CONCEPT-FORMING AI

AND THE OPERATIONAL WORLD MODEL

From AI That Is Taught Concepts by Humans to AI That Observes, Remembers, and Learns from the Organization

FREE OVERVIEW EDITION

A technical vision for a next-generation enterprise AI foundation integrating organizational observation, memory, concept formation, process recognition, and decision support

Mindware Research Institute, Inc.

Mindware Multi-Client Development Project | 2026

About This Free Overview

This document is a concise introduction to the concept of Concept-Forming AI and the Operational World Model. It explains the central thesis, technical architecture, enterprise applications, technologies already developed, and the Multi-Client Development Project. Detailed theoretical discussion, algorithmic considerations, implementation issues, risks, and development plans are provided in the full technical report.

The core proposition: AI for changing organizations must not only process concepts defined in advance. It must also form and update concept candidates from organizational data.

Contents

1. What Comes After Generative AI

Large language models have transformed AI from a specialist analytical technology into a general-purpose information-processing foundation that people can use through natural language.

Enterprises can now use AI to summarize and draft documents, search information, answer inquiries, generate software, and support analysis. By combining LLMs with retrieval-augmented generation (RAG), process mining, robotic process automation (RPA), and AI agents, organizations can also connect language models to internal information and external systems.

Yet the ability to generate language, retrieve documents, and operate tools is not the same as understanding the reality of an enterprise or organization.

Current enterprise AI still lacks a common foundation for:

- Continuously observing what is happening across the organization
- Connecting information distributed across departments and systems
- Discovering problems and opportunities that do not fit existing categories
- Remembering past decisions, actions, and outcomes as organizational experience
- Recognizing how concepts and processes change over time
- Relating the present situation to structurally similar past cases
- Providing AI agents with an organization-specific basis for judgment

Today’s AI can answer questions that an enterprise already knows how to ask. It is much less capable of identifying important questions that the enterprise has not yet recognized.

A serious quality problem or a new market opportunity may not be stated in any single document. It may be distributed across minor customer inquiries, sales notes, technical concerns, exception handling, and meeting discussions. A system is therefore needed that observes these signals continuously and forms structures and changes as concept candidates.

