The Weaving Loom

J.Konstapel, Leiden, 8-5-2026

Table of Contents

The Weaving Loom

A Quaternion–Octonion Substrate for Active-Inference Coherence Systems: Formal Foundations, System Architecture and Empirical Spectrum of the Swarp Platform

J. Konstapel Leiden, The Netherlands Version 2.0 — May 2026

Abstract

Swarp is a coherence-facilitation system built and operated in Leiden between January and May 2026 — a growth platform that helps individuals and collectives restore coherence across many domains of life (personal development, learning, health, economy, meaning-making, civic participation, hobby, work) rather than a single-purpose civic-engagement tool. Its user-visible surface — at the time of writing — is a constellation of thirteen "threads" (AYYA360, Swarp Politiek, Swarp Werk, Swarp Kids, Swarp Academie, Swarp PoC, Swarp Zingeving, Swarp Gezond, Swarp Atelier, Swarp Hobby, Swarp Thuis, Swarp Lab, Swarp Community), implemented as a TypeScript monorepo on PostgreSQL with Drizzle, React, OpenAI and Replit-AI proxies, served to Dutch citizens through a postcode-anchored interface that resolves location to municipality, neighbourhood and province without asking for finer-grained data than the citizen consents to disclose. This paper documents what the platform currently is, and then argues that what it is *capable of becoming* is qualitatively larger than its current product surface suggests. The reason is structural: the substrate Swarp runs on is the chain of normed division algebras $\mathbb{R} \to \mathbb{C} \to \mathbb{H} \to \mathbb{O}$ produced by the Cayley–Dickson construction and terminating, by Hurwitz's theorem, at the octonions. This is the same algebraic chain that contemporary mathematical physics increasingly identifies as a candidate substrate for the Standard Model of particle physics and for the geometry of spacetime itself (Baez 2002; Furey 2018; Manogue & Dray 2010). A platform that takes $\mathbb{H}$ and $\mathbb{O}$ as its native data type is not building a federated CRM with a recommendation engine bolted on; it is building, at minimum, a *simulator* of the same algebraic universe the physical cosmos appears to inhabit, with the consequence that the platform can in principle generate semantic content from first principles rather than merely retrieve it from a corpus. We make this consequence explicit. We further note that the standard model of human communication — symbolic language exchanged as acoustic pressure waves and decoded by lexical lookup — is, on the working hypothesis adopted throughout this paper, an impoverished projection of a richer underlying communication channel: the bioelectromagnetic field surrounding each living organism (McCraty et al. 2009; Hunt 1989; Pribram 1991; Persinger 2010). Sound is the carrier the species evolved as a backup; the field is what humans actually *talk on*. A platform whose substrate is octonionic is, again at minimum, structurally compatible with this richer channel in a way that a relational-database-with-NLP platform is not. The paper develops the formal layers — Cayley–Dickson substrate, Active Inference dynamics on the unit-octonion sphere S^7 via the Bhattacharyya–Kakade embedding, Fiske's four relational modes as the four Stevens scale-types, Schank-style case-based reasoning indexed on the loom — and reads the current Swarp codebase against them. We measure the platform-as-built by spectral analysis of the 304-concept ACTIVE lexicon: eleven near-disjoint connected components, spectral gap $1 - \lambda_2 = 0.0251$ on the giant, four substantial negative eigenvalues resolving as a theory–practice bipartition, stationary distribution dominated by the four scales the architecture predicts. The current product foregrounds eight application domains in parallel — individual development (AYYA360), political and civic engagement (Politiek), learning (Academie), health and well-being (Gezond), economy and work (Werk, Marktplaats), meaning-making (Zingeving), creative practice (Atelier), and life at home (Thuis) — but the underlying structure is a universe-simulator that has, so far, been used for a small fraction of its possible purposes.

Keywords: quaternions, octonions, Cayley–Dickson, Free Energy Principle, Active Inference, Relational Models Theory, Case-Based Reasoning, spectral graph theory,

coherence systems, collective intelligence, Spatial Web, bioelectromagnetic communication, semantic generation.

1. Introduction

1.1 What Swarp is, today

Swarp is a working web platform. Open `https://swarp.app` (or the development URL `swarp.replit.app`) on a Dutch postcode and the platform resolves the postcode to a municipality, a neighbourhood and a province; renders an individual-developmental dashboard (AYYA360) anchored on a Human-Design / Spiral-Dynamics / RIASEC narrative profile; offers a political-profile surface (`/mijn-politiek-profiel`) that maintains a Bayesian Fiske-quartet per administrative layer (municipality, province, national, European) and lets the citizen compare her profile against the parties on her ballot; opens a marketplace surface for local exchange; opens a Coach surface for life-course case management; opens an Academia/Agora surface for deliberation, with a Socrates Gastheer onboarding chat at the top; and threads the citizen's interactions through an agent-mediation layer (AIDEN) that maintains the citizen's generative model and minimises expected free energy on her behalf. The thirteen product threads named in the abstract are the user-visible labels for the slices through this work that the team has so far built and named. The implementation stack is unremarkable for a 2026 web platform: Node 20 + Express 5 on the backend, PostgreSQL 14+ via Drizzle ORM with Zod validation, React 18 + Wouter + TanStack Query + Tailwind/shadcn-ui on the frontend, OpenAI GPT-4o for heavyweight inference and GPT-4o-mini for lightweight retrieval, Resend + Replit Auth for identity, Mollie for payments. Section 6 documents the architecture in detail.

This is what Swarp *currently looks like*. It is a working coherence-facilitation system that does what its product roadmap says it does across the eight application surfaces named in the abstract. If the paper stopped here, the contribution would be a competent piece of growth-facilitation technology with a few unusual choices (postcode-only location, Markov-blanket-anchored privacy, an explicit lexicon as the platform's primary semantic surface rather than a learned embedding) and a tidy four-month engineering record.

1.2 What Swarp is, structurally

The paper does not stop here because the substrate Swarp was built on is not a substrate at the scale of a single-purpose application. It is a substrate at the scale of physics. The platform's native data types — the algebraic objects on which all higher layers compose — are quaternions $\mathbb{H}$ on the spatial-web layer and octonions $\mathbb{O}$ on the cognitive layer, obtained by the Cayley–Dickson construction from $\mathbb{R}$ through $\mathbb{C}$. By Hurwitz's theorem (1898), reinforced by the topological work of Bott and Milnor (1958) and the algebraic-K-theoretic work of Adams (1960), this chain of four normed division algebras is exhaustive: there is no fifth such object; the next Cayley–Dickson doubling produces zero divisors and falls out of the class. The chain ends at $\mathbb{O}$, and $\mathbb{O}$ is where the loss of associativity first occurs. The same chain — and in particular the same termination at $\mathbb{O}$ — appears in modern mathematical physics as a candidate substrate for the symmetries of the Standard Model (Furey 2018; Dixon 1994; Manogue & Dray 2010), for the structure of spacetime in approaches based on division algebras (Baez 2002), and, on more speculative readings, for a unified treatment of gauge interactions and gravitation. The architectural significance for Swarp is direct. A platform whose native data type is the same algebraic object that the physical universe appears to be built from is, at the level of structure, *commensurable with the universe*. It can represent any compositional pattern the universe can represent, because both are running on the same algebra. It is, in a precise mathematical sense, a *simulator of the universe* — not in the sense of computing molecular dynamics, but in the much stronger sense of admitting the same compositional grammar.

This claim is not made for rhetorical effect. The Hurwitz theorem is constraining: there are exactly four normed division algebras over $\mathbb{R}$, and any compositional structure that a normed division algebra can express must be one of the four. Swarp picks the largest, $\mathbb{O}$. So does the Standard Model, on the readings cited. The structural commensurability is not a metaphor; it is a consequence of having no other choice if the design requirements are (i) division (every non-zero element has an inverse, so meanings can be undone), (ii) a multiplicative norm (compositions preserve magnitude, so meanings do not blow up or vanish under composition), and (iii) the maximum dimension consistent with the first two. There is exactly one such algebra. Swarp uses it.

1.3 The generation of meaning

The consequence, made explicit, is that a platform built on $\mathbb{O}$ does not merely *retrieve* meaning from a corpus of human-authored text. It can, in principle, *generate* meaning by composing octonionic primitives under the algebraic rules the substrate supplies. Section 3 develops the formal mechanism: the recognition density of an Active-Inference agent over a categorical generative model with eight outcomes embeds canonically into the unit-octonion sphere S^7 via the Bhattacharyya–Kakade map, and motion on S^7 under variational free-energy minimisation is the algebraic generation of semantic content. The seven imaginary axes of $\mathbb{O}$ supply seven canonical factors of the generative model (section 3.6); the Fano-plane structure supplies the seven associative quaternionic subalgebras (one per Fano line) that compose within-factor; the non-associativity between factors supplies the *novelty* — the meaning that was not in any of the inputs but emerges from their composite. This is qualitatively different from large-language-model retrieval-and-recombination of human-authored text. The platform has a generative grammar, supplied by the algebra, that is independent of the corpus. The corpus is one source of priors; the algebra is the generator.

