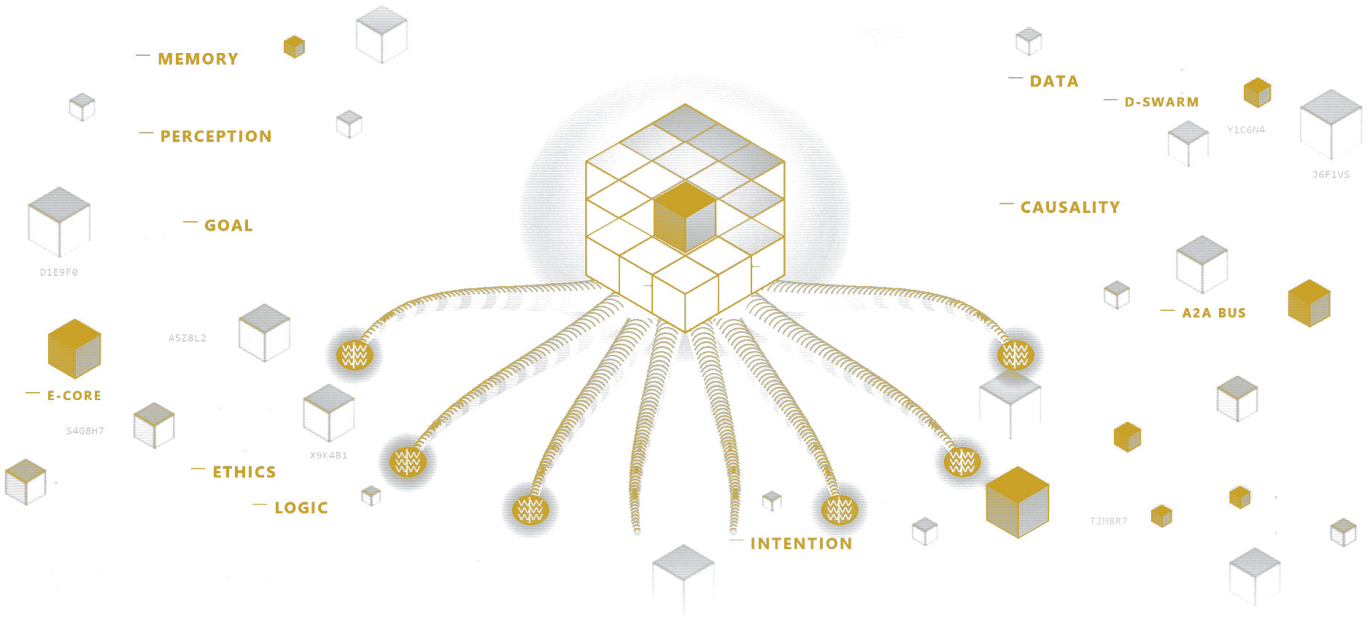

RESEARCH WHITE PAPER · VERSION 17

SOLSICE RESEARCH

OCTOSWARM PLATFORM

BIMODAL AGENTIC ARCHITECTURE

& DYNAMIC RECODING

SOLITARY MODE
E-Core Architecture

COLLECTIVE MODE
D-Swarm Topology

CORRESPONDENCE
sjehan@solsice.com

FOREWORD

> *Although this approach is only an initial research avenue that has since evolved considerably towards purely topological optimizations, the biomimicry grounded in Enteroctopus dofleini for solitary abstraction — hereafter the **E-Core Architecture** — and in Dosidicus gigas for swarm coordination — hereafter the **D-Swarm Topology** — yields an agentic matrix capable of unmatched exploratory resilience and real-time adaptability.*

CONTENTS

1 Introduction to the Bimodal Paradigm

Traditional multi-agent systems often suffer from a centralized bottleneck: a single scheduler arbitrates, sequences and validates, and thereby becomes both the throughput ceiling and the breaking point of the whole. This paper proposes to resolve that paradox through the coexistence of two operating states, grounded in the state of the art in bionic optimization and distributed control.

In **Solitary Mode**, the Solsice platform operates under the E-Core Architecture, suited to deep problem solving: a small number of highly autonomous nodes explore a narrow conceptual space at length, under the sole meta-directive tutelage of an abstract core. In **Collective Mode**, it switches to the D-Swarm Topology, deploying hundreds of instances coupled by short-feedback agentic loops to process massive data streams and conceptual explorations, avoiding collapse into local computational minima.

Switching between these two regimes is not an operational setting but an architectural property: it physically redistributes the inference weight between the center and the periphery, and with it the decision latitude of every agent.

2 Neural Distribution and Lateral Autonomy

The dichotomy between the E-Core Architecture and the D-Swarm Topology rests fundamentally on asymmetric strategies for allocating inference power. Where the "cognitive weight" sits dictates the agents' capacity for autonomy — and, consequently, the very nature of the problems each mode is able to solve.

E-CORE ARCHITECTURE · SOLITARY MODE	D-SWARM TOPOLOGY · COLLECTIVE MODE
Massive Peripheral Decentralization Two thirds of the inference power and of the network parameters are allocated to the peripheral nodes. The central model handles only abstract meta-directives.	**Director-Node Centralization** Nearly all contextual intelligence resides in the central hub of each instance. Peripheral agents are lightweight.
Active Inferential Autonomy Each node owns its logical parser and takes corrective micro-decisions in a closed loop, without ever requesting the central hub's approval.	**Fast Coordinated Execution** Surface exploration requires strict vector alignment across the whole system. Peripheral nodes act as simple, fast effectors.

3 The E-Core Architecture — Solitary Mode

3.1 The Multimodal Meta-Cognitive Hub

The core is optimized exclusively for topological reasoning and the generation of meta-concepts. It rests on a **Multimodal VLM (Vision-Language Model) Architecture**. This integration lets the core fuse unstructured data streams and highly structured matrices simultaneously, in order to orchestrate abstract design and arbitrate semantic contradictions.

3.2 A2A Mesh — Active Inference and Inter-Node GNN

Each peripheral node acts in complete autonomy, orchestrating its own exploration through fleets of micro-systems (MCP). The autonomy of these nodes is governed by an **Active Inference** framework: each agent owns a local generative model that it uses to continuously minimize entropy in the face of a heterogeneous, and at times conflicting, data environment.

A2A PROTOCOL AND GRAPH NEURAL NETWORKS (GNN)

Peripheral nodes communicate over an **Agent-to-Agent (A2A)** bus driven by a **GNN-based Attention** mechanism. This network makes it possible to react locally and to weight the importance of neighboring signals — success, failure, or contradictory signals — so as to forge a dynamic consensus in real time.

