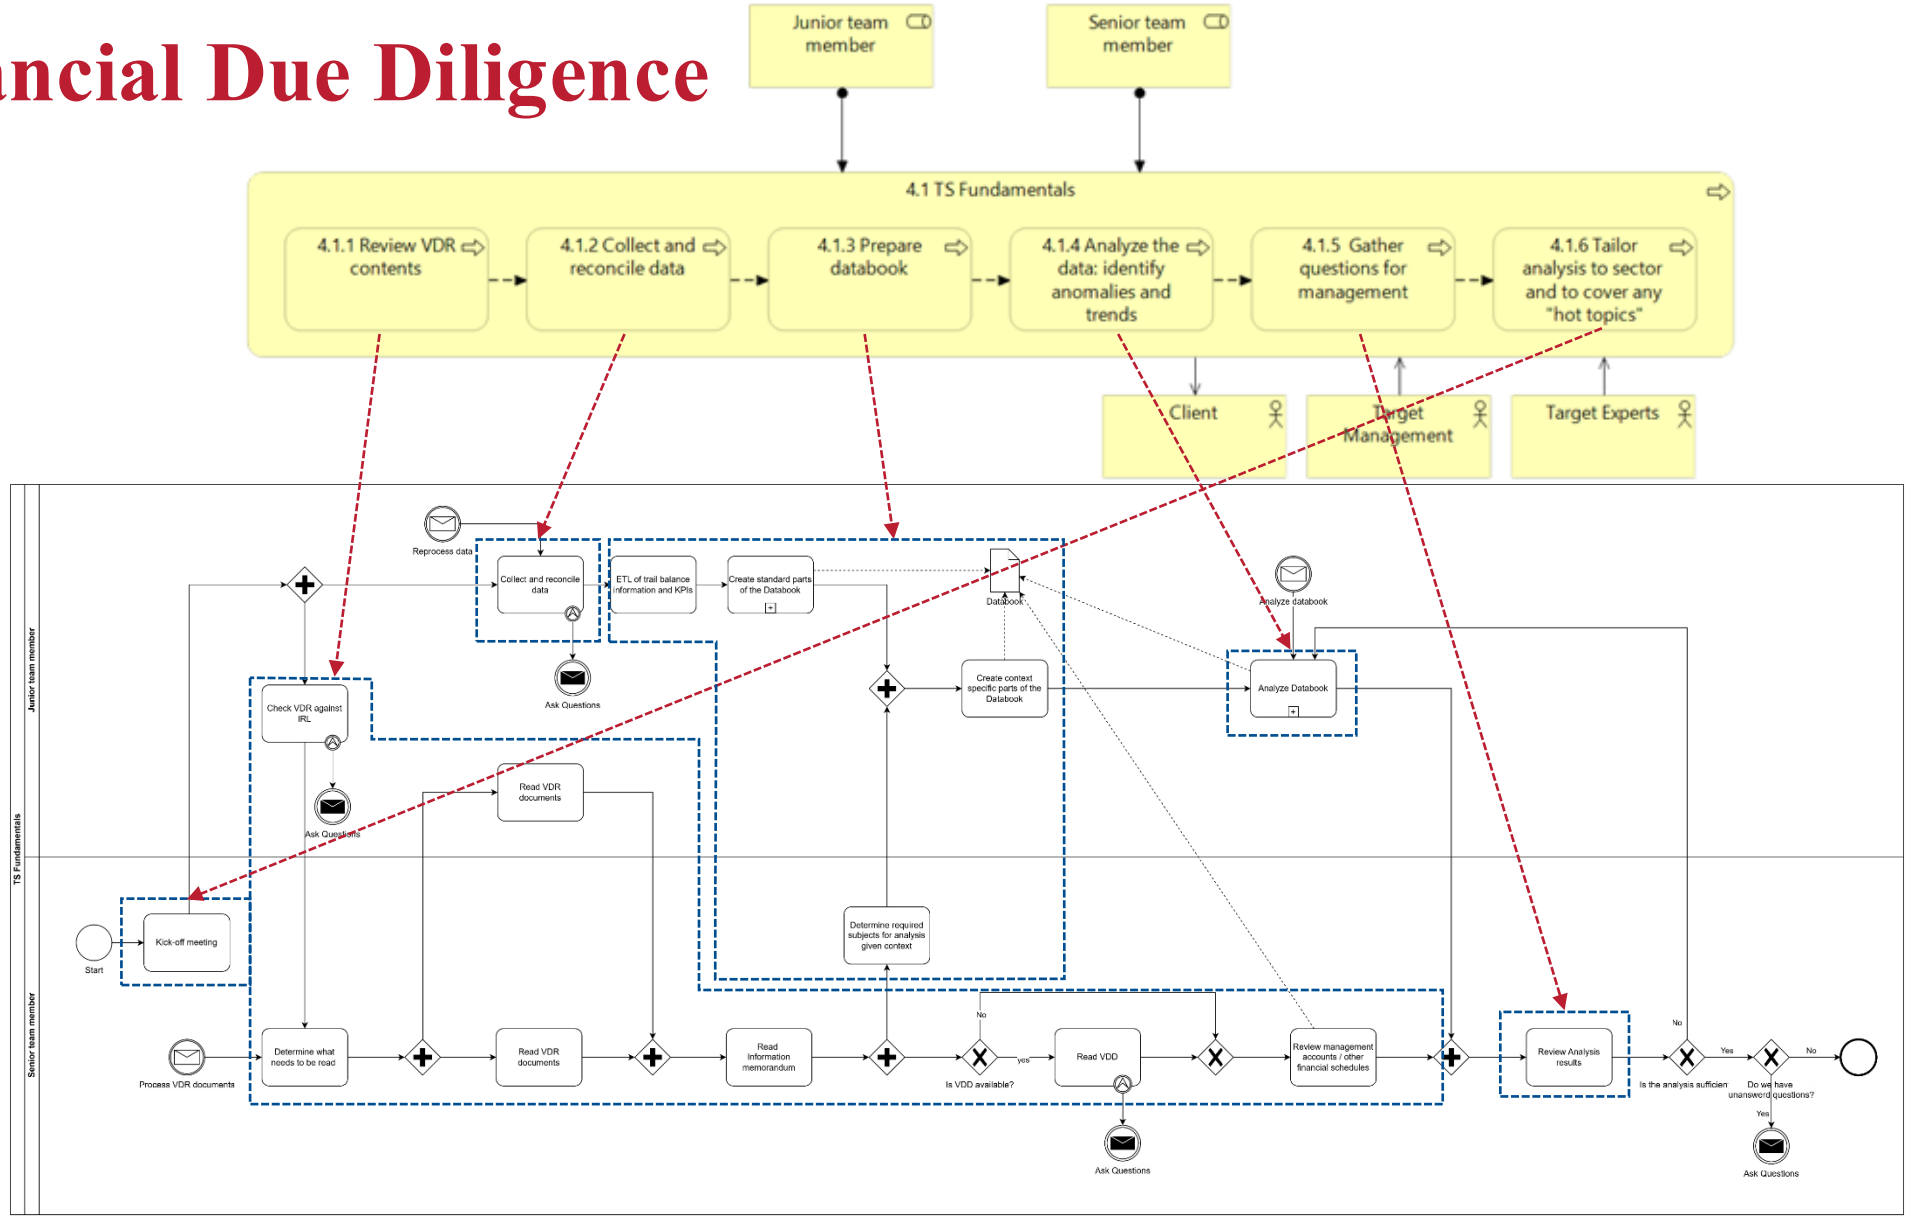