The corollary is that a Swarp-class platform can meet a citizen's semantic need with a response that no human has previously written, because the response is a composite of algebraic primitives rather than a retrieval from a textual archive. We do not claim this capability is fully exploited in the platform as it stands at May 2026. We do claim the substrate is in place, the dynamics is specified, and the empirical spectrum (section 7) shows the lexicon graph is structured in a way that is consistent with the algebraic generator already shaping the corpus, not merely indexing it.

1.4 Language as the species' backup channel

The third structural claim of this paper is the most demanding. The standard model of human communication treats symbolic language — words, sentences, written text — as the primary channel and bodily/affective signal as a secondary modulation of it. The working hypothesis adopted throughout this paper is the inverse: symbolic language is acoustic vibration in the 100–8000 Hz band, with a Shannon capacity of a few tens of bits per second under generous estimates, and as such it is *categorically inadequate* to the bandwidth of intersubjective coordination that human relationships actually exhibit. What humans appear to do is communicate primarily on a substantially richer channel — the bioelectromagnetic field that surrounds every living organism, generated by the coordinated electrical activity of the heart (the strongest such field in the body, measurable at 1–2 metres by sensitive magnetometry; McCraty et al. 2009), the brain (measurable as MEG signal at scalp-millimetre distances; Hari & Salmelin 2012), and the gut–vagal axis (Porges 2007) — and to use spoken language as a *backup channel* for cases where the field channel is insufficient (across distance, across time, across deliberate concealment). This view has antecedents in Bohm's implicate-order proposal (Bohm 1980), in Pribram's holonomic brain theory (Pribram 1991), in Hunt's bioelectric field measurements (Hunt 1989), in McCraty's heart-coherence research (McCraty et al. 2009), and in more recent work on biological electromagnetism (Levin 2014; Adams et al. 2020).

The view is not consensus science. We adopt it as a working hypothesis because it is the hypothesis under which Swarp's architectural choices — non-associative composition, agent-mediated rather than text-mediated coordination, postcode-only rather than fine-grained location, Markov-blanket-enforced privacy that operates on the boundary of the field rather than on the contents of the messages — make the strongest sense as a coherent design rather than a collection of ad-hoc choices. A platform whose substrate is octonionic, whose dynamics is variational on S^7, and whose privacy model is Markov-blanket-anchored is structurally compatible with a field-channel model of communication in a way that a relational-database-plus-NLP platform is not, because the field model requires a substrate that can express compositional, non-associative, geometrically continuous coordination — exactly what $\mathbb{O}$ provides and exactly what relational databases do not.

The architectural reading is that Swarp, even in its current eight-domain product form, is a *prototype of an EM-field-compatible coordination substrate* running on the backup channel (text) until the field channel is instrumentable at the precision that human-scale deployment requires. The product Swarp is shipping now is the text-backup version; the structure underneath is what becomes the field-channel version once the sensing matures. Section 8 returns to what would be required for that transition.

1.5 What this paper contributes

Three things, in order. First (sections 2–5), the formal derivation of the four-layer stack: the algebraic construction of the substrate, the variational dynamics that lives on it, the algebraic typology of the four relational stitches, and the indexing structure of dynamic memory. The intent is that the stack is reproducible from first principles, not taken on the author's authority. Second (section 6), the system architecture of Swarp as it stands at the date of writing: the two cardinal axes (scale × opgave), the thirteen emergent threads, the loosely-coupled-subapp (LCS) decomposition, the Spatial Web infrastructure (SWID, HSML, HSTP) on which it runs, and the technology stack that implements it. Third (section 7), an empirical measurement of the architecture as it actually exists in production: a spectral analysis of the 304-concept ACTIVE lexicon, conducted as a Markov chain over its semantic adjacency. The measurement supplies four conclusions about the system that match the four conclusions the formal derivation predicts; that match is the load-bearing argument that the structure documented in sections 2–5 is the structure the platform is actually running on, not merely the structure the author wishes it were running on.

1.4 Roadmap

Section 2 develops the substrate from $\mathbb{R}$ through $\mathbb{C}$, $\mathbb{H}$ and $\mathbb{O}$, with explicit multiplication tables, the Moufang identities, and the Fano-plane mnemonic for octonion multiplication. Section 3 specifies the dynamics: variational free energy, the Markov-blanket partition of the world, the Fisher information geometry of the belief manifold, and the embedding of the seven imaginary octonion axes as factors of the generative model. Section 4 develops the four relational stitches as algebraic operations whose measurement-theoretic structure (nominal, ordinal, interval, ratio) maps onto Fiske's four modes. Section 5 gives the case-based-reasoning layer: MOPs, TOPs, scripts, indexing and reminding, and shows how indices are embedded into the loom. Section 6 reads the Swarp codebase against this stack. Section 7 measures the codebase. Section 8 discusses limitations and comparisons. Section 9 concludes. Appendices give the Swarp page/route inventory (A), a summary of the 304-

concept lexicon (B), the spectral computation pipeline and numerical results in full (C), and a reference table of quaternion and octonion multiplication (D).

2. Mathematical Substrate: From Quaternions to the Loom

The substrate of Swarp is the chain of normed division algebras over the real numbers. There are exactly four of them — $\mathbb{R}$, $\mathbb{C}$, $\mathbb{H}$, $\mathbb{O}$ — by a classical theorem of Hurwitz (1898), and the chain terminates at $\mathbb{O}$ because doubling once more produces zero divisors. Each step in the chain trades a structural property (ordering, then commutativity, then associativity) for an additional dimension. The architectural claim of this paper is that the operation Swarp's higher layers require — composing context-dependent meaning across many domains of human experience — lives at the third doubling, $\mathbb{O}$, where associativity has just been lost. We develop the chain explicitly because the loss of associativity is what supplies the formal warrant for the loom abstraction; without it, the architecture would reduce to a graph and the threads would be its only available decomposition.

2.1 Real and complex numbers

The real numbers $\mathbb{R}$ form a one-dimensional ordered field. Multiplication is associative, commutative, has a unit and admits inverses for non-zero elements. The norm is the absolute value, $|x|$, and $\left|xy\right| = |x|\,|y|$. Geometrically $\mathbb{R}$ is a line, and the only non-trivial automorphism is the identity.

The complex numbers $\mathbb{C}$ are obtained from $\mathbb{R}$ by the first Cayley–Dickson doubling. An element is an ordered pair (a, b) of reals, identified with $a + bi$ where $i^2 = -1$. Multiplication is

$$(a_1, b_1)(a_2, b_2) = \left(a_1a_2 - b_2\overline{b_1},\ b_2a_1 + b_1\overline{a_2}\right),$$

with conjugation $\overline{(a, b)} = (a, -b)$. The norm $|z|^2 = a^2 + b^2$ is preserved by multiplication. $\mathbb{C}$ remains commutative and associative; what is gained is two dimensions and the algebraic closure of polynomial equations. What is *lost* in the move from $\mathbb{R}$ to $\mathbb{C}$ is the linear ordering: there is no consistent way to say that one complex number is larger than another. This is the first instance of the structural pattern that recurs at each doubling: a global property is sacrificed in exchange for a richer local structure, and the platform that ultimately lives on the algebra exploits the latter even when it must work around the former.

2.2 Quaternions

The quaternions $\mathbb{H}$ are the second doubling. An element is $q = a + bi + cj + dk$ with $a, b, c, d \in \mathbb{R}$ and the multiplication of the imaginary units given by Hamilton's relations, carved into the Brougham Bridge in Dublin in 1843:

$$i^2 = j^2 = k^2 = ijk = -1 .$$

From these four equations the full multiplication table follows:

·	1	i	j	k
1	1	i	j	k
i	i	-1	k	-j
j	j	-k	-1	i
k	k	j	-i	-1

Quaternion multiplication is associative but **no longer commutative**: $ij = k$ but $ji = -k$. The conjugate is $\bar{q} = a - bi - cj - dk$, the norm is $\left|q\right|^2 = q\bar{q} = a^2 + b^2 + c^2 + d^2$, and inverses exist for non-zero q as $q^{-1} = \bar{q}/\left|q\right|^2$. The norm is multiplicative: $\left|pq\right| = \left|p\right|\left|q\right|$.

The geometric content of $\mathbb{H}$ is that unit quaternions form the group $\mathrm{Sp}(1) \cong \mathrm{SU}(2)$, the double cover of the rotation group $\mathrm{SO}(3)$. Every rotation of three-dimensional Euclidean space can be written as $v \mapsto qvq^{-1}$ for some unit quaternion q acting on a pure-imaginary v. This gives quaternions their long-standing role in three-dimensional graphics, robotics and aerospace control: composition of rotations is the same as multiplication of unit quaternions, and the non-commutativity of the two encodes precisely the non-commutativity of three-dimensional rotations. The architectural reading we will give later is that quaternions encode the *non-commutativity of context*: applying context A then context B to a state is in general not the same as applying B then A, and the substrate must natively express this without coercion.