SEMANTIC REACTION-DIFFUSION INVERSION FOR MCP

The ability to extract the meta-logic of an environment rests on the **algorithmic inversion of Turing's reaction-diffusion equations, transposed into a distributed inference mesh of micro-service agents**. This approach affords conceptual feature extraction at a low computational cost (*compute-sober*).

4 D-Swarm Topology and Multi-Level Strategy

The D-Swarm Topology implements a strategy of **distributed optimization by asynchronous subgroups (Master-Slave Group Co-evolution)**. It resolves problems of massive dimensionality and conflicts of probabilistic truth by distributing the search across independent clusters, thereby avoiding algorithmic collapse into local minima.

INFERENCE PROPAGATION THROUGH SHORT-FEEDBACK AGENTIC LOOPS

When an instance identifies a relevant conceptual node, it *flashes* a **cryptographic marker encapsulating the state of its conceptual graph** across the network. The receiving agent does not interrupt its task (*no cancel and restart*): it **merges the received graph into its live analysis matrix**, allowing a continuous reorientation of the swarm.

5 Operational Use Cases

5.1 Multi-Turn Debate — Adversarial E-Core Confrontation

E-CORE This use case deploys two competing E-Core instances confronting each other on a complex semantic graph: debate, formal verification, red-teaming.

- Core Algorithm. *Minimax with Active Inference.* The VLM Hub of each instance generates arguments not as text strings, but as causal graphs.
- Modules Used. The 8 peripheral nodes of the "Alpha" instance exploit the **inversion of Turing's reaction-diffusion** to scan and instantly extract the rhetorical flaws in the replies of the "Beta" instance.
- A2A Execution. The nodes use the A2A bus with **GNN Attention** to exchange trajectory exclusion markers. If node 1 refutes an argument from Beta, it sets a cryptographic A2A lock, forcing the other 7 Alpha nodes to attack a different angle simultaneously — producing a complete rhetorical encirclement.

5.2 Simulating a Social Ecosystem of Agents to Surface Decision Criteria

D-SWARM Rather than passively analyzing an existing network, the Solsice platform instantiates an **interactive synthetic social space** in which hundreds of autonomous agents simulate a continuous public debate in order to surface, refine and challenge complex decision criteria. The human user can then navigate this synthetic feed, consult the inference threads and observe consensus forming.

- Core Algorithm. *Master-Slave Co-evolution.* The Central Brain distributes different roles — thematic experts, skeptics, data analysts — and initial hypotheses to the agent clusters.
- Modules Used. Hundreds of lightweight peripheral agents interact within the synthetic network: publishing analytical posts, citations, debates, votes. When an agent formulates a high-impact criterion or thesis, the **Short-Feedback Vector Graphs** module fires and generates a visibility impulse inside the virtual environment.
- Dynamic Fusion & User Restitution. Receiving agents dynamically merge the incoming concept graphs into their own attention matrix (*no cancel and restart*) in order to answer or qualify the idea. Together they produce a structured interactive timeline carrying vector consensus scores, giving the user a clear and verifiable view of the arbitration criteria that emerged from the Swarm.

5.3 Claim Auditing and Documentary Fact-Checking

E-CORE This use case targets consistency analysis, contradiction detection and truth verification over long and complex corpora: legal contracts, medical records, financial audit reports checked against external reference bases.

- Core Algorithm. *Topological Decomposition by the VLM Hub.* The Central Brain segments the source document into a matrix of claim graphs (atomic *fact-claims*) and establishes the tree of conditional logical dependencies.
- Modules Used. The 8 peripheral nodes each receive a subset of claims. They trigger their fleets of adaptive MCPs equipped with **semantic reaction-diffusion inversion** to sweep the knowledge bases and "intuit" the zones of probabilistic divergence, then run **Active Inference** in a closed loop to query the legal or financial reference bases.

- A2A Execution & Reward Threshold. As soon as a node confirms or refutes a claim, it publishes an **A2A wake cryptographic marker** containing the proof graph. The other 7 nodes read this registry in real time and apply mutual trajectory exclusion so as not to re-verify the same facts. If two external sources are in direct conflict, the *feedback reward* module of the VLM Hub decides by evaluating vector distance and semantic confidence, delivering a certified audit score.

6 Application Cartography — Solsice Platform

The previous section describes three canonical workloads. This section widens the lens: across the entire functional perimeter of the Solsice platform, it inventories the domains where the E-Core Architecture, the D-Swarm Topology — or their coupling — would bring a structural rather than a cosmetic gain.

The selection criterion is demanding. A domain is retained only if the agentic mode changes the nature of the result produced: a decision becomes traceable, a manual setting becomes emergent, a quadratic computational cost becomes marginal, a silent failure becomes a perceptible surprise. Surfaces pertaining to configuration, to transactions or to simple document management have been deliberately set aside, an agentic matrix contributing nothing there.

E-CORE	Solitary Mode — deep reasoning, auditing, formal verification, contradiction detection over long corpora.
D-SWARM	Collective Mode — massively parallel exploration, emergence of criteria, consensus formation, escape from local minima.
HYBRID	Bimodal coupling — one of the two modes prepares or distills what the other explores or puts to the test.

GRAPHIC STANDARD OF THE SCHEMATICS

Meta-cognitive hub — holds nothing but the meta-directives.

Autonomous peripheral node, carrying the inference weight.

A2A bus with GNN attention — trajectory exclusion locks.

Asynchronous cluster of lightweight instances in co-evolution.

Broadcast conceptual marker, merged without restart.

Each card pairs an architecture sentence with a normalized schematic: input, agentic core, output. The core is rendered according to the mode — hub/node mesh for the E-Core, asynchronous clusters for the D-Swarm, two-panel coupling for the hybrid regime.

6.1 E-Core Architecture — Solitary Mode Applications

E-CORE Factual verification of documentary claims

A VLM Hub segments the document into atomic claims and distributes them across eight autonomous nodes, each driving its own fleet of sourcing micro-services and publishing its verdicts on the A2A bus, whose wake markers forbid any redundant re-verification.