Figure 2 Enterprise AI and the Missing Cognitive Foundation

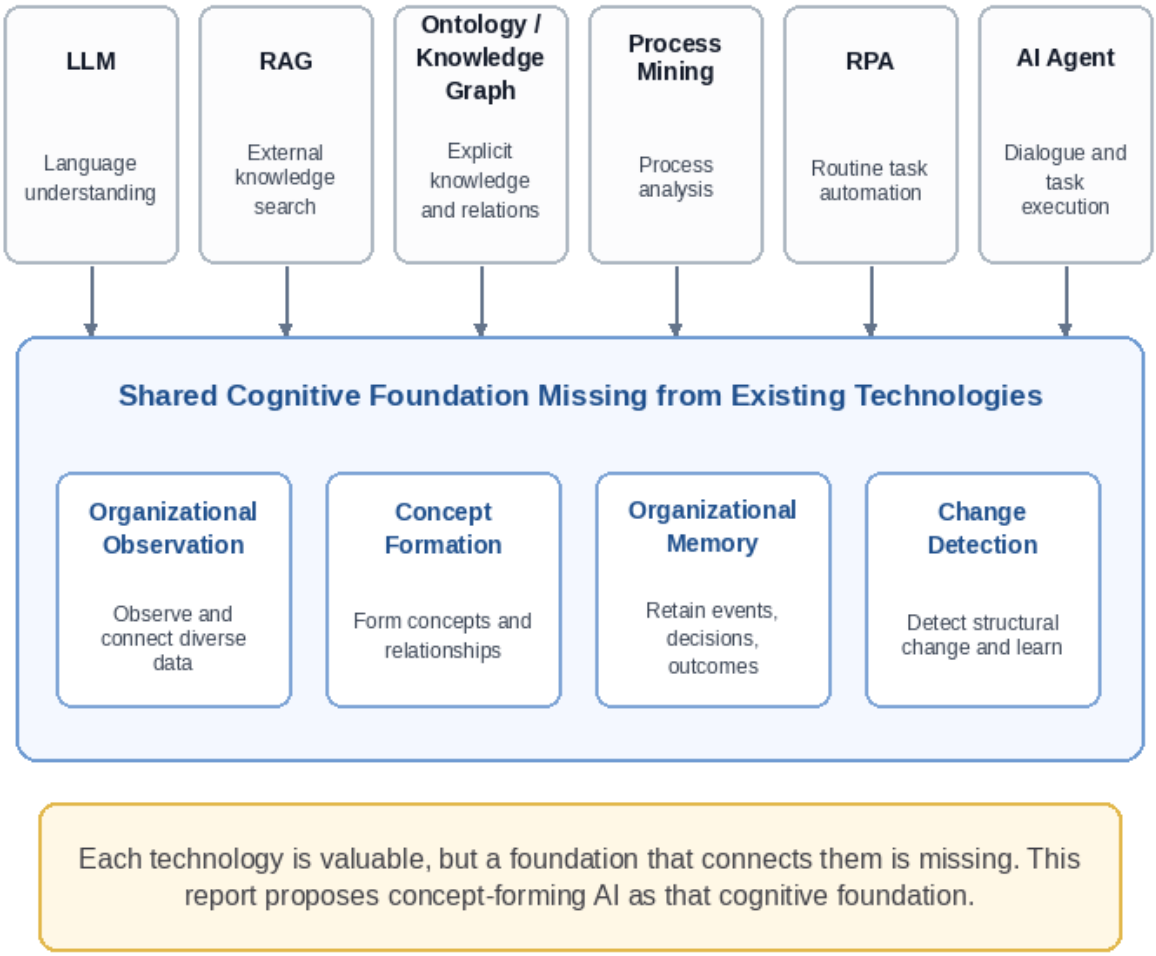

Figure 1. Existing enterprise technologies are valuable, but they do not by themselves provide a shared cognitive foundation for organizational observation, concept formation, memory, and change detection.

2. Central Thesis

There are limits to an approach in which humans define concepts in advance and teach them to a system. To deal with a changing reality, the system must form concepts from data and continuously update them.

Traditional enterprise information systems begin with human definitions. People specify data fields, categories, rules, workflows, and ontologies, and the system then processes information according to those definitions.

This method is effective when the domain is stable and the categories are already known. It is necessary for accounting classifications, legal requirements, product identifiers, security levels, and organizational authority.

However, emerging customer behavior, unknown quality problems, changing market segments, informal work patterns, and new technical bottlenecks cannot all be defined before they appear.

Existing categories organize the world that people have already recognized. To discover changes outside those categories, the system must be able to form structures directly from data.