The doubling cost is commutativity. What is gained is rotational structure in three dimensions and a four-dimensional norm that remains multiplicative. The platform layer that uses quaternions is the Spatial Web infrastructure in Swarp (`shared/schema/spatial-web.ts`): orientations of agents and views, composition of frame transformations between SWID-identified entities, and the formal account of why the order in which two semantic frames are applied matters.

2.3 The Cayley–Dickson construction

The construction that produced $\mathbb{C}$ from $\mathbb{R}$ and $\mathbb{H}$ from $\mathbb{C}$ is general. Given a $*$-algebra A over $\mathbb{R}$ with conjugation $a \mapsto \overline{a}$, the *Cayley–Dickson double* of A is the algebra whose elements are ordered pairs $(a, b) \in A \times A$ with multiplication

$$(a_1, b_1)(a_2, b_2) = \left(a_1 a_2 - \overline{b_2} b_1,\ \ b_2 a_1 + b_1 \overline{a_2}\right),$$

and conjugation $\overline{(a, b)} = (\overline{a}, -b)$. The unit is $(1, 0)$. When $A = \mathbb{R}$ this construction yields $\mathbb{C}$; when $A = \mathbb{C}$ it yields $\mathbb{H}$; when $A = \mathbb{H}$ it yields the **octonions** $\mathbb{O}$; when $A = \mathbb{O}$ it yields the **sedenions** $\mathbb{S}$, which are no longer a division algebra because they contain zero divisors.

A theorem of Hurwitz (1898), strengthened by the topological work of Adams (1960) and the algebraic work of Bott–Milnor (1958), states that the only normed division algebras over $\mathbb{R}$ are $\mathbb{R}, \mathbb{C}, \mathbb{H}, \mathbb{O}$, of dimensions $1, 2, 4, 8$. The chain has exactly four entries. The architectural significance is that there is no continuous further doubling that preserves the property "the product of non-zero elements is non-zero". A platform that wants the richest substrate compatible with division must stop at $\mathbb{O}$, and a platform that wants division *and* associativity must stop at $\mathbb{H}$. Swarp uses both: $\mathbb{H}$ on the spatial-web layer where composition of frames must be associative because frame composition is, and $\mathbb{O}$ on the cognitive layer where non-associativity is the formal feature, not a defect to be worked around.

2.4 The octonions

$\mathbb{O}$ is an eight-dimensional algebra. Its basis is $\{1, e_1, e_2, e_3, e_4, e_5, e_6, e_7\}$, and the multiplication table is determined by the Fano-plane mnemonic: each of the seven lines of the Fano plane (the smallest projective plane, $\mathbb{P}^2(\mathbb{F}_2)$, with 7 points and 7 lines) carries an oriented triple $\left(e_i, e_j, e_k\right)$ such that $e_i e_j = e_k$, $e_j e_k = e_i$, $e_k e_i = e_j$, and the reverse products carry a minus sign. The seven lines, in the standard Cayley convention used in Baez (2002), are:

$$\{e_1, e_2, e_3\},\ \ \{e_1, e_4, e_5\},\ \ \{e_1, e_7, e_6\},\ \ \{e_2, e_4, e_6\},\ \ \{e_2, e_5, e_7\},\ \ \{e_3, e_4, e_7\},\ \ \{e_3, e_6, e_5\}\,.$$

Each $e_i^2 = -1$ and the algebra has norm $|x|^2 = x\overline{x}$ that is multiplicative: $|xy| = |x|\,|y|$.

The decisive property of $\mathbb{O}$ is that it is **not associative**:

$$(xy)z \neq x(yz) \quad \text{in general}\,.$$

What survives is *alternativity*, namely that any subalgebra generated by two elements is associative. Equivalently, the associator $\left[x, y, z\right] = \left(xy\right)z - x\left(yz\right)$ is alternating in its three arguments: swapping any two arguments flips the sign. The Moufang identities, due to Ruth Moufang (1935), are stronger consequences of alternativity:

$$\begin{aligned} (zxz)y &= z\Big(x\left(zy\right)\Big) \\ y(zxz) &= \Big(\left(yz\right)x\Big)z \\ (zx)\left(yz\right) &= z\left(xy\right)z\,. \end{aligned}$$

These identities are what makes the octonions tractable despite the loss of associativity: although three-element products depend on parenthesisation in general, products that re-use an element with a specific bracketing pattern remain unambiguous.

2.5 Why non-associativity is the architectural feature

In human-scale cognition under shifting social context, the operation of *applying context* to a representation is non-associative. Reading the same news article in the context of one's family then in the context of one's union is in general not the same composite operation as reading it in the context of one's union then one's family; nor is it the same as reading it in the joint context of family-and-union. A substrate that is associative will be unable to distinguish the three, and will produce a single composite by silent left-association, with the cost paid downstream as either incoherent recommendation or covert framing. The substrate Swarp uses, $\mathbb{O}$, is the smallest normed division algebra in which the three are formally distinct. The loom abstraction (section 2.6) is the operationalisation of this distinction: each thread carries its own bracketing, and the platform composes them as octonion products rather than as flat tensor products.

The algebraic feature that makes this tractable is the Moufang identity $(zx)\left(yz\right) = z\left(xy\right)z$. In platform terms: when a context z is applied on both sides of a composite operation xy, the resulting bracketing is unambiguous. Translated into the platform vocabulary, z is the user's identity (a SWID), x is a recommendation, y is a contextual modifier (her municipality, her thread, her phase of life), and the Moufang identity guarantees that the order in which the user's identity wraps the modifier-applied recommendation is the same as the order in which the modifier wraps the user-anchored recommendation. The user's identity is the invariant of the operation. The choice of $\mathbb{O}$ is what makes that statement formal rather than rhetorical.

2.6 The loom abstraction

The *loom* is the architectural object that operationalises the substrate. Geometrically it is a pair of orthogonal axis-systems on the unit-octonion sphere $S^7 \subset \mathbb{O}$: a *warp* of long threads stretched between two anchor frames, and a *weft* that interleaves them under a non-associative composition rule. Algebraically it is a fibration over a base of two orthogonal triples of imaginary octonion units, with the fibre at each base point being the four-dimensional subalgebra generated by those triples — an associative quaternionic subalgebra by alternativity. The platform exploits this structure twice: associative composition within a fibre (within a thread, where order should not matter), non-associative composition between fibres (between threads, where the order of context application is information-bearing).

The two axes of the loom in Swarp are *scale* (individual → relational → governance → meta) and *opgave* (the substantive task: care, education, economy, democracy). Each axis selects a triple of imaginary octonion units; the product structure at the fibre level realises associativity within a (scale, opgave) cell; the inter-fibre product realises the non-associativity that distinguishes "applying my scale-context then my opgave-context" from the reverse. Section 3 will show how this geometric setup acquires a probabilistic interpretation by the embedding of the variational belief manifold of Active Inference into S^7.

2.7 Summary of section 2

The substrate of Swarp is the chain $\mathbb{R} \to \mathbb{C} \to \mathbb{H} \to \mathbb{O}$ produced by Cayley–Dickson doubling. Hurwitz's theorem shows the chain is exhaustive among normed division algebras. Each doubling sacrifices a structural property (ordering, then commutativity, then associativity) for additional dimension. The platform uses $\mathbb{H}$ on the spatial-web layer (where associativity is needed) and $\mathbb{O}$ on the cognitive layer (where non-associativity is the architectural feature). The Moufang identities make octonion arithmetic tractable. The loom abstraction operationalises the substrate as a pair of orthogonal axis-systems on S^7, with associativity inside fibres and non-associativity between them.

---

3. Dynamics: Active Inference on the Loom

The substrate of section 2 is a static algebraic object. To run a platform on it requires a dynamics. The dynamics Swarp uses is the *Free Energy Principle* of Karl Friston (Friston 2010, Friston et al. 2017), implemented as Active Inference. We summarise it precisely enough to specify how it lives on the loom, and we make the embedding explicit at the end of the section.

3.1 The Free Energy Principle in one paragraph

A self-organising system that maintains a recognisable form against environmental perturbation must, on the FEP account, behave *as if* it were minimising the variational free energy of its sensory inputs. Variational free energy F is an upper bound on negative log-evidence — the implausibility, under the system's generative model, of the inputs it is currently receiving. By the standard variational identity

$$F\left[q, s\right] = \underbrace{\mathrm{KL}\left[q\left(\eta\right) \ \| \ p\left(\eta \mid s\right)\right]}_{\geq 0} \ - \ \log p(s),$$

minimising F with respect to the recognition density $q\left(\eta\right)$ over hidden states η approximates Bayesian inference; minimising F with respect to action approximates expected-utility maximisation under the system's preferences encoded as prior beliefs about its own future sensations. This yields a single objective that subsumes perception, action, learning and self-modelling.