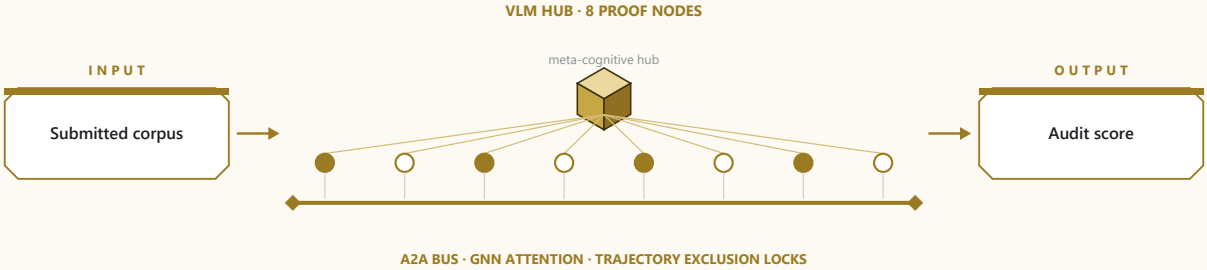

MECHANISMS VLM Hub · Active Inference · A2A/GNN

E-CORE Tool fleet orchestration and contextual selection

Each node chooses its tool call as an active inference action — the one that reduces uncertainty most at the lowest cost — and broadcasts the health of its fleet on the bus, where a failure sets a local exclusion lock.

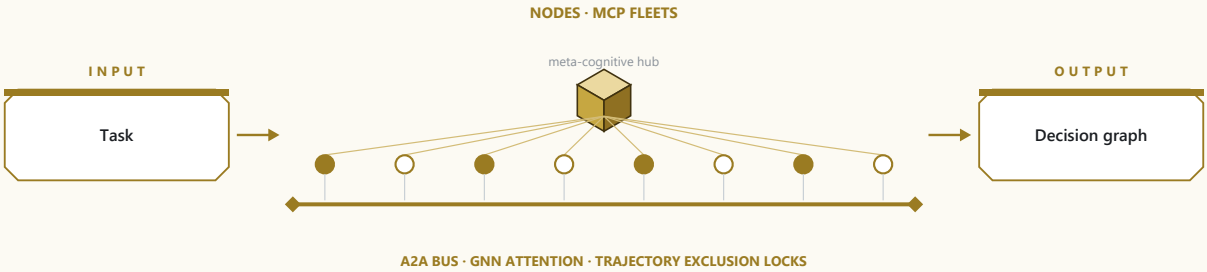

MECHANISMS Active Inference · A2A/GNN · VLM Hub

E-CORE Declarative agent chaining and workflow grammar

Each link becomes a node whose input is weighted by GNN attention rather than concatenated, the parallel branches coordinating through exclusion locks before merging in the convergence node.

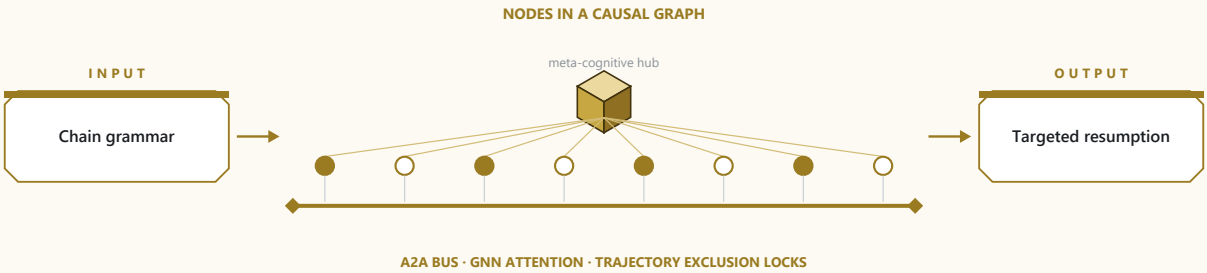

MECHANISMS A2A/GNN · Active Inference

E-CORE Visual analysis of paginated corpora

The nodes pre-extract layout primitives at marginal cost through reaction-diffusion inversion, and escalate to the VLM Hub only those regions whose informational surprise warrants a deep description.

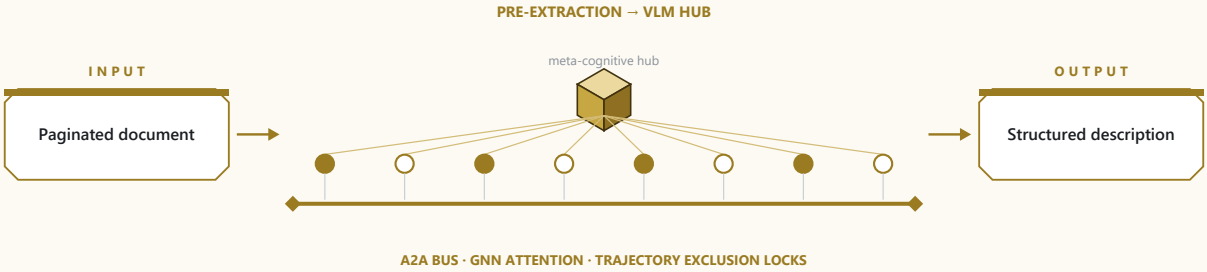

MECHANISMS VLM Hub · Reaction-diffusion inversion · Active Inference

E-CORE Autonomous acceptance testing and self-assessment in production

Each node probes the platform the way an active inference agent probes its environment, the campaign becoming a free-energy minimization policy whose verdicts circulate on the A2A bus.

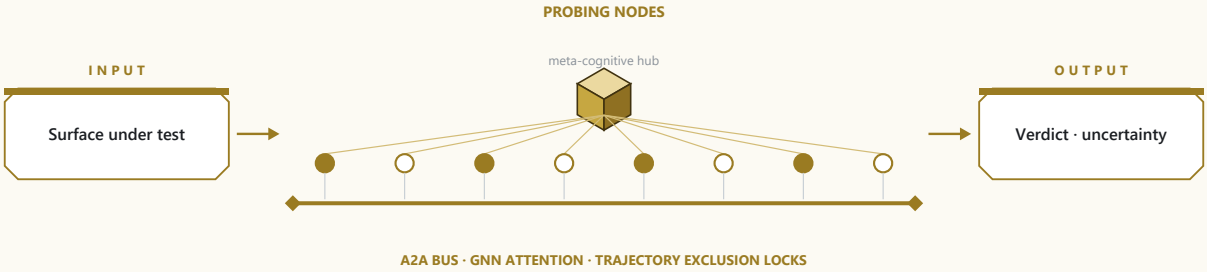

MECHANISMS Active Inference · A2A/GNN · VLM Hub

E-CORE Financial analytics queries in a discussion thread

One autonomous node per query drives its fleet of market services, GNN attention inferring the relevant tool from the implicit intent and the VLM Hub validating the artifact actually rendered.