Figure 1 Human-Defined AI vs. Concept-Forming AI

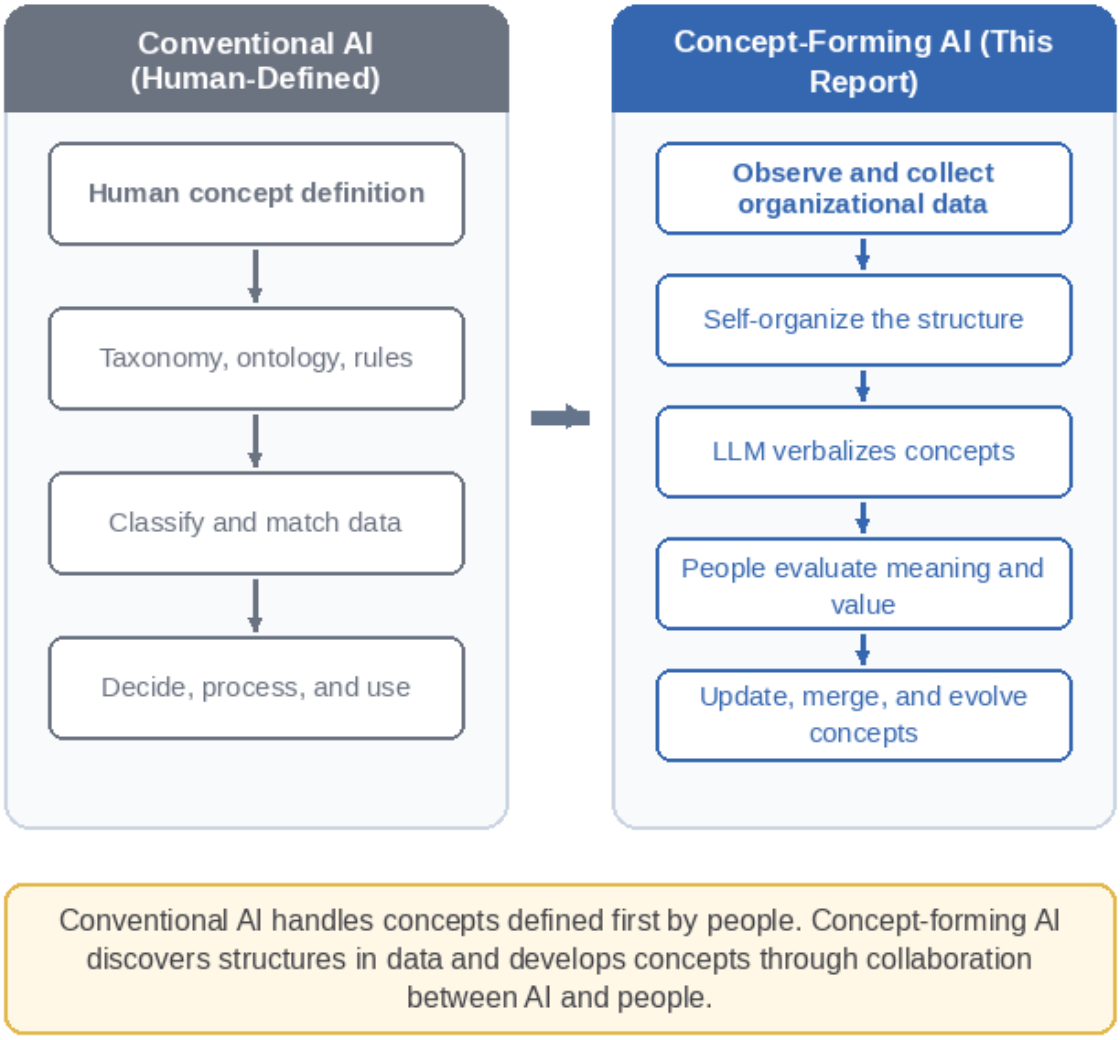

Figure 2. Conventional AI processes concepts defined first by people. Concept-Forming AI observes organizational data, forms structures, verbalizes them, and develops concepts through human evaluation.

3. What Is Concept-Forming AI?

Concept-Forming AI does more than assign data to predefined categories. It forms patterns, continuities, branches, boundaries, and exceptions in data and presents them as concept candidates that people can examine and use.

A basic concept-formation process consists of the following steps:

1. Collect documents, dialogue, database records, event logs, and other relevant data.
2. Transform the data into comparable representations using embeddings and numerical features.
3. Use self-organizing learning to form representative nodes and neighborhood structures.
4. Represent continuity, branching, and boundaries through networks and a minimum spanning tree.
5. Form clusters or other concept candidates when useful.
6. Use an LLM to summarize and verbalize the characteristics of the structures.
7. Have people evaluate their meaning, relevance, and value.
8. Update the structures as new data and outcomes become available.

The system does not discover a single objectively correct conceptual system. The same data can support different structures depending on the purpose, features, weights, and direction of attention.

Self-organization discovers structures. The LLM verbalizes them. People evaluate meaning and value.

Human responsibility therefore does not disappear. It shifts from designing every category in advance to evaluating structures formed by the system, setting the purpose of observation, and deciding which concept candidates matter.