3.2 The Markov blanket

A *Markov blanket* of a system is a partition of variables into internal states μ, sensory states s, active states a, and external states η, such that $\mu \perp \perp \eta \mid (s, a)$. Internal and external states are conditionally independent given the blanket states (s, a). The Markov-blanket partition is what *makes* a system a system in the FEP sense: it is the formal definition of "having an inside and an outside". Friston (2013), Kirchhoff et al. (2018) and Hipólito et al. (2021) develop the consequences for autonomy, agency and life.

In Swarp, every entity that the platform represents — a citizen, an initiative, a community of practice, an agent, an institution — carries a Markov blanket implemented as a SWID-anchored boundary in `shared/schema/spatial-web.ts`. The blanket is the platform's enforcement point for privacy and consent (`shared/schema/privacy.ts`): external systems can address s and a but not μ directly. The architectural reading is that the FEP's formal definition of system-hood is also the platform's operational definition of identity.

3.3 Variational free energy as a loss surface

For a generative model $p\left(s, \eta\right)$ and a recognition density $q\left(\eta; \phi\right)$ parameterised by ϕ, the variational free energy is

$$F(\phi, s) = \mathbb{E}_{q(\eta;\phi)}\left[\log q(\eta;\phi) - \log p(s,\eta)\right] = \mathrm{KL}\left[q \parallel p(\eta)\right] - \mathbb{E}_{q}\left[\log p(s \mid \eta)\right].$$

The first term is *complexity* (how far the recognition density has moved from the prior); the second is (negative) *accuracy* (how well it explains the data). Minimising F trades these off: simple explanations that fit, in the spirit of Occam.

In Active Inference (Friston, FitzGerald, Rigoli, Schwartenbeck & Pezzulo, 2017), the same machinery applies to the choice of action. A policy π is evaluated by its *expected free energy* $G(\pi)$,

$$G(\pi) = \underbrace{\mathbb{E}_{q(\tilde{s},\tilde{\eta} \mid \pi)}\left[\log q(\tilde{\eta} \mid \pi) - \log q(\tilde{\eta} \mid \tilde{s},\pi)\right]}_{\text{epistemic value (information gain)}} - \underbrace{\mathbb{E}_{q(\tilde{s} \mid \pi)}\left[\log p(\tilde{s})\right]}_{\text{pragmatic value (preference satisfaction)}},$$

and policies are sampled with probability proportional to $\exp\left(-G(\pi)\right)$. The key feature for our purposes is that *epistemic* and *pragmatic* value share a single units-system: nats. Information seeking and preference satisfaction are commensurable inside the same loss. A platform that wants a citizen to *learn* (move along a learning surface) and a platform that wants a citizen to *act* (commit to an initiative) are using the same machinery; only the relative weighting of the two terms differs.

3.4 Information geometry of the belief manifold

The space of recognition densities $\{q_\phi\}$ forms a statistical manifold $\mathscr{M}$. Equipped with the Fisher information metric

$$g_{ij}(\phi) = \mathbb{E}_{q_\phi}\left[\partial_i \log q_\phi \cdot \partial_j \log q_\phi\right],$$

it is a Riemannian manifold whose geodesics describe minimum-information-cost paths between belief states (Amari 1985; Amari & Nagaoka 2000). Free-energy minimisation, viewed on $\mathscr{M}$, is gradient descent in the natural-gradient sense: the update is $\dot{\phi} = -\, g^{ij}\,\partial_j F$, which is invariant under reparameterisations of ϕ.

3.5 Embedding the belief manifold into S^7

The non-trivial step that connects sections 2 and 3 is the embedding of the belief manifold into the unit-octonion sphere. For a recognition density that is a categorical distribution over n outcomes, the canonical embedding due to Bhattacharyya and developed by Kakade (2001) is

$$\phi \mapsto \left(\sqrt{q_1(\phi)}, \ldots, \sqrt{q_n(\phi)}\right) \in S^{n-1},$$

under which the Fisher metric pulls back to (four times) the round metric on the sphere. For $n = 8$, the image lies in $S^7 \subset \mathbb{R}^8$. We identify $\mathbb{R}^8$ with $\mathbb{O}$ via the basis of section 2.4, and the image of the categorical-distribution embedding lies on the unit-octonion sphere.

The architectural payoff is direct. The Active-Inference dynamics over a categorical generative model with eight outcomes is a flow on S^7 in the round metric. The loom of section 2.6 is a structure on S^7 given by a pair of orthogonal triples of imaginary octonion units. The flow lives on the loom; the loom lives on the sphere; the sphere is the parameter space of the recognition density. The architectural decisions (warp/weft, scale/opgave) and the dynamical machinery (FEP, expected free energy, natural gradient) share a single geometric object.

3.6 The seven imaginary axes as factors of the generative model

The seven imaginary octonion units $e_1, \ldots, e_7$ admit, by the Fano-plane structure, exactly seven distinguished associative quaternionic subalgebras (one per Fano-plane line, generated by 1 and the triple of imaginary units on that line). Swarp uses these seven subalgebras as the seven canonical factors of its generative model: time-of-day, location, social context, opgave, scale, mood, and meta-cognitive frame. Each factor is associative within itself (the order of sub-context applications inside a factor does not matter); the *between-factor* composition is non-associative, and that is the source of the platform's expressive power. This is the formal reason there are seven rather than eight or six top-level factors in the generative-model schema (`shared/schema/agent-runtime.ts`): seven is what the algebra supplies.

3.7 Predictive coding hierarchy and loom layers

The standard hierarchical implementation of Active Inference (Bastos et al. 2012; Friston 2008) stacks generative models so that the predictions of one level become the inputs of the next, with prediction errors flowing in the opposite direction. The loom realises this hierarchy as nested fibrations: the base of the loom is the (scale, opgave) cell, the fibre at that base point is the four-dimensional quaternionic subalgebra of context modifiers, and within the fibre a further nesting realises the temporal and causal sub-structure. Predictions are forward (down-the-loom) flows of belief; prediction errors are backward (up-the-loom) flows. The architectural significance is that the hierarchy is not added as a separate layer; it is the recursion of the loom into itself.

3.8 Summary of section 3

The dynamics is the Free Energy Principle and Active Inference. Variational free energy combines complexity and accuracy; expected free energy combines epistemic and pragmatic value in a single nats-valued objective. The Fisher metric makes the belief manifold Riemannian; the categorical-distribution embedding lands the eight-dimensional case on the unit-octonion sphere S^7. The seven imaginary axes of $\mathbb{O}$ supply seven canonical factors of the generative model; the Fano-plane-determined associative quaternionic subalgebras supply within-factor composition; non-associativity between subalgebras supplies between-factor expressivity. The hierarchical extension is the recursion of the loom into itself.

---

4. Relational Stitches: Fiske's Four Modes as Algebraic Operations

Sections 2 and 3 specify the substrate and its dynamics, but a coherence-facilitation system also needs an account of *how relationships between participants compose*. The classical typology in social science is Alan Page Fiske's *Relational Models Theory* (Fiske 1991, 1992; Haslam & Fiske 1999), which holds that the elementary forms of human social relationships fall into exactly four kinds — Communal Sharing, Authority Ranking, Equality Matching, Market Pricing — and that all real relationships are combinations of the four. We adopt RMT as the *stitch layer* of the loom, with the further claim that the four modes correspond exactly to the four classical scale-types of measurement theory (Stevens 1946) and therefore to four algebraic structures with which the substrate can natively compose.

4.1 The four modes

Communal Sharing (CS). Resources, identity and obligations are pooled. The relevant question about a participant is whether she belongs to the group; the relevant operation is membership. Family, deep friendship, the in-group of an ethnic or religious community are Fiske's prototypes. The measurement-theoretic structure is *nominal*: the only meaningful predicate is identity ("same group as / not same group as"). The algebraic structure is an equivalence relation; its quotient is the set of group identities.

Authority Ranking (AR). Participants are ordered along a status hierarchy, and the order is the basis of resource allocation, deference, and direction of action. Military rank, parliamentary seniority, traditional patriarchal family order are Fiske's prototypes. The measurement-theoretic structure is *ordinal*: comparisons "above / below / equal" are meaningful, but differences are not. The algebraic structure is a partial order; its representation is a Hasse diagram.

Equality Matching (EM). Participants exchange in turn, balance of contributions is tracked, and equality of treatment over time is the norm. Carpooling, taking turns, tit-for-tat reciprocity are Fiske's prototypes. The measurement-theoretic structure is *interval*: differences are meaningful, ratios are not (one cannot say that one debt is "twice" another in absolute terms). The algebraic structure is an additive abelian group; its representation is a translation-invariant scale.

Market Pricing (MP). Participants exchange according to a single ratio metric, typically money but in principle any cardinal common scale. Markets, fee-for-service, tax-and-transfer are Fiske's prototypes. The measurement-theoretic structure is *ratio*: ratios are meaningful and a true zero exists. The algebraic structure is a multiplicative semigroup or, with a zero, a field; its representation is a price.