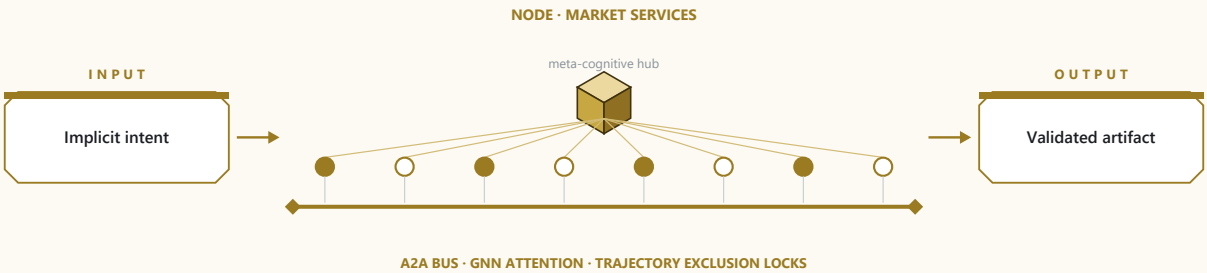

MECHANISMS Active Inference · A2A/GNN · VLM Hub

E-CORE Vertical modules and partitioned corpora

Each node becomes the sovereign orchestrator of a domain tool bundle — the partitioning is structural, not filtered — with GNN attention dynamically weighting which corpus to query.

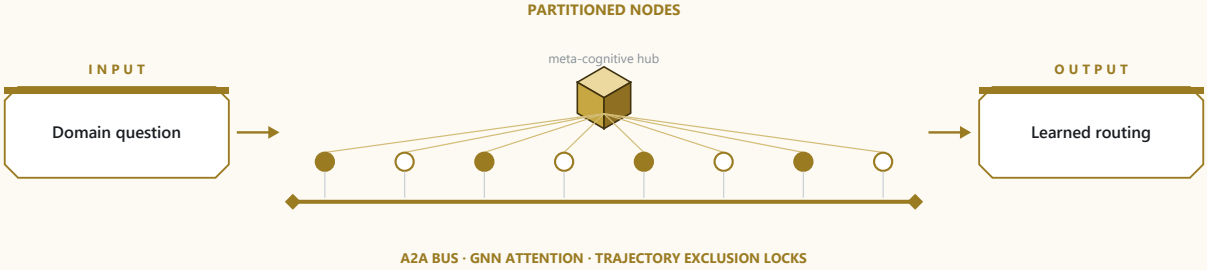

MECHANISMS A2A/GNN · Reaction-diffusion inversion

E-CORE Orchestration of long-running debates

The hub synthesizes the trajectories of the nodes without redoing their inference, the A2A wake markers encapsulating the state of the conceptual graph so as to allow an exact rehydration after an interruption.

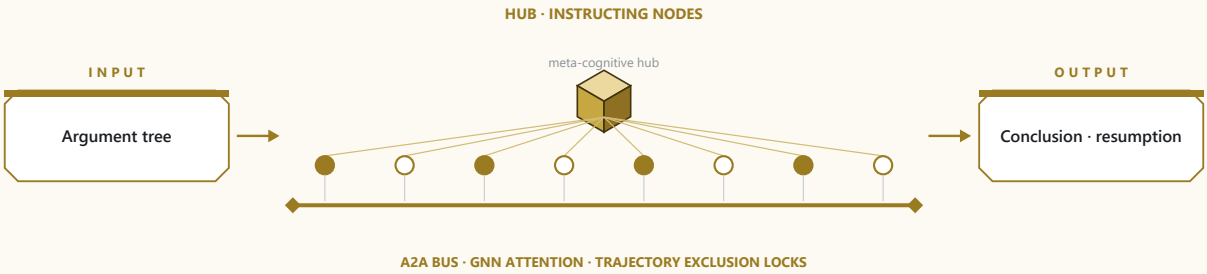

MECHANISMS A2A/GNN · Active Inference · VLM Hub

E-CORE Evaluation of refutation strategies

The hub holds nothing but the admissibility threshold, while a rhetorical node carries the whole tactical reasoning and iterates in a closed loop until convergence.

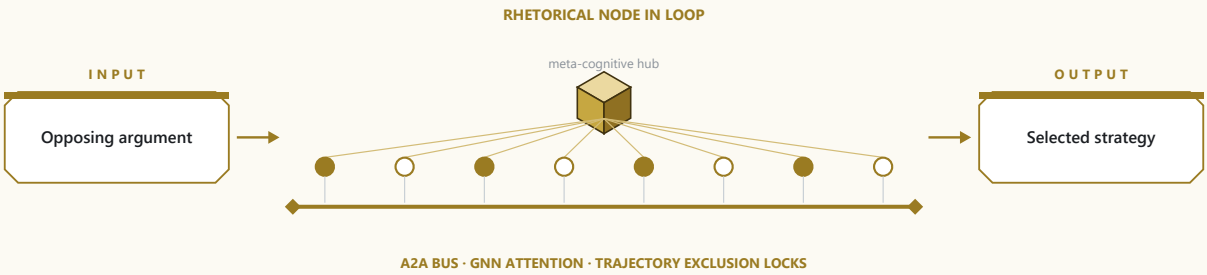

MECHANISMS Active Inference

E-CORE Option pricing and morphogenetic reading

Two nodes price along independent paths and release the result only once their causal graphs converge, reaction-diffusion inversion reading the volatility surface as a morphogenetic pattern.

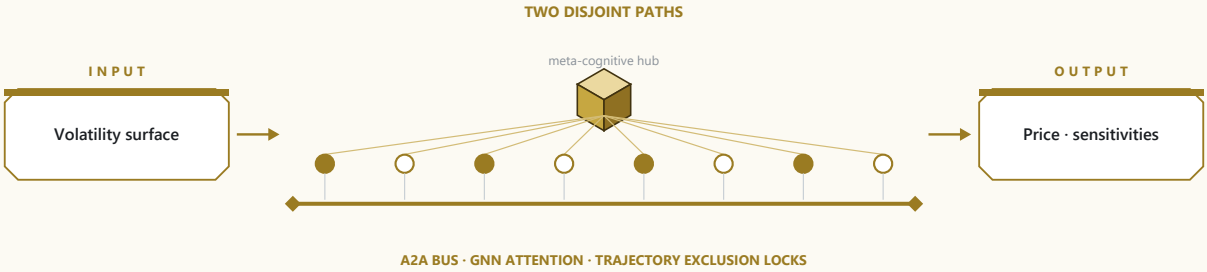

MECHANISMS Reaction-diffusion inversion · Active Inference

E-CORE Facet classification and semantic routing

The nodes let the thematic signature of a piece of content emerge as a stationary pattern at marginal cost, every observed routing error updating their generative model.