Agentic Design

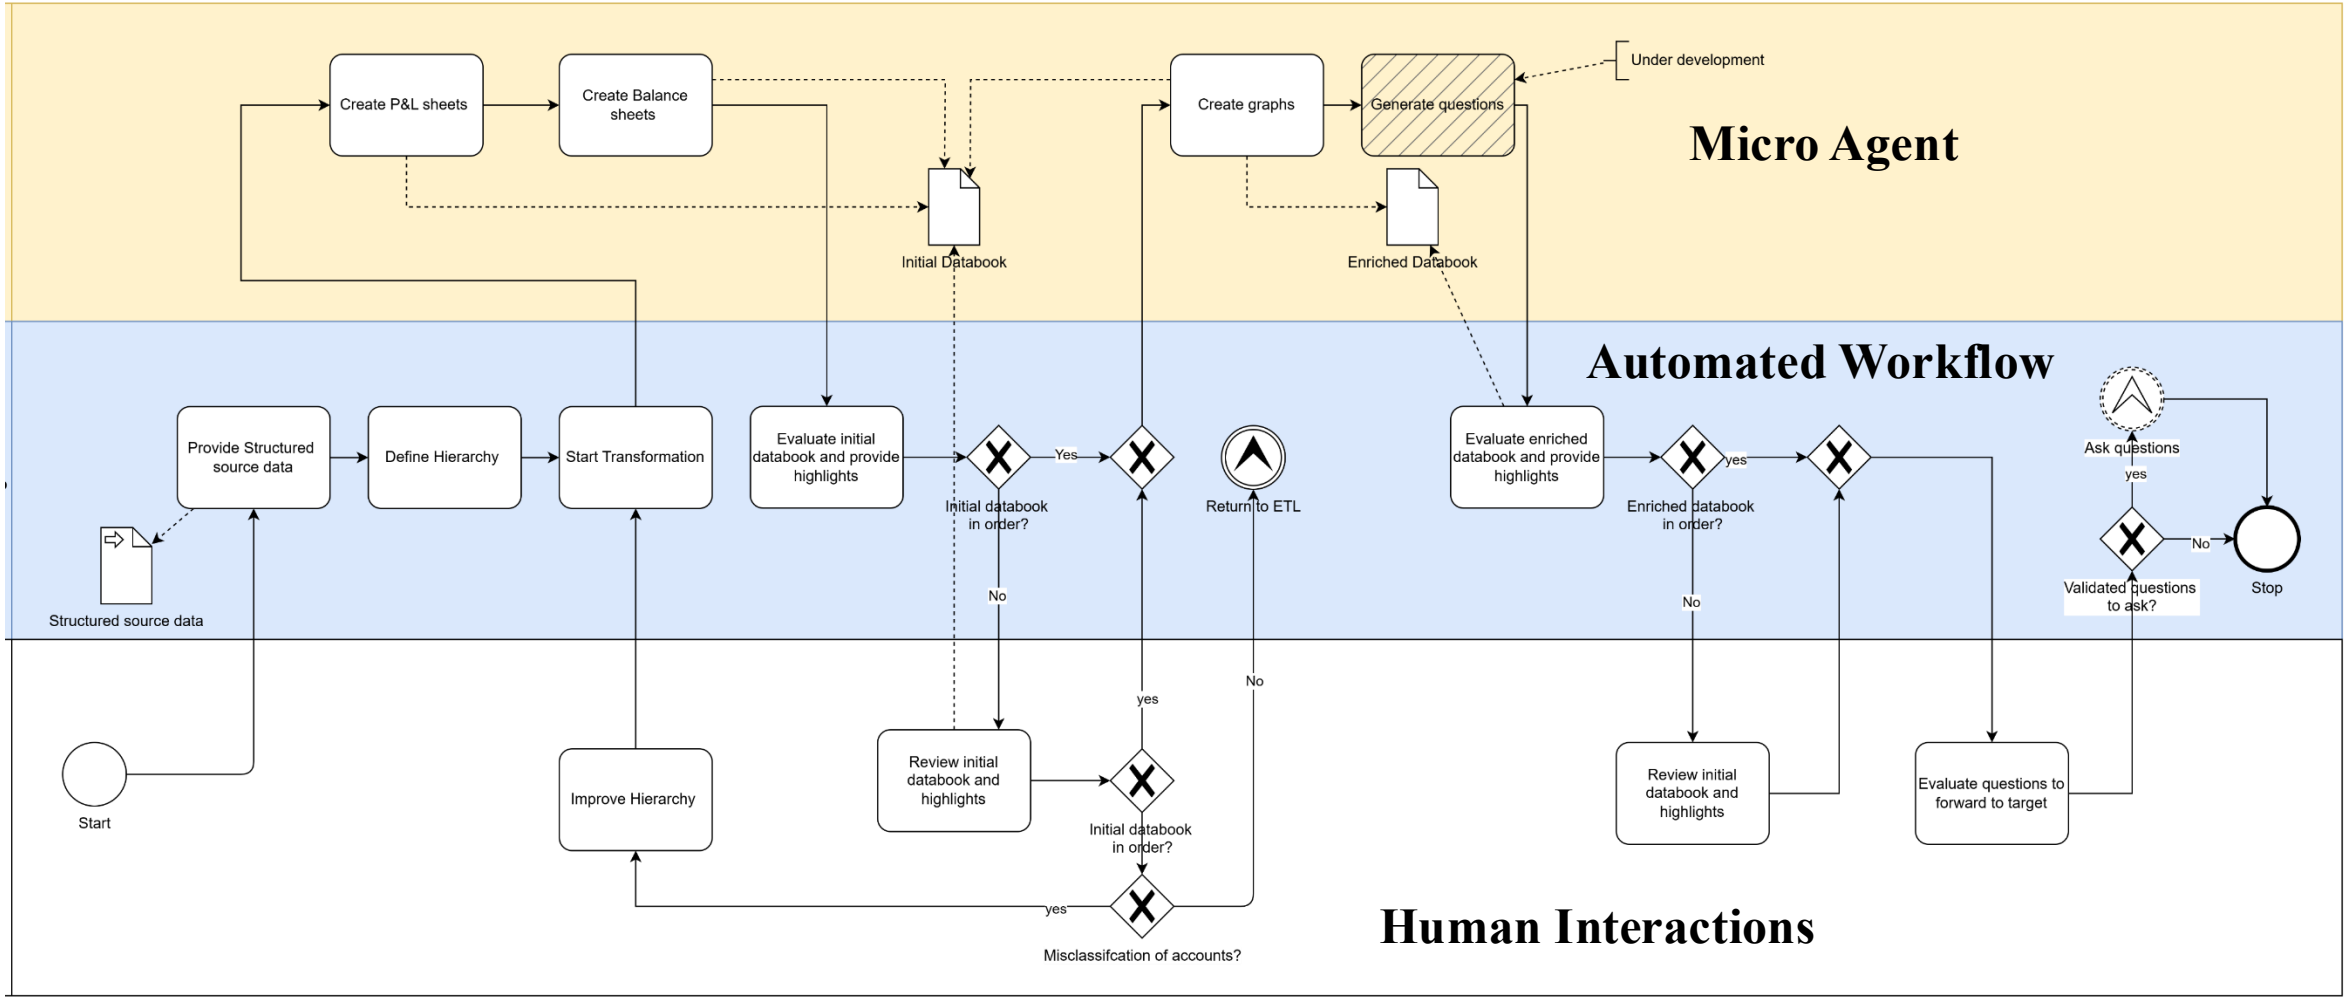