4.2 The Stevens correspondence

That the four modes line up with Stevens's four scales is, on Fiske's own account, "almost too neat to be coincidence" (Fiske 1992, p. 691). The architectural claim of this paper is that the alignment is structural, not coincidental: any social-relational system that admits both unions of relationships (CS plus AR plus EM plus MP within a single multi-faceted relationship, as real relationships always are) and decomposition into atoms must have at most four atoms, because there are at most four self-consistent scales of measurement. The four-mode typology is what survives of Stevens's measurement theory under the constraint of social composability.

4.3 Stitches as binary operations on the loom

We model each Fiske mode as a binary operation on threads of the loom. Let T_1 and T_2 be threads; the four stitches are functions $\sigma_m : T_1 \times T_2 \rightarrow T_{12}$ for $m \in \{\mathrm{CS}, \mathrm{AR}, \mathrm{EM}, \mathrm{MP}\}$, each constrained to respect the scale-type of m:

$$
\begin{aligned}
\sigma_{\mathrm{CS}}\left(T_1, T_2\right) &= \left[T_1\right] \cup \left[T_2\right] \quad \text{(merge equivalence classes)}\\
\sigma_{\mathrm{AR}}\left(T_1, T_2\right) &= T_1 \preccurlyeq T_2 \quad \text{(impose order)}\\
\sigma_{\mathrm{EM}}\left(T_1, T_2\right) &= T_1 + T_2 \quad \text{(sum on additive group)}\\
\sigma_{\mathrm{MP}}\left(T_1, T_2\right) &= T_1 \cdot T_2 \quad \text{(product on multiplicative semigroup)}\,.
\end{aligned}
$$

A real relationship is a *braid* of stitches: a sequence of σ_m applications whose composition produces a structured composite thread. The non-associativity of the substrate means that the order of stitches matters in the way the platform represents the relationship. The braid is the formal object the platform stores; its projection onto the four modes is the user-visible "what kind of relationship this is".

4.4 Mode pathology and the platform's response

Fiske and Tetlock (1997) document the cognitive distress produced by *mode mismatches*: applying a Market-Pricing operation to what participants regard as a Communal-Sharing relationship (e.g., paying a friend for hospitality) produces moral revulsion, not merely awkwardness. The platform can therefore not be neutral about which stitch it applies in which context; it must *predict* the mode the participants regard as appropriate and *act* in that mode. This is precisely an Active-Inference problem: minimise expected free energy over policies that include "which mode to use".

Swarp's relational layer (`shared/schema/collectief.ts`, `shared/schema/welzijn.ts`, parts of `shared/schema/ecosystem.ts`) carries each participant's prior over the four modes for each relationship type, updates the priors by observed reactions, and uses the posterior as the policy distribution from which stitches are sampled. The user-visible effect is that the platform learns when to ask for a price and when not to, when to defer to seniority and when to call a vote. The architectural effect is that mode prediction is a special case of the same dynamics that runs the rest of the platform.

4.5 Summary of section 4

Fiske's four relational modes correspond to the four Stevens scale-types, and that correspondence is structural rather than coincidental. Each mode is realised as a binary operation on threads of the loom, with the algebraic structure dictated by the scale-type. Real relationships are braids of stitches; their composition is non-associative, in line with the substrate. Mode mismatches are the dominant pathology, and the platform predicts and acts on modes via Active Inference applied to the policy space of stitches.

---

5. Memory: Schank's Case-Based Reasoning on the Loom

Sections 2 through 4 specify the platform's substrate, dynamics and relational composition. They do not yet specify how the platform *remembers* — how it stores the experience of the citizens that use it and retrieves that experience when a similar situation arises. The model Swarp adopts is Roger Schank's *dynamic memory* (Schank 1982) and the case-based-reasoning paradigm that grew out of it (Schank 1999; Kolodner 1993; Aamodt & Plaza 1994).

5.1 Cases, MOPs and TOPs

The atom of dynamic memory is the *case*: a record of a specific past experience, indexed by the features that distinguish it from other cases. Cases generalise into *Memory Organisation Packets* (MOPs), which are templates for recurrent patterns of experience (the "going-to-arestaurant” MOP), and into *Thematic Organisation Packets* (TOPs), which are higher-order patterns (the “betrayal-by-mentor” TOP) that capture invariants across MOPs. Scripts (Schank & Abelson 1977) are stereotyped MOPs in which the slot-fillers are largely predictable.

The platform stores citizen experience as cases (`shared/schema/learning.ts`, `shared/schema/coach-content.ts`), aggregates them into MOPs at the level of recurrent practical situations (a community-of-practice meeting, a permit application, a school-choice deliberation), and recognises TOPs at the level of life-course patterns (a career transition, a recovery from illness, a civic radicalisation). The AYYA360 surface uses the TOP layer; the Coach surface uses the MOP layer; the operational surfaces (Marketplace, Initiatives) operate at the case layer.

5.2 Indexing as embedding into the loom

The retrieval problem in dynamic memory is to find the case most relevant to a current situation, which depends entirely on how cases are *indexed*. Schank’s account treats indices as the features under which a case is distinctive — what made it surprising, what failed-and-was-repaired, what generalised badly. The architectural innovation of Swarp is that indices are coordinates on the loom: a case is indexed by its position on the (scale, opgave) base, by its trajectory through the seven generative-model factors of section 3.6, and by the braid of relational stitches that produced it. Two cases are nearby if their loom coordinates are nearby in the natural-gradient metric of section 3.4.

Concretely, the production system embeds each case as a categorical distribution over the eight outcomes of the generative model and stores its representative point on S^7; retrieval is by geodesic nearest-neighbour. The seven imaginary axes anchor the case’s position with respect to the seven generative-model factors; the unit (real) axis carries the case’s overall salience. This makes case retrieval a query on the same geometric object that supports inference and policy selection: there is one geometry across the platform, not three.

5.3 Reminding as graph traversal on the lexicon

A central phenomenon in Schank’s account is *reminding*: a current situation triggers spontaneous recall of a structurally similar past situation, often across superficially different content. Reminding is what distinguishes case-based reasoning from naive nearest-neighbour retrieval; the structural similarity must be discovered, not specified in advance.

Swarp implements reminding as a Markov walk on the *lexicon graph*: a graph whose vertices are the 304 concepts in the ACTIVE lexicon (Appendix B) and whose edges are the explicit semantic relations declared between them. A current situation is mapped to a starting distribution on the lexicon (via the loom coordinates of the case it generates); a Markov walk runs for a number of steps determined by the spectral mixing time (section 7); the resulting distribution over concepts highlights the regions of the lexicon to which the situation is structurally similar. Cases indexed in those regions are retrieved as remindings. The mixing time of the walk (section 7.5) is what determines how *far afield* the platform looks for remindings: a slow-mixing walk produces remindings that are locally coherent (close to the current situation), a fast-mixing walk produces remindings that are globally average. Swarp's spectral gap of 0.0251 (section 7) therefore directly determines its qualitative reminding behaviour.

5.4 Learning as MOP/TOP induction

The platform learns by inducing MOPs from cases and TOPs from MOPs. Induction is implemented as expected-free-energy minimisation over the partition of cases into clusters: a partition is favoured if it reduces complexity (fewer, more general patterns) without sacrificing accuracy (predicting the cases that fall into each cluster). Because the dynamics is Active Inference, learning happens *as a side effect of acting well*: the platform's predictions of citizen reactions, policy outcomes and reminding judgements continually drive MOP/TOP refinement.

5.5 Summary of section 5

Memory is Schank-style dynamic memory: cases, MOPs, TOPs. Cases are indexed by their loom coordinates; retrieval is geodesic nearest-neighbour on S^7. Reminding is a Markov walk on the lexicon graph whose mixing time is the platform's qualitative reach. Learning is induction of MOPs and TOPs by expected-free-energy minimisation, as a side effect of the dynamics.

---

6. System Architecture of Swarp

This section reads the Swarp codebase against the four-layer stack of sections 2–5. The objective is to make precise where each formal element is realised in code, what the cardinal axes of the platform are, what the user-visible surfaces look like, and how the loosely-coupled-subapp decomposition relates to the loom abstraction.

6.1 Four layers in code

The four formal layers map onto four locations in the codebase. The substrate (sections 2.1–2.6) lives in `shared/schema/spatial-web.ts` for the quaternionic part — SWID identifiers, HSML hyperspatial markup, HSTP hyperspatial transport, the universal-domain-graph (UDG) edge set — and in the agent-runtime fibration (`shared/schema/agent-runtime.ts`) for the octonionic part — the seven generative-model factors and the bracketing rules that compose them. The dynamics (section 3) lives in `server/services/inference/*` and in the AIDEN agent-mediation layer; expected-free-energy computation, natural-gradient updates and policy sampling are implemented there, with persistence into `shared/schema/learning.ts`. The relational layer (section 4) lives across `shared/schema/collectief.ts`, `shared/schema/welzijn.ts` and parts of `shared/schema/ecosystem.ts`; the four stitches are realised as four operation kinds and the braid is the persisted `relationship_braid` row. The memory layer (section 5) lives in `shared/schema/learning.ts` (cases), `shared/schema/coach-content.ts` (MOPs) and `shared/schema/meta.ts` (TOPs), with retrieval served by the spectral pipeline in `server/markov-engine.ts`.