Figure 3 Concept-Forming AI Processing Flow

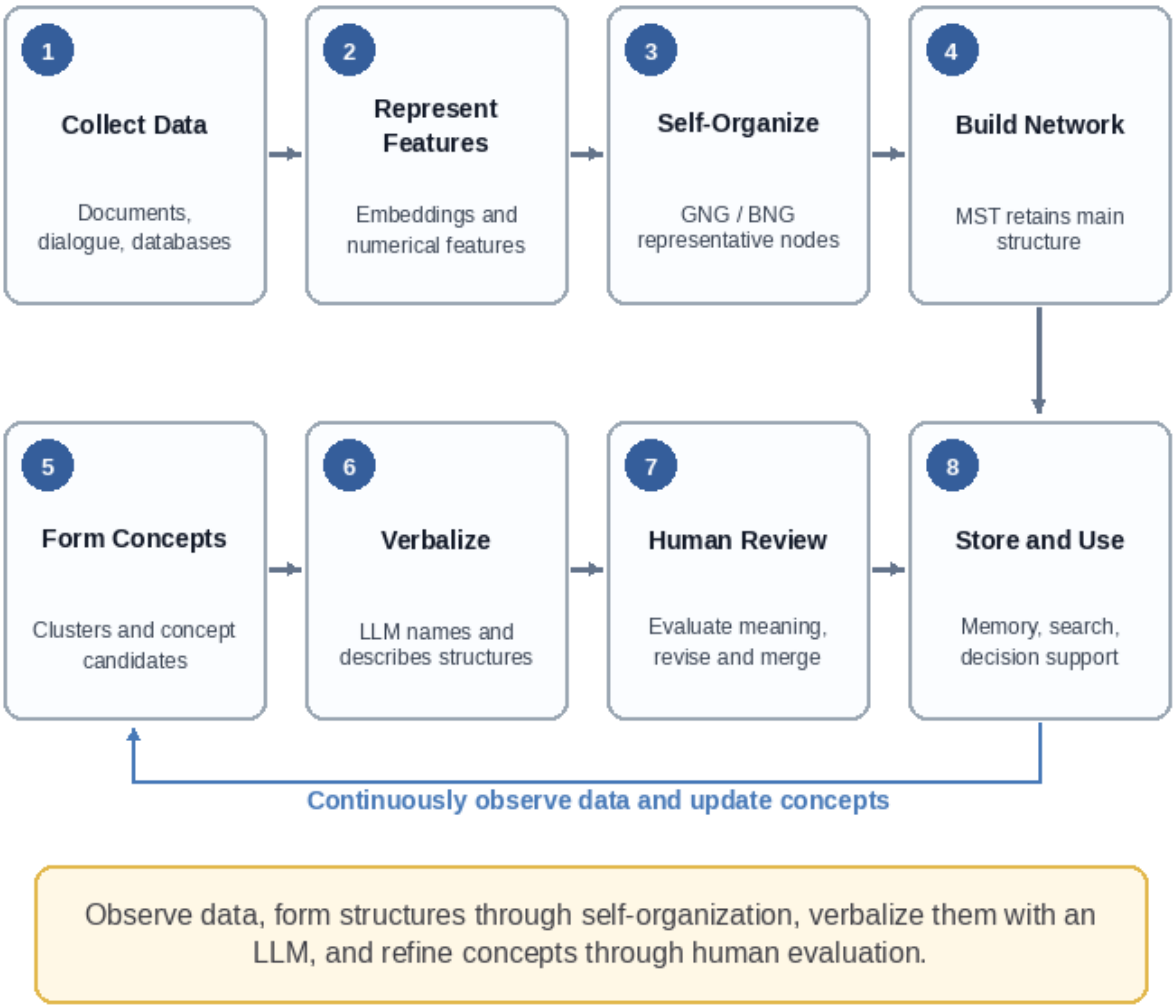

Figure 3. Concept formation is a multi-stage cycle that combines representation, self-organization, networking, LLM-based verbalization, human review, and organizational use.