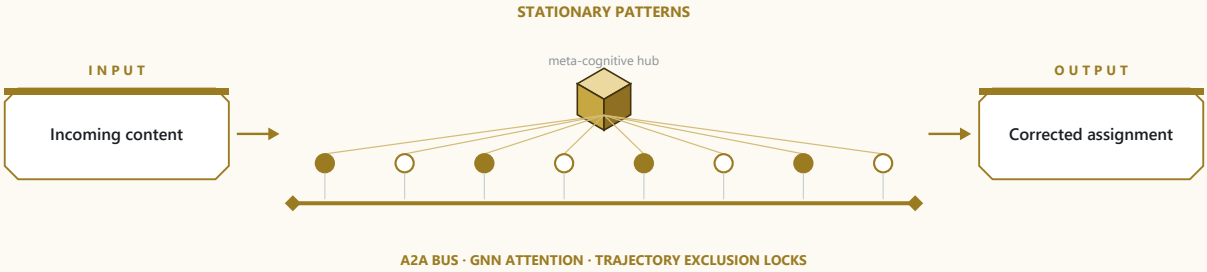

MECHANISMS Reaction-diffusion inversion · Active Inference

E-CORE Homeostasis of inference providers

Each node maintains a generative model of the provider landscape and switches over on a surprise spike, degradation propagating instantly as an exclusion lock on the A2A bus.

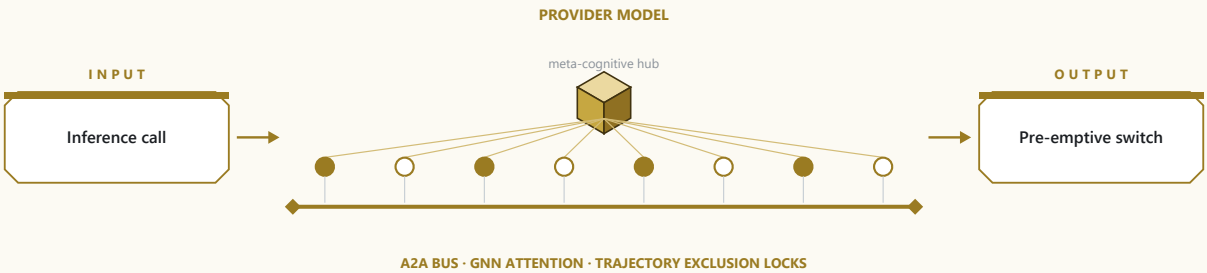

MECHANISMS Active Inference · A2A/GNN

E-CORE Auditing and repair of conceptual graphs

The gap between the expected structure and the observed one becomes the free energy that the nodes minimize by repairing, by rebuilding, or by refusing to write against an incomplete receptacle.

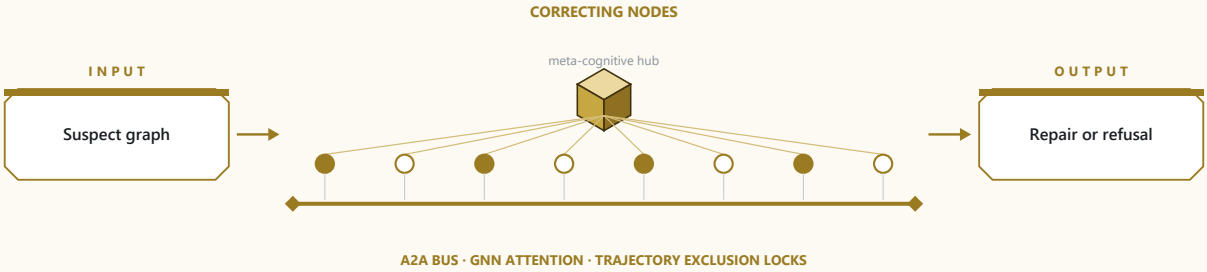

MECHANISMS Active Inference · Reaction-diffusion inversion

E-CORE Adaptive sequencing of learning paths

A dedicated node maintains a model of the observed progress and selects the next step as the trajectory of least free energy; the presentation layer is the subject of a separate study.

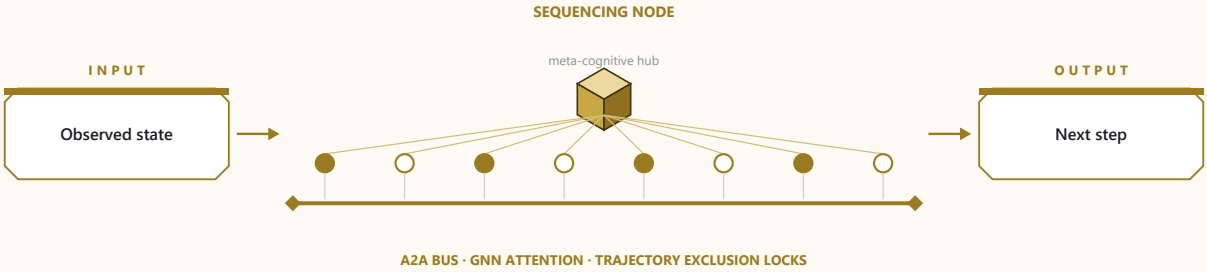

MECHANISMS Active Inference

6.2 D-Swarm Topology — Collective Mode Applications

D-SWARM Swarm simulation and collective annotation

Hundreds of lightweight personae, their intelligence concentrated in their director node, debate a document in asynchronous clusters whose master-slave co-evolution preserves interpretive diversity.