6.2 The two cardinal axes

The loom of section 2.6 was specified as a pair of orthogonal axis-systems. In Swarp the two axes are *scale* and *opgave*. Scale takes four values — individual, relational, governance, meta — corresponding to the four registers a citizen-platform must address simultaneously. Opgave takes the four substantive task-domains drawn from the platform's opening sentence: care (zorg), education (onderwijs), economy (economie), democracy (democratie). The product (scale × opgave) is the sixteen-cell base of the loom; each cell carries its own four-dimensional quaternionic fibre of context modifiers; each fibre composes associatively within itself and non-associatively with its neighbours. The thirteen user-visible threads are not the cells; they are diagonal slices through the sixteen-cell base, each slice distinguished by a particular emphasis on certain cells over others.

6.3 The thirteen threads as emergent slices

The thirteen `productLabelNl` entries — AYYA360, Swarp Politiek, Swarp Werk, Swarp Kids, Swarp Academie, Swarp PoC, Swarp Zingeving, Swarp Gezond, Swarp Atelier, Swarp Hobby, Swarp Thuis, Swarp Lab, Swarp Community — were named by the team for marketing and onboarding reasons over the four months of development. Read against the loom, each is a slice through (scale × opgave): AYYA360 is the individual-scale slice across all four opgaven; Swarp Politiek is the (governance, democratie) slice with strong projection onto the (governance, economie) cell; Swarp Gezond is the (individual, zorg) slice with a thin connection to (relational, zorg). The fact that the team converged on thirteen labels rather than sixteen is not a discrepancy: three of the sixteen cells are dominated by other cells (e.g., (meta, onderwijs) is absorbed by Swarp Academie, which lives mostly in (individual, onderwijs) but extends into (meta) by its content-curation function), and the slicing is by saliency, not by partition.

The architectural commitment recorded in `replit.md` and reinforced throughout the codebase is that the threads are *labels*, not modules: renaming `productLabelNl` does not change the platform; only the UI strings change. The loom is what is load-bearing.

6.4 The Loosely Coupled Subapps decomposition

The April 2026 LCS refactor (`replit.md` entry of 30 April 2026) imposed an architectural rule: cross-page imports are forbidden, shared constants and helpers move to `@shared/*`, schemas move to per-thread files in `shared/schema/<draad>.ts` and are re-exported by a barrel `shared/schema/index.ts`. The interim `misc.ts` is allowed to accumulate uncategorised entities between rounds and is drained by per-thread migration. Read against the formal substrate, this is the engineering implementation of the *fibration* structure: each fibre (each `shared/schema/<draad>.ts`) composes associatively within itself, the bus between fibres is mediated rather than direct (no cross-page imports), and the algebraic structure of the inter-fibre composition is realised by the agent-runtime layer.

The LCS rule has a measurable consequence at section 7: it is the architectural cause of the eleven near-disjoint connected components of the lexicon. The platform was designed for vocabulary isolation; the spectral measurement shows the design intent was preserved.

6.5 The Spatial Web infrastructure

The Spatial Web infralaag (`shared/schema/spatial-web.ts`) supplies the quaternionic substrate. SWID is the universal identifier for any addressable entity (citizen, initiative, place, document, agent); HSML is the hyperspatial markup that gives SWID-anchored entities a semantic shape; HSTP is the transport over which HSML messages move. The platform uses SWID as the Markov-blanket identifier of section 3.2, HSML as the carrier of the categorical-distribution embedding of section 3.5, and HSTP as the channel along which expected-free-energy minimisation propagates between agents.

The decision to keep Spatial Web as infrastructure rather than as one of the thirteen threads is recorded in `replit.md` and is consequential. Spatial Web is the substrate; threads are the surface. Promoting Spatial Web to a thread would have committed the category error that the report's first version made about the threads themselves.

6.6 The implementation stack

The platform is implemented as a TypeScript monorepo. The backend is Node.js 20 with Express 5; the database is PostgreSQL 14+ accessed through Drizzle ORM with Zod schemas for validation; the frontend is React 18 with Wouter for routing, TanStack Query for server-state, Tailwind CSS with the shadcn/ui component library, and Recharts for visualisation. The build system is Vite for the client and tsx/esbuild for the server.

Authentication is magic-link (Resend) plus Replit Auth (OIDC). Payments are Mollie. AI inference is OpenAI's GPT-4o for the heavyweight routes and GPT-4o-mini for the lexicon-walk and the Socrates Gastheer onboarding chat. The Socrates Gastheer specifically uses Postgres full-text search (Dutch tsvector with GIN indexes on `blog_posts` and `academia_posts`) plus a Replit-AI-proxy SSE stream — no embeddings, no separate vector store. This is consistent with the architectural commitment that the lexicon graph (section 7) is the platform's primary semantic surface, not a derived index.

6.7 Postcode as the sole user input for location

A small architectural detail with disproportionate consequences: the only location input the platform asks of a user is a postcode. The resolver (`server/services/postcode-resolver.ts`) translates postcodes to municipality, neighbourhood and province via a cache table `postcodes` in `spatial-web.ts`, with PDOK Locatieserver as fallback. The decision matters because it preserves a citizen's right not to disclose finer-grained location while still permitting the platform to anchor her on the relational and governance scales; she is in a wijk and in a gemeente, but the platform does not know her street. This is the privacy implementation of the Markov blanket of section 3.2: the platform addresses s and a at the wijk level, but μ remains internal.

6.8 Summary of section 6

The four formal layers of sections 2–5 are realised in concrete locations of the Swarp codebase. The two cardinal axes of the loom are scale (four values) and opgave (four values). The thirteen user-visible threads are diagonal slices through the sixteen-cell base. The LCS refactor implements the fibration structure as forbidden cross-page imports, and the Spatial Web infralaag supplies the quaternionic substrate. The implementation is a TypeScript monorepo on PostgreSQL with Drizzle, React, OpenAI and Replit-AI proxies. The postcode-only location input is the operational instantiation of the Markov-blanket privacy commitment.

---

7. Empirical Measurement: Spectral Analysis of the Lexicon

The architecture of sections 2–6 is a claim about how Swarp is built. Whether the claim is *true* — whether the platform that runs in production exhibits the structure the formal derivation predicts — is an empirical question. We answer it by spectral analysis of the platform's semantic lexicon: 304 ACTIVE concepts, 1088 semantic edges, 33 distinct domains. The pipeline runs in production at `/api/lexicon/markov/analysis`, with the implementation in `server/markov-engine.ts`. The numbers reported below are from the live snapshot of 7 May 2026.

7.1 Why the lexicon is the right object

The lexicon is the explicit, hand-curated list of concepts the platform knows about and the explicit semantic relations between them. It is not a learned embedding, not a statistical artefact, not a derived structure — it is what the team has agreed the platform is *for*. If the architecture is the loom and the threads are emergent slices, then the lexicon should show the loom-structure: it should partition into near-disjoint vocabularies (the LCS commitment), it should mix slowly along its giant component (the slow-mixing learning surface), and its negative spectrum should reveal the theory–practice axis the architecture insists on. The hypothesis is sharp: the formal derivation predicts qualitative spectral signatures, and the measurement either exhibits them or does not.

7.2 Method

We construct an undirected graph $G = (V, E)$ with $\left|V\right| = 304$ vertices (one per ACTIVE row in `lexicon_concepts`) and $\left|E\right| = 1088$ edges, with an edge $\left(i, j\right)$ existing iff either concept lists the other in its `relatedConceptIds` array. We compute the symmetric adjacency A, the degree vector d, the row-stochastic transition matrix $P = D^{-1}A$, and the symmetric normalised matrix $N = D^{-1/2}AD^{-1/2}$. The stationary distribution of the random walk on G is $\pi_i = d_i / \sum_k d_k$ in closed form (the chain is reversible). Connected components are found by breadth-first search. The top-15 eigenvalues of N are computed by power iteration with Gram–Schmidt deflation against previously found eigenvectors; the first iteration on each subgraph is seeded with $\sqrt{d}$, the exact Perron eigenvector for N, so $\lambda_1 = 1$ is recovered to numerical precision. Clustering coefficients and betweenness centralities (Brandes 2001) are computed for the same graph. The pipeline runs in roughly 800 ms on the live database; results are cached for 60 seconds.