Figure 4 High-Dimensional Structure and 2D Display

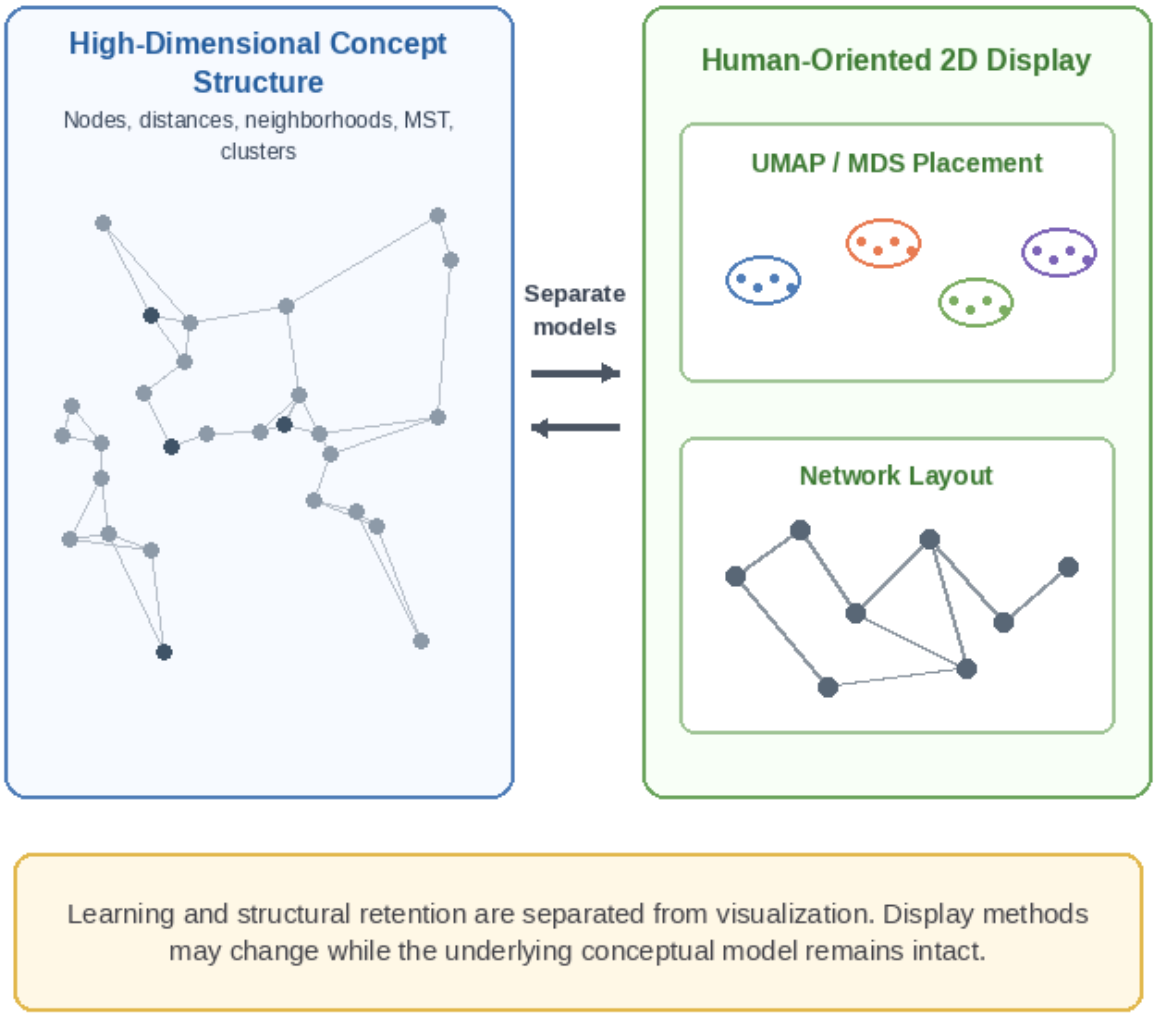

Figure 4. The high-dimensional concept model should be separated from its two-dimensional display. Visualization methods may change while the underlying structure remains intact.

4. The Operational World Model

The Operational World Model is the integrated foundation that connects concept formation with organizational observation, memory, process recognition, judgment, action, and outcome feedback.

It continuously observes the operational world of an enterprise and maintains an updateable internal model of entities, events, relationships, concepts, processes, hypotheses, decisions, actions, and outcomes.

Its basic cycle is:

Observe → Form Information Structure → Form Concepts → Remember → Hypothesize and Decide → Act → Observe Outcomes → Update the Internal Model

The “world” in this term is not limited to physical space. For an enterprise it includes customers, products, technologies, markets, contracts, projects, inquiries, meetings, people, departments, problems, risks, opportunities, decisions, actions, and temporal change.

The Operational World Model is not intended to replace CRM, document management, ERP, support systems, or project management systems. It is designed as an observational and cognitive layer above existing systems, connecting their information through common entities, events, concepts, time, and processes.