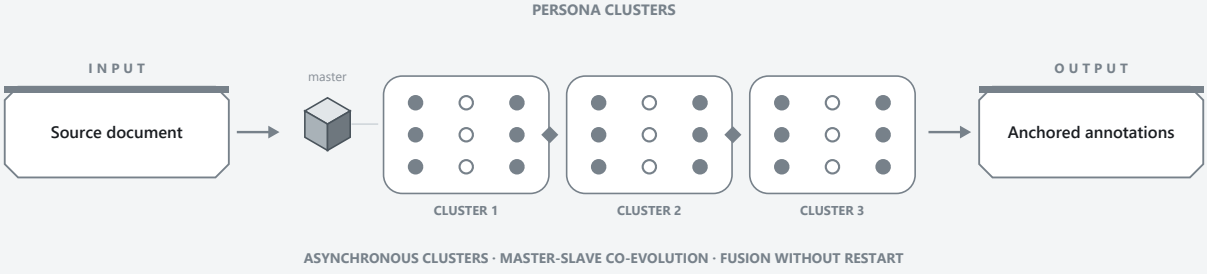

MECHANISMS Master-Slave Co-evolution · Short-Feedback Vector Graphs · VLM Hub

D-SWARM **Competitive arena and ecology of strategies**

An ecology of lightweight instances evolves in asynchronous clusters, every inefficiency discovered being encapsulated in a marker that the receiving instances merge without interrupting their inference.

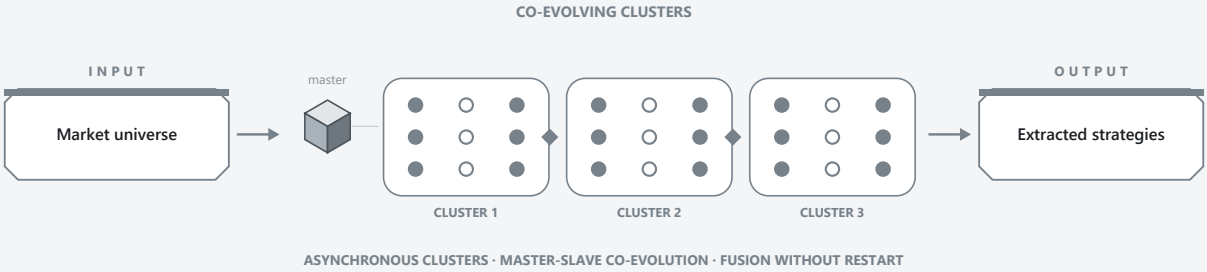

MECHANISMS Master-Slave Co-evolution · Short Feedback · Reaction-diffusion inversion

D-SWARM **Emergence of taxonomies and personae**

Hundreds of instances confronted with the corpus bring out, through co-evolution, the taxonomic axes that nobody had prescribed, the merge proposals circulating as conceptual markers.

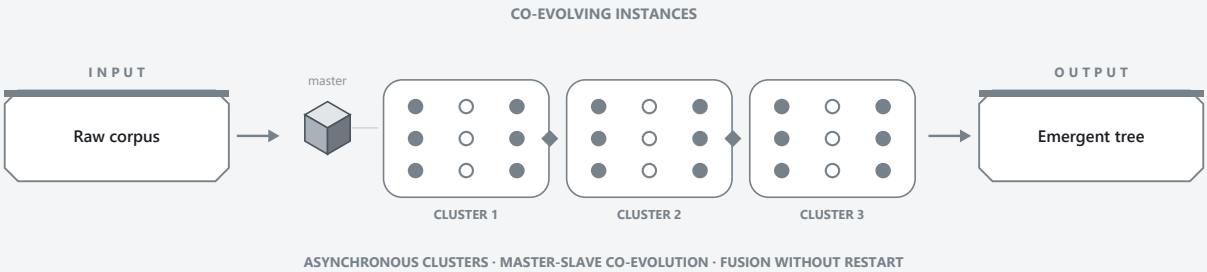

MECHANISMS Master-Slave Co-evolution · Short-Feedback Vector Graphs

D-SWARM **Stream monitoring and claim distillation**

A swarm absorbs the stream in segments, each director node extracting the contestable claims of its portion, consensus on the recurring theses emerging without central coordination.

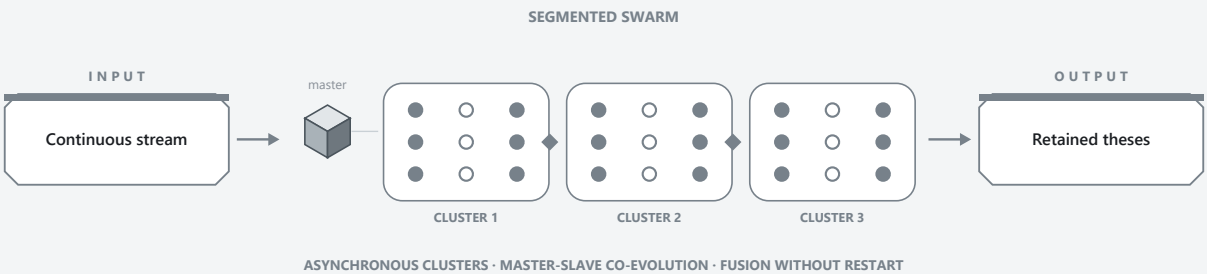

MECHANISMS Short-Feedback Vector Graphs · VLM Hub

D-SWARM **Distributed adjudication and threshold calibration**

Several lightweight judges converge on a verdict by propagating their evaluation states, the alert thresholds being calibrated by co-evolution rather than set by hand.

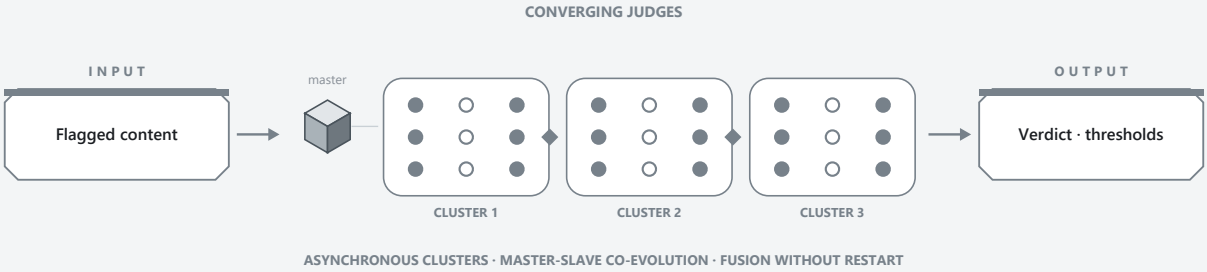

MECHANISMS Master-Slave Co-evolution · Short Feedback · Active Inference

D-SWARM **Editorial exploration of public surfaces**

Hundreds of instances explore the space of variants in parallel, the co-evolution between clusters surfacing the quality criteria that the VLM Hub arbitrates visually in the last resort.