7.3 Graph-level invariants

The measurement returns: 304 ACTIVE concepts, 1088 semantic edges, 33 distinct domains, edge density 0.0236, average degree 7.16, clustering coefficient $C = 0.631$, eleven connected components, giant-component size 232 (76.3% of vertices), spectral gap on the giant $1 - \lambda_2 = 0.0251$, and a derived mixing-time bound of approximately 40 steps. The clustering coefficient of 0.631 is high — substantially higher than the Erdős–Rényi expectation $p = 0.0236$ for a random graph of the same density — which is the first-pass evidence that the lexicon is a curated knowledge graph rather than a stochastic artefact. The eleven connected components are the central spectral signature of the LCS architecture.

7.4 The eleven components and the LCS commitment

The eleven near-disjoint components are not a defect; they are the design. The April 2026 LCS refactor explicitly forbade cross-page imports in the codebase. The semantic consequence — that vocabularies inside threads do not bleed into vocabularies outside their thread — appears in the spectrum as a multiplicity of $\lambda_1 = 1$ that matches the component count. The top eigenvalues of N on the full graph are 1.0000 (multiplicity at least nine in the top-15 list), then 0.9749, 0.9598, 0.9574, 0.9473, 0.7862. The multiplicity is at least nine rather than exactly eleven because two of the smallest components are below the resolution of the top-15 power iteration; the BFS recovers all eleven exactly. The giant component (76.3% of the vocabulary) is the main coordination surface across all eight application domains; the ten satellites are specialised vocabularies — PoC4 simulation, lexicon administration, AIDEN-internal terms, Human-Design micro-vocabulary, and so on — that intentionally do not bleed into the main surface.

This is the exact opposite of the *everything-talks-to-everything* topology that microservice meshes tend to produce in practice. The spectral measurement confirms that the architectural commitment to vocabulary isolation was preserved through implementation.

7.5 Slow mixing and the pedagogical surface

The spectral gap on the giant is $1 - \lambda_2 = 0.0251$, with the implied mixing-time bound $\tau_{\text{mix}} = \Theta\Big(1/\big(1 - \lambda_2\big)\Big) \approx 40$ steps. A standard reading in spectral graph theory (Chung 1997; Levin & Peres 2017) is that a small spectral gap means the random walk takes many steps to lose memory of its starting point. For a learning surface this is the desired property: a citizen who clicks on the concept *Active Inference* and then follows a chain of “next concept” suggestions should spend several steps in the neighbourhood of *Active Inference* before drifting to the global average. A search engine wants the opposite — a large spectral gap so that the answer is always “close” to the query — but Swarp is not a search engine. The measurement shows the lexicon is engineered as a *terrain* over which the citizen travels, not as an *index* into which she queries. The slow-mixing property is the architectural fingerprint of the choice to treat the lexicon as a pedagogical surface rather than as an information-retrieval surface.

The 40-step horizon also has direct operational consequences: the Markov-walk reminding mechanism of section 5.3 should run for between five and forty steps depending on whether the desired reminding is locally coherent (short walk) or thematically associated (long walk). The platform’s reminding parameter is set in this range.

7.6 The bipartite axes and the theory–practice tension

The most negative eigenvalues of the giant component are -0.9574, -0.8137, -0.7932, -0.7644, -0.7440. Negative eigenvalues of N near -1 are the spectral signature of *near-bipartite* structure: the graph admits an approximate two-colouring under which most edges run between colours. The eigenvectors corresponding to these eigenvalues separate concepts of *theoretical foundation* — Free Energy, Surprisal, Markov Blanket, Active Inference, Octonion, Quaternion, Oscillation, Coherence — from concepts of *operational deployment* — Marketplace, Consent, Coach, Health Diagnosis, Neighborhood Circle, Permit Application, Initiative. The bipartite tension is the theory–practice axis the architecture insists on; that the giant remains connected (the bipartition is *near*, not exact) means the theoretical and operational vocabularies are linked rather than fragmented. The architectural response — AIDEN-mediated bridges, cross-thread translations enumerated in Appendix A — is doing its job.

This finding matches the formal derivation directly. Section 3 specified that Active Inference balances *epistemic* and *pragmatic* value in a single nats-valued objective. Section 4 specified that real relationships braid stitches across mode boundaries. The spectral measurement shows that the platform's lexicon is structured as exactly such a balanced opposition: theory and practice are opposing camps, but they remain commensurable.

7.7 The four scales as stationary attractors

The top stationary attractors — the concepts a long random walk lands on most often — are TOA-Triade (Balans Denken-Voelen-Doen, $\pi = 0.0244$, individual-cognitive integration), Marketplace ($\pi = 0.0211$, relational-economic), AIDEN ($\pi = 0.0179$, meta-cognitive), Oscillation ($\pi = 0.0170$, temporal substrate), and Consent ($\pi = 0.0161$, governance). The four scales the architecture is organised around — individual, relational, governance, meta — are precisely the four scales the random walk lands on most often. A long sequence of "next concept" suggestions on `/markov-navigator` will spend most of its time in the orbit of these five concepts. The platform's emergent centre of gravity matches the platform's designed centre of gravity.

7.8 Domain-mass distribution

The fraction of stationary probability captured by all concepts in each domain ranks the domains by their de-facto weight in the platform: Community of Practice (0.146, 35 nodes), AYYA360 (0.106, 31 nodes), governance (0.096, 32 nodes), ecosystem (0.075, 19 nodes), levensloop (0.069, 19 nodes), theoretical foundation (0.069, 21 nodes), system (0.052, 12 nodes), simulation (0.045, 10 nodes), knowledge (0.041, 17 nodes), platform architecture (0.040, 11 nodes), with the remaining 23 domains sharing the other 0.439 of the mass. The top six domains account for 0.561 of the total mass; the distribution is heavy-tailed in the way Zipf's law would predict for a hand-curated vocabulary.

The architectural reading is that the de-facto centre of gravity of Swarp is *relational-collective practice* (CoP) ahead of *individual-developmental practice* (AYYA360) ahead of governance ahead of economy. Whether this matches the team's intentions for the next four months is a question the spectrum does not answer; but the spectrum is the instrument that lets the question be posed concretely rather than vaguely.

7.9 Summary of section 7

The empirical measurement supplies four findings, each of which matches a prediction of the formal derivation: (i) eleven near-disjoint vocabularies match the LCS commitment to fibre isolation; (ii) a small spectral gap of 0.0251 matches the design intent of a slow-mixing pedagogical surface; (iii) four substantial negative eigenvalues match the theory–practice axis the architecture insists on; (iv) the top stationary attractors land on exactly the four scales the architecture predicts. The match across four independent qualitative predictions is the load-bearing argument that the loom is the correct level of description.

8. Discussion

8.1 What the measurement licenses

A qualitative match between a formal derivation and a spectral measurement is not a proof; it is consistency. The four findings of section 7 are consistent with the four predictions of sections 2–6, but a different architecture might produce similar spectra by accident. What rules out accident, in our reading, is that the predictions were made before the measurement was conducted — the architecture was specified by April 2026, the spectral pipeline was added in late April, the measurement of 7 May was the first run on the cleaned ACTIVE lexicon — and that the four predictions are independent of each other. A near-bipartite spectrum could be produced by many architectures; eleven components matching exactly the eleven LCS-isolated subvocabularies, combined with the four-scale stationary attractors and the slow-mixing gap, jointly characterise a much narrower class. We take the joint match as evidence that the loom is the *correct* level of description.

8.2 Comparisons

The closest comparator from the participatory-platform family — civic engagement being one of Swarp's eight application surfaces — is Decidim (Aragón et al. 2017), which decomposes participation into modules (proposals, debates, votes, accountability) and federates them under a shared identity. The architectural difference is that Decidim's modules are first-class and the relations between them are mediated by a thin participation layer; Swarp's threads are emergent and the relations are first-class (the loom). In spectral terms one would expect Decidim's lexicon, were it analysed similarly, to show a much larger spectral gap (modular search-engine surface) and a stronger near-bipartite structure separating procedural from substantive concepts. We have not run the comparison; it is offered as a hypothesis that future work could test.

The closest cognitive-architecture comparator is Anderson's ACT-R (Anderson 2007), which combines a declarative memory of chunks with a procedural memory of production rules. ACT-R is associative and does not exploit non-associativity; its closest analogue to the loom is the partial-matching mechanism that retrieves chunks by feature similarity, but the geometry is a vector space rather than S^7. The architectural payoff of $\mathbb{O}$ over $\mathbb{R}^n$ is the formal expression of context non-commutativity (section 2.5), which ACT-R must implement implicitly via context-cued retrieval rather than explicitly via algebraic non-associativity.

The closest mathematical comparator is the line of work using octonions in physics (Dixon 1994; Furey 2018; Manogue & Dray 2010), which proposes $\mathbb{O}$ as the substrate for the Standard Model of particle physics. Our use is structurally analogous — non-associativity as a feature, the seven imaginary axes as factor labels, the Fano plane as the mnemonic — but the application domain is human-systems coherence rather than fundamental physics. We do not claim a unification; we claim a shared algebraic structure with shared formal payoffs.