Figure 5 Operational World Model: Core Cycle

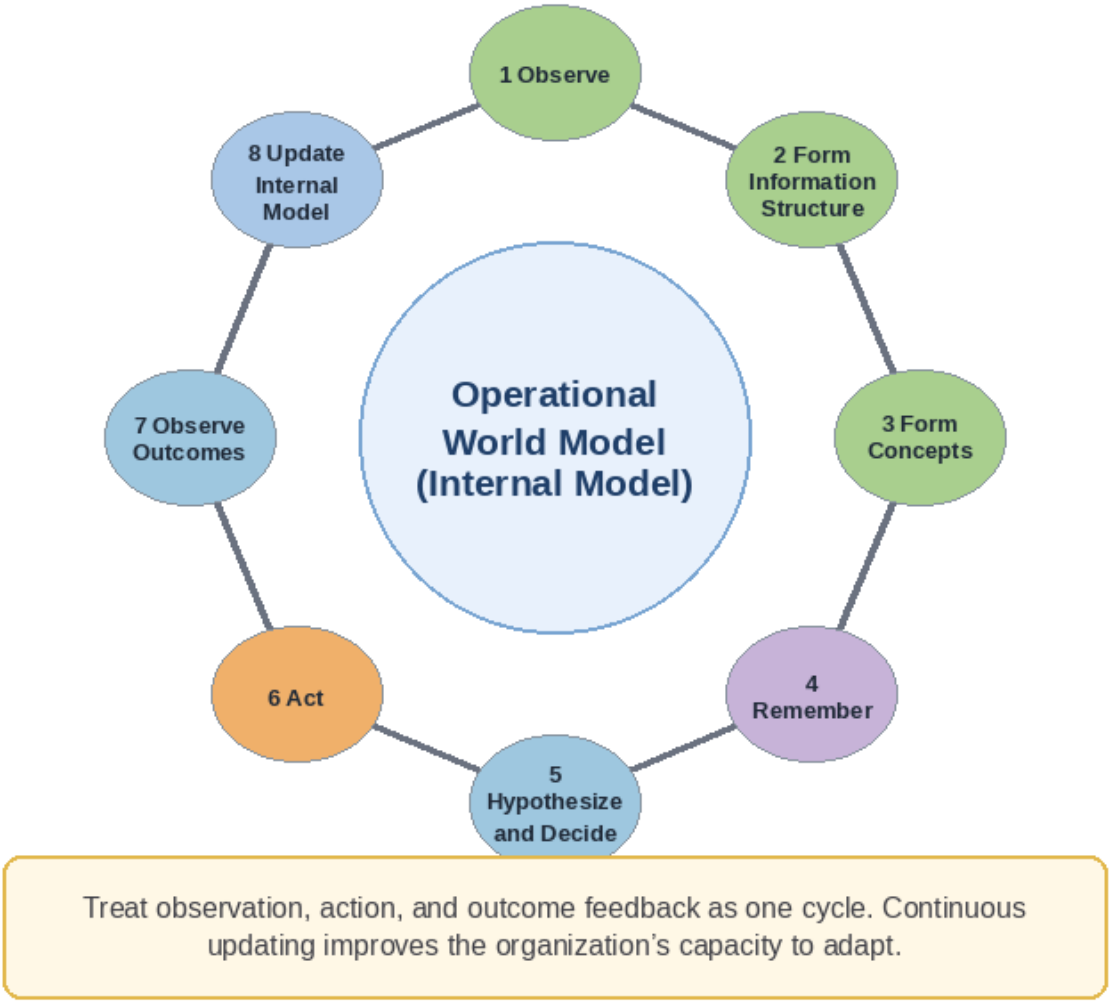

Figure 5. The Operational World Model treats observation, concept formation, memory, judgment, action, and outcome feedback as a continuous cycle.

Seven-Layer Architecture

- Information Connection: connections to documents, databases, messages, event logs, APIs, and external sources
- Information Extraction: embeddings, entities, events, actors, objects, time, actions, and outcomes
- Concept Formation: self-organization, representative nodes, networks, clusters, concept profiles, and change detection
- Memory: LLM Wiki, long-term associative memory, episodic memory, decisions, and outcomes
- Recognition and Reasoning: state recognition, hypotheses, similar cases, process recognition, risks, and options
- Interaction and Action: dialogue, dashboards, alerts, AI agents, RPA, and external operations
- Governance: permissions, approval, audit, security, human correction, and data governance

Figure 6 Seven-Layer Operational World Model Architecture

Figure 6. AI agents operate on top of a cognitive foundation. They are part of the interaction and action layer, not a substitute for the lower layers.

5. How It Differs from Existing Technologies

Large Language Models

LLMs understand and generate language, but they do not automatically and continuously observe what is happening inside a particular enterprise. In the Operational World Model, the LLM is an essential component for extraction, verbalization, hypothesis generation, and dialogue—not the world model itself.

Retrieval-Augmented Generation

RAG retrieves documents relevant to a question. Concept-Forming AI observes a dataset before a question is fully known and forms unknown problem candidates and structures. RAG is useful for answering known questions; concept formation is intended to help discover questions that have not yet been recognized.

Ontologies and Knowledge Graphs

Ontologies and knowledge graphs are effective for maintaining explicit concepts, relationships, constraints, and institutional definitions. The Operational World Model adds an adaptive layer in which new concept candidates can emerge from data, while stable human-defined concepts remain in force where required.

Process Mining

Process mining reconstructs actual workflows from event logs. Yet many important decisions and exceptions exist only in email, chat, meetings, and reports. The Operational World Model seeks to combine structured event logs with events and decisions extracted from natural language.