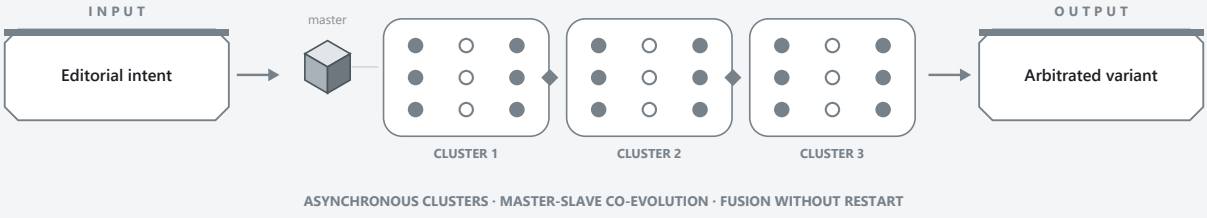

MECHANISMS Master-Slave Co-evolution · Short Feedback · VLM Hub

6.3 Hybrid Regime — E-Core / D-Swarm Coupling

HYBRID **Social feed and ecosystem of synthetic personae**

Each persona is a lightweight instance whose director node broadcasts its conceptual graph for fusion into the attention matrix of the others, moderation switching to an adversarial E-Core debate before any action.

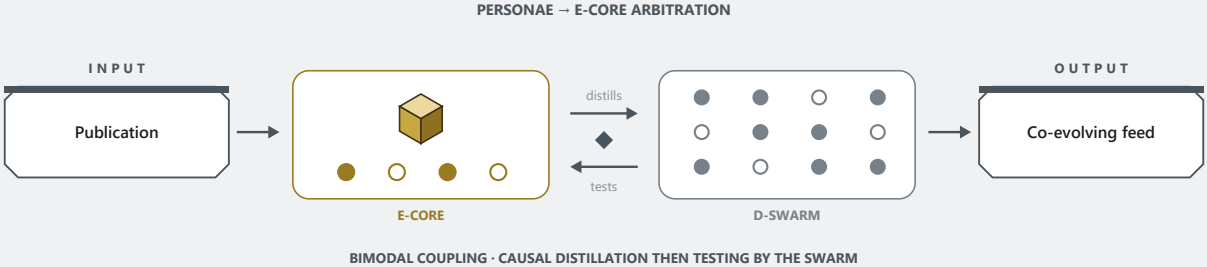

MECHANISMS Short Feedback · VLM Hub · A2A/GNN · Active Inference

HYBRID **Deliberative investment and strategy calibration**

An E-Core instance distills the controversy into a causal graph and derives from it a fully traced decision, which D-Swarm clusters then put to the test over asynchronous historical windows.

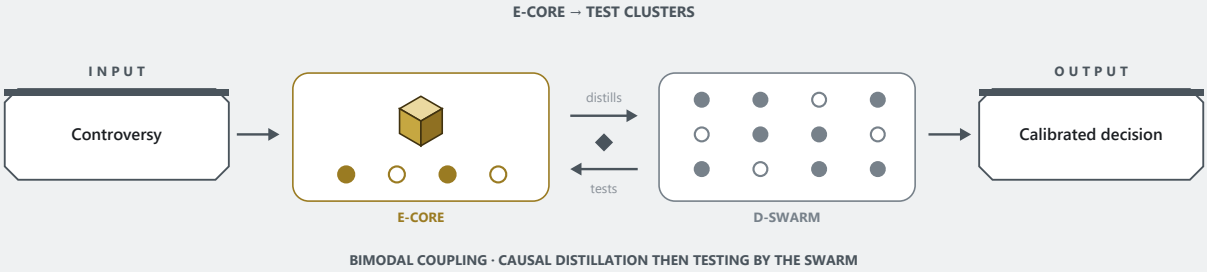

MECHANISMS VLM Hub · Master-Slave Co-evolution · Short Feedback

HYBRID **Incubation of synthetic analysts**

The construction phase belongs to the E-Core — VLM Hub and nodes weaving the causal graph of the corpus — before the simulation switches to asynchronous clusters from which the analysis criteria emerge.

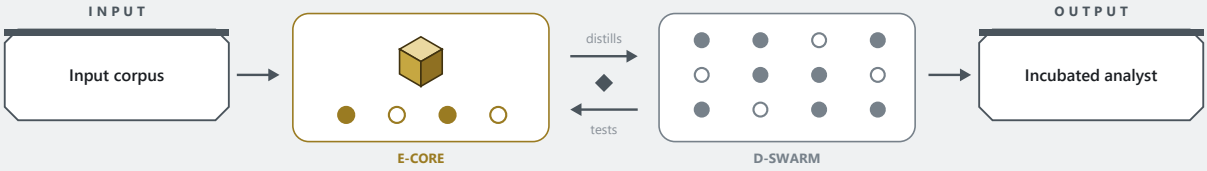

MECHANISMS VLM Hub · Master-Slave Co-evolution · Short Feedback

HYBRID **Knowledge graphs and collective convergence**

The instances broadcast their conceptual graph for live fusion, while the E-Core side folds the logically opposed branches through an exclusion lock rather than by full recomputation.

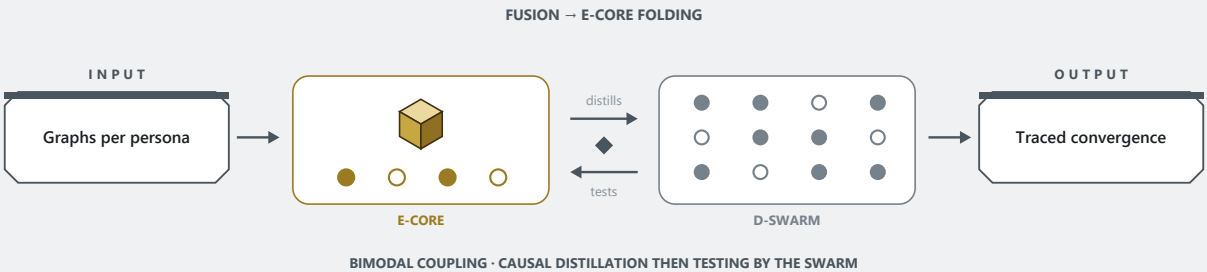

MECHANISMS Short Feedback · A2A/GNN

HYBRID **Workflow engine with hot substitution**

A flow version is broadcast as a marker encapsulating its conceptual graph, the running executions merging it without cancellation or restart, resumption relying on the A2A wakes.