Ordering a Piece of Art

Autonomous AI-Workflow

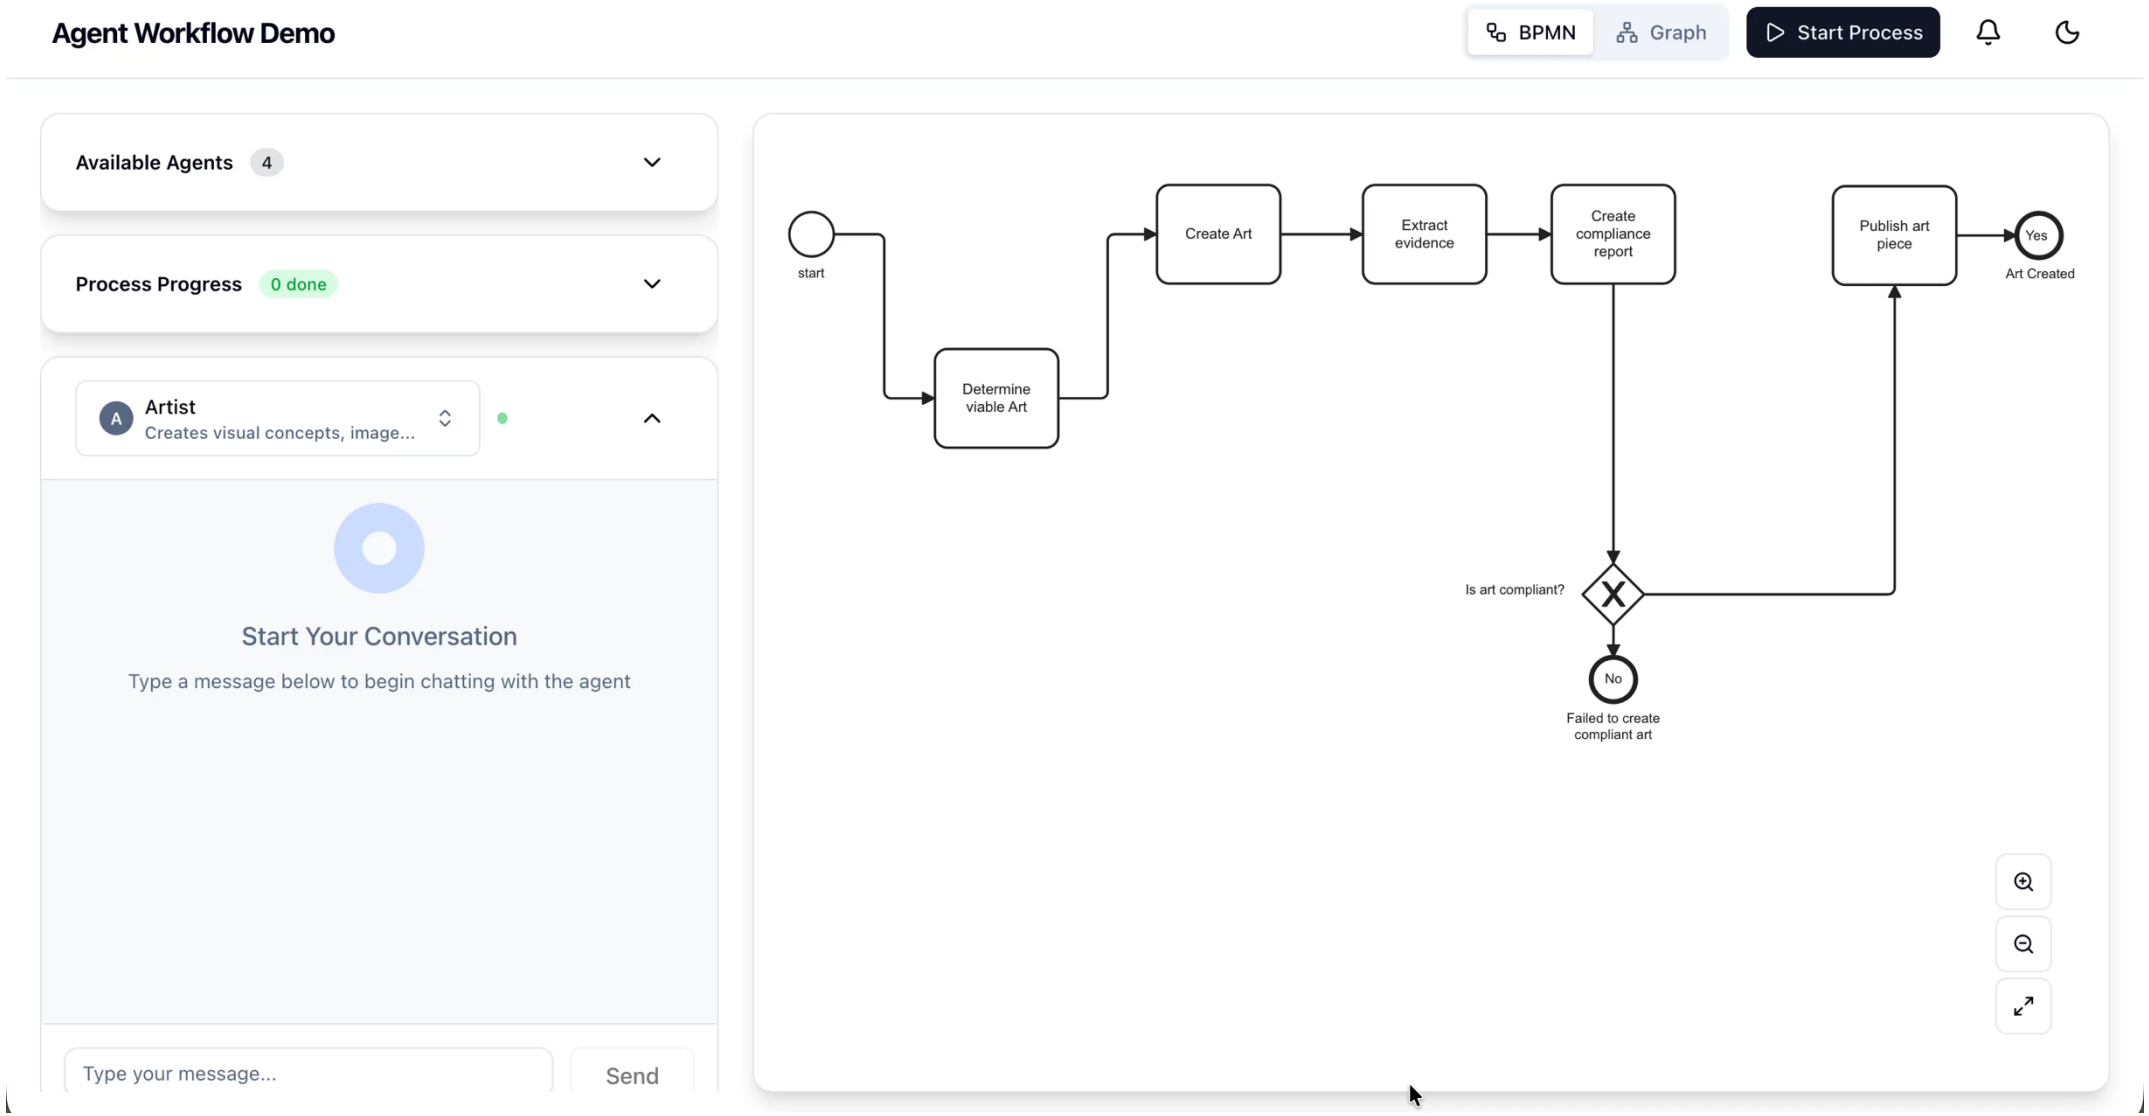

Resulting Artwork