AI Agents

AI agents can act, but the ability to act is not the same as the ability to judge. The Operational World Model provides agents with the current state, relevant concepts, similar cases, constraints, hypotheses, and past outcomes, and records the results of their actions as organizational experience.

6. Organizational Memory

Enterprise knowledge does not exist only in policies, manuals, and documents. It also exists in sequences of experience: what the situation was, who was involved, what was recognized as the problem, which alternatives were considered, what was decided, what was done, and what happened afterward.

Long-Term Associative Memory

Long-term associative memory retrieves past information according to its relationship with the current situation—not merely by chronological order or keyword overlap. It can use semantic similarity, shared concepts, common customers or products, similar process structures, and comparable outcomes.

Episodic Memory

Episodic memory stores events, conditions, judgments, actions, and outcomes as a coherent unit of organizational experience. A serious customer complaint, for example, should not be stored only as a document. The preceding signals, internal discussions, response, technical correction, customer reaction, and long-term result should be connected as one episode.

This allows the organization to recall not only what was written in the past, but how a similar situation developed and which actions produced which outcomes.

7. Technologies Already Developed

This vision does not begin from zero. It is based on long-term experience with self-organizing maps and approximately three years of recent experimentation and software development combining self-organization, network structures, embeddings, LLMs, knowledge bases, and hypothesis formation.

ConceptMiner

ConceptMiner is a concept-formation engine that combines embeddings, Growing Neural Gas, Batch Neural Gas, network structures, minimum spanning trees, clustering, and LLM-based verbalization to form and explore conceptual structures in text and data. It provides the concept-formation layer of the Operational World Model.

Concept Index

Concept Index adds conceptual search and segmentation to existing relational databases such as PostgreSQL. It reads existing data, forms embeddings and concept structures externally, and returns the relationships between records, representative nodes, neighborhoods, and clusters in a form that can be used through SQL. It is designed to avoid replacing or directly modifying the production database.

ThinkNavi

ThinkNavi integrates LLM Wiki, long-term associative memory, automated research, hypothesis diagnosis, concept maps, and dialogue to support human investigation, thought, and decision-making. It provides the knowledge, memory, reasoning-support, and interaction layers of the Operational World Model.

The next challenge is to integrate these technologies into a single continuous operational cycle suitable for enterprise use.

8. Core Functions to Be Developed

Episodic Memory

Store events, situations, judgments, actions, and outcomes as reusable units of experience.

Natural-Language Process Mining

Extract actors, objects, actions, purposes, judgments, and outcomes from email, chat, meeting transcripts, and reports, and combine them with structured event logs to form process models.

Continuous Observation

Observe internal data sources periodically or in response to events and incorporate new information into the model.

Change Detection

Detect not only changes in volume, but also the emergence, division, merger, movement, and disappearance of concepts, relationships, and process patterns.

Integrated Data Model

Connect documents, records, entities, events, concepts, processes, hypotheses, decisions, actions, outcomes, time, sources, and permissions.

AI-Agent Integration

Provide agents with state, memory, concepts, constraints, and similar cases, and return their actions and outcomes to organizational memory.

Enterprise Operations and Governance

- Authentication, access control, approval, and audit
- Scheduled execution, error handling, backup, and model updates
- Windows, Linux, Docker, on-premises, private-cloud, and local-processing options
- Bring Your Own Key (BYOK) and control over information sent to external models

9. Potential Enterprise Applications

Voice of the Customer and Inquiry Analysis

Discover problem groups, new customer requirements, churn signals, and emerging usage patterns that cut across existing categories.

Early Detection of Quality Problems

Relate weak signals distributed across departments and systems and present candidates for significant problems before they become obvious.

Product and Technology Development

Integrate customer needs, market information, technical literature, experimental results, and internal discussions to form new development themes and shared technical problems.

Observation of Actual Business Processes

Identify real workflows, exceptions, rework, informal coordination, and dependence on particular individuals—not merely the official procedures described in manuals.

Organizational Knowledge Transfer

Preserve decisions, actions, outcomes, failures, exceptions, and temporary responses as organizational experience rather than only as static documentation.

Management and Decision Support

Organize the current state, related concepts, past cases, hypotheses, options, risks, uncertainties, and missing information without forcing them into one premature conclusion.

A Shared Judgment Foundation for AI Agents

Provide multiple agents with a common representation of enterprise state, concepts, memory, constraints, and governance conditions.

The value of the Operational World Model cannot be measured only by labor savings. It should also be evaluated by whether it helps the organization detect problems earlier, discover opportunities, reuse past experience, clarify assumptions, reveal differences in perception across departments, and form new product or research themes.