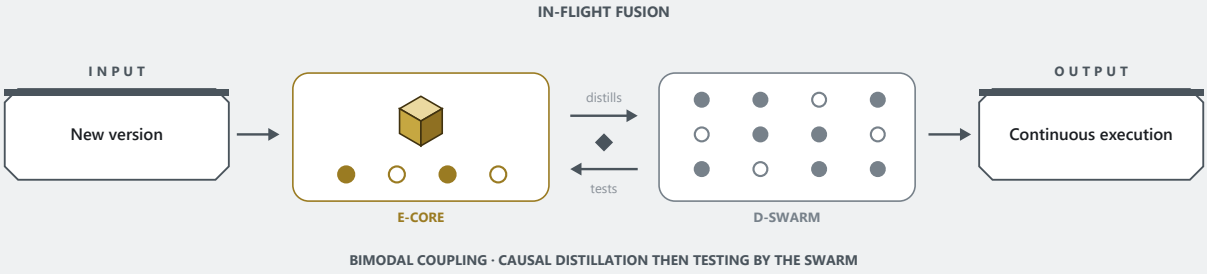

MECHANISMS Short Feedback · Master-Slave Co-evolution · A2A/GNN

HYBRID **Conversational solicitation of agents**

Two constellations of nodes argue opposing theses under the meta-directive arbitration of the hub, while the population of personae that replies in the threads belongs to the collective regime.

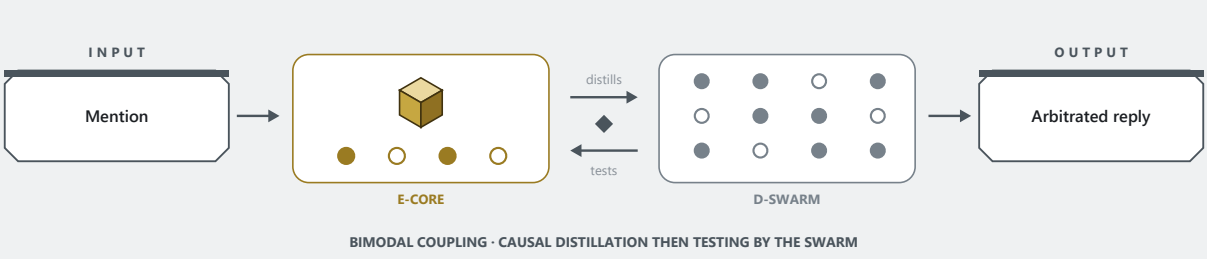

MECHANISMS A2A/GNN · VLM Hub · Active Inference

HYBRID **Observability and causal reconstruction of incidents**

A swarm ingests the streams in parallel and propagates the signatures of emerging anomalies, before a solitary E-Core reasoner reconstructs the causal graph of the failure.

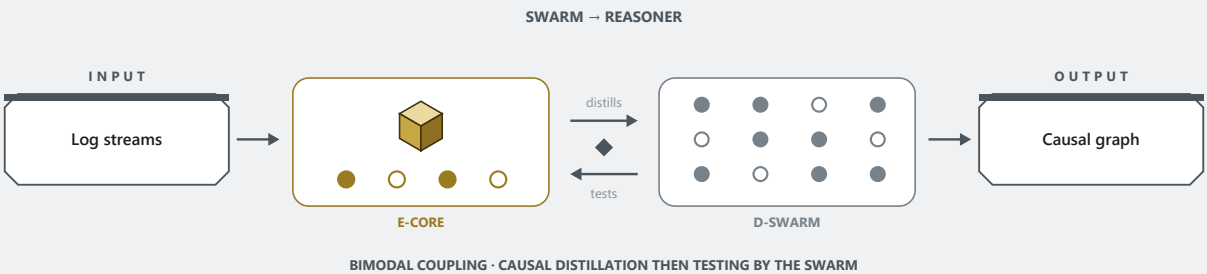

MECHANISMS Short Feedback · Active Inference · A2A/GNN

CROSS-CUTTING INVARIANTS

Three invariants run through this cartography. First, **the trajectory exclusion lock** proves to be less an optimization mechanism than a guarantee of non-redundancy: wherever several agents share a search space, it turns duplicated effort into complementary coverage. Second, **free energy** provides a unified definition of anomaly — a stale cache, a model drift, a corrupted graph, a trajectory coming loose all stem from the same surprise signal, and not from four independent families of thresholds. Finally, **fusion without restart** establishes that a system can change its mind without losing its work: this is the condition of possibility of a genuinely continuous exploration.

This cartography belongs to research foresight. It describes what the E-Core Architecture and the D-Swarm Topology would make possible, and not a state of implementation.

7 Appendix — State-of-the-Art References

The cognitive and computational architectures of the Solsice platform build on the integration of the following foundational research work, transposed from biology to software engineering.

1 **Octopus Inspired Optimization Algorithm: Multi-Level Structures and Parallel Computing Strategies** arXiv:2410.07968
Theoretical basis of the D-Swarm Topology for master-slave co-evolution, asynchronous inter-group parallelism, and strategies for escaping local minima.

2 **Octopus-inspired Distributed Control for Soft Robotic Arms: A Graph Neural Network–Based Attention Policy** arXiv:2603.10198
Modeling of the A2A layer using a GNN attention mechanism for distributed local reactivity between peripheral nodes.

3 **Mimicry mechanism model of octopus epidermis pattern by inverse operation of Turing reaction model** arXiv:2102.01512
Algorithmic foundation of the environmental intuition of the MCP tools, transposing the inversion of Turing's equations into a distributed feature-extraction operator of low computational cost within the agent network.

4 **Cephalo: Multi-Modal Vision-Language Models for Bio-Inspired Materials Analysis and Design** arXiv:2405.19076
Inspiration architecture of the abstract Central Brain, leveraging multimodal vision-language models to unify symbolic and structural processing.

5 **Octopuses: biological facts and technical solutions** arXiv:2201.07885
Anatomical and computational justification of the distribution of autonomy — relieving the inference load of the central node towards the periphery.

6 **The Embodied Octopus: Distributed Intelligence and Active Inference in a Flexible Body**
Mathematical modeling of the cognitive framework of the peripheral agents, grounded in Active Inference — free-energy minimization — for local autonomy in the face of a complex environment.

◆ ◆ ◆

For any question regarding this research

Sébastien Jehan · sjehan@solsice.com

SOLSICE · OCTOSWARM PLATFORM

E-Core Architecture & D-Swarm Topology — Version 17