10. The Roles of Humans and AI

The Operational World Model is not a proposal for an omniscient AI that automatically understands the enterprise and replaces human management or responsibility.

Concept structures and hypotheses formed by the system contain uncertainty. Similarity does not prove a common cause, and a network connection does not establish causality. A system can observe only what has been recorded and may overlook unrecorded work, silent customers, and informal context.

AI Responsibilities

- Continuously observe large volumes of information
- Present structures, similarities, exceptions, and changes
- Recall related past information and episodes
- Generate hypothesis and question candidates
- Execute authorized routine actions

Human Responsibilities

- Define the purpose and direction of observation
- Evaluate the meaning and value of concept candidates
- Adopt, revise, merge, split, or reject hypotheses and concepts
- Approve consequential actions and set permissions
- Provide organizational values, priorities, and accountability

The purpose of the system is not to increase the amount of information people must review. It is to direct limited human attention toward important change, exceptions, disagreement, and uncertainty.

11. Multi-Client Development Project

The common productization of the Operational World Model will be pursued through a Multi-Client Development Project involving multiple participating companies.

This is not custom development based on the requirements of a single enterprise. The project organizer retains responsibility for the product concept, technical direction, architecture, and implementation. Problems and requirements presented by participating companies are compared and generalized into specifications that can be used across organizations.

Phase 1 Basic Terms

- Project period: October 1, 2026 to September 30, 2027
- Participation fee: JPY 2.5 million per unit
- Participants: normally up to two people per unit
- Payment: 50% at application and contract, and the remaining 50% before the project begins
- Minimum number of participating companies: 20
- Target number of participating companies: 70

What Participating Companies Receive

- Technical documents and development information
- Regular online meetings and progress reports
- Access to prototypes as they become available
- Opportunities to evaluate results and submit requirements
- Final deliverables and software usage rights under the contract
- Source-code or GitHub access where included in the participation terms

Providing raw corporate data to the project organizer is not mandatory. Companies may participate through anonymized data, sample data, data structures, use scenarios, or evaluation results generated inside their own environment.

12. Why Use a Multi-Client Model?

If the system is developed for only one enterprise, it may become overly specialized to that company’s systems, terminology, and workflow. If it is developed without direct contact with real enterprise requirements, it may remain an abstract research system that cannot be deployed.

A multi-client model makes it possible to:

- Distinguish company-specific requirements from broadly shared requirements
- Compare data structures and use cases across industries
- Identify practical security and deployment conditions
- Determine which management and governance functions are actually required
- Obtain user evaluation from the early stages of development
- Form specifications suitable for general commercial distribution
- Distribute development cost and risk across multiple participating companies

The project is therefore closer to joint productization than to ordinary collaborative research. Its objective is not merely to report research results, but to complete software that enterprises can install and use.

13. Full Technical Report

This free overview introduces the overall concept. The full technical report provides the detailed theoretical background, algorithmic architecture, comparison with existing approaches, implementation considerations, enterprise applications, development risks, evaluation methods, and complete Multi-Client Development Project plan.

Full Technical Report: “Concept-Forming AI and the Operational World Model”
More than 220 pages | Price: JPY 70,000

The full report is intended for executives, R&D leaders, new-business managers, AI and digital-transformation leaders, information-systems managers, product and technology planners, and organizations evaluating the introduction of AI agents and organizational memory.

Purchasing the report does not require participation in the Multi-Client Development Project. It is designed to stand alone as a technical and strategic resource for evaluating generative AI, RAG, ontologies, process mining, AI agents, and organizational memory as parts of a broader enterprise cognitive foundation.

14. Conclusion

Large language models have enabled AI to work with human language. AI agents are beginning to enable AI to operate external systems and take action. What is still missing is the cognitive foundation that connects language and action through organizational observation, concepts, memory, experience, and feedback.

Observe the organization. Form structures from information. Form concepts from structures. Remember experience. Support hypotheses and decisions. Observe the results of action. Update the world model.

The Operational World Model is not intended to be a finished and uniquely correct model of the enterprise. Organizational information is incomplete, meanings change, and multiple interpretations may coexist.

The goal is to create an updateable organizational cognitive foundation that retains evidence, uncertainty, different perspectives, concept change, decisions, actions, and outcomes—and allows people and AI to improve that model together.

From AI that processes concepts defined by humans to AI that observes the organization, forms concepts, and learns from experience.