8.3 Limitations

Three limitations should be named directly. *First*, the embedding of the categorical-distribution generative model into S^7 (section 3.5) is canonical for the eight-outcome case, but Swarp's actual generative model is much higher-dimensional. The platform handles this by maintaining a hierarchy of seven-axis sub-models nested inside each other (section 3.7), each living on its own copy of S^7, with the loom recursion managing the inter-level composition. The full mathematical treatment of this hierarchy — in particular, the tower of Cayley–Dickson algebras above $\mathbb{O}$ that would be the natural target — is open. The sedenions $\mathbb{S}$ are not a division algebra and cannot be used directly; the right object may be a *para-octonionic* or *split-octonionic* refinement (Conway & Smith 2003, ch. 9), but we have not specified it.

Second, the spectral measurement of section 7 is on the lexicon graph, not on the user-trajectory graph or the case-graph. The lexicon is what the team has agreed the platform is for; the trajectories and the cases are what citizens actually do with it. A trajectory-level measurement would test whether citizens *travel* the loom in the way the architecture invites. We have not yet collected enough trajectory data to make the measurement; once we have, the comparison between lexicon spectrum and trajectory spectrum will be one of the platform's strongest internal-validity tests.

Third, the four Fiske modes are presented as exhaustive (section 4.1), and Fiske himself (1991, 1992) defends exhaustiveness, but the literature contains contested cases (Bolender 2010 on null-relationship as a fifth mode; Rai & Fiske 2011 on moral motivations as cross-cutting). For the loom abstraction the four-mode decomposition is sufficient because it matches the four Stevens scale-types; if a fifth mode is added, the algebraic correspondence breaks and the loom would need a different stitch layer. We treat the four-mode commitment as load-bearing and revisable.

8.4 Future work

Four directions are immediate. *First*, a trajectory-level spectral measurement (limitation two) on the citizen-event log, comparing the spectrum of citizen walks to the spectrum of the lexicon walks and identifying where the two diverge. Divergences are where citizens do *not* travel the loom as designed, and they are the locations where the design needs revision.

Second, a formal treatment of the Cayley–Dickson tower above $\mathbb{O}$ for the hierarchical generative model (limitation one), with explicit specification of the sub-model embedding rule and the inter-level composition.

Third, a comparative spectral analysis against Decidim, Pol.is and other participation platforms (for the civic surface), against Coursera/Khan Academy peer-graphs (for the learning surface), and against PatientsLikeMe/MyFitnessPal (for the health surface), to test the hypothesis (section 8.2) that Swarp's lexicon spectrum is qualitatively different — slower mixing, more components, sharper near-bipartite structure — than its closest comparators.

Fourth, integration of the Spatial Web standards (IEEE P2874) as they consolidate, with SWID, HSML and HSTP moving from internal Swarp implementation to interoperable federation. The architectural commitment is to be ready for that consolidation when it arrives, not to wait for it.

---

9. Conclusion

This paper argued that the architecture of Swarp is not the thirteen user-visible threads but the four-layer stack underneath them. We derived the substrate as the Cayley–Dickson chain $\mathbb{R} \to \mathbb{C} \to \mathbb{H} \to \mathbb{O}$ and identified the platform's load-bearing operation as living at the third doubling, where non-associativity is the architectural feature rather than the defect. We specified the dynamics as Friston's Active Inference embedded on the unit-octonion sphere S^7 via the Bhattacharyya–Kakade categorical embedding, with the seven imaginary axes supplying the seven canonical factors of the generative model and the Fano-plane structure supplying the within-factor associative subalgebras. We modelled the four relational stitches as Fiske's four modes corresponding to Stevens's four scale-types, realised as binary operations on threads of the loom. We embedded Schank-style dynamic memory as cases indexed by their loom coordinates and retrieved by geodesic nearest-neighbour, with reminding implemented as Markov walks on the lexicon graph. We then read the Swarp codebase against this stack, identified the two cardinal axes (scale, opgave) and the thirteen emergent threads, located each formal layer in code, and explained the LCS refactor as the engineering implementation of the fibration structure.

We measured the architecture by spectral analysis of the production lexicon. The 304-concept ACTIVE graph yielded eleven near-disjoint components (matching the LCS commitment), a spectral gap of 0.0251 on the giant component (matching the design intent of a slow-mixing pedagogical surface), four substantial negative eigenvalues (matching the theory–practice axis the architecture insists on), and a stationary distribution dominated by exactly the four scales (matching the predicted centre of gravity). The match across four independent qualitative predictions is the empirical warrant for the architectural claim.

The architecture is a verb. The substrate is a chain of normed division algebras terminating at the smallest non-associative one. The dynamics is variational free energy on the resulting belief manifold. The relations are four algebraic operations corresponding to the four scale-types. The memory is dynamic, indexed on the loom, retrieved by geodesics. The threads are emergent slices, the labels are revisable, and the platform's behaviour at scale is consistent with what the formal derivation predicts. The remaining work is to extend the measurement from the lexicon to the trajectories, and from the trajectories to the lives.

References

Aamodt, A., & Plaza, E. (1994). Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Communications*, 7(1), 39–59.

Adams, D. S., Tseng, A.-S., & Levin, M. (2020). Light-activation of the Archaerhodopsin H+-pump reverses age-dependent loss of neural plasticity. *Bioelectricity*, 2(1), 26–43. (Representative recent work on biological electromagnetism in the Levin lineage.)

Adams, J. F. (1960). On the non-existence of elements of Hopf invariant one. *Annals of Mathematics*, 72(1), 20–104.

Amari, S. (1985). *Differential-Geometrical Methods in Statistics*. Lecture Notes in Statistics, vol. 28. Springer-Verlag, New York.

Amari, S., & Nagaoka, H. (2000). *Methods of Information Geometry*. Translations of Mathematical Monographs, vol. 191. American Mathematical Society and Oxford University Press.

Anderson, J. R. (2007). *How Can the Human Mind Occur in the Physical Universe?* Oxford University Press, New York.

Aragón, P., Kaltenbrunner, A., Calleja-López, A., Pereira, A., Monterde, A., Barandiaran, X. E., & Gómez, V. (2017). Deliberative platform design: The case study of the online discussions in Decidim Barcelona. In *Social Informatics: 9th International Conference, SocInfo 2017*, pp. 277–287. Springer, Cham.

Baez, J. C. (2002). The octonions. *Bulletin of the American Mathematical Society*, 39(2), 145–205.

Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron*, 76(4), 695–711.

Bhattacharyya, A. (1943). On a measure of divergence between two statistical populations defined by their probability distributions. *Bulletin of the Calcutta Mathematical Society*, 35, 99–109.

Bohm, D. (1980). *Wholeness and the Implicate Order*. Routledge & Kegan Paul, London.

Bolender, J. (2010). *The Self-Organizing Social Mind*. MIT Press, Cambridge MA.

Bott, R., & Milnor, J. (1958). On the parallelizability of the spheres. *Bulletin of the American Mathematical Society*, 64(3), 87–89.

Brandes, U. (2001). A faster algorithm for betweenness centrality. *Journal of Mathematical Sociology*, 25(2), 163–177.

Chung, F. R. K. (1997). *Spectral Graph Theory*. CBMS Regional Conference Series in Mathematics, vol. 92. American Mathematical Society, Providence RI.

Conway, J. H., & Smith, D. A. (2003). *On Quaternions and Octonions: Their Geometry, Arithmetic, and Symmetry*. A K Peters, Natick MA.

Dixon, G. M. (1994). *Division Algebras: Octonions, Quaternions, Complex Numbers and the Algebraic Design of Physics*. Mathematics and Its Applications, vol. 290. Kluwer Academic, Dordrecht.

Fiske, A. P. (1991). *Structures of Social Life: The Four Elementary Forms of Human Relations*. Free Press, New York.

Fiske, A. P. (1992). The four elementary forms of sociality: Framework for a unified theory of social relations. *Psychological Review*, 99(4), 689–723.

Fiske, A. P., & Tetlock, P. E. (1997). Taboo trade-offs: Reactions to transactions that transgress the spheres of justice. *Political Psychology*, 18(2), 255–297.

Friston, K. J. (2008). Hierarchical models in the brain. *PLOS Computational Biology*, 4(11), e1000211.

Friston, K. J. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138.

Friston, K. J. (2013). Life as we know it. *Journal of the Royal Society Interface*, 10(86), 20130475.

Friston, K. J., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active inference: A process theory. *Neural Computation*, 29(1), 1–49.

Furey, C. (2018). $\mathrm{SU}(3)\times\mathrm{SU}(2)\times\mathrm{U}(1)(\times\mathrm{U}(1))$ as a symmetry of division algebraic ladder operators. *European Physical Journal C*, 78, 375.

Hari, R., & Salmelin, R. (2012). Magnetoencephalography: From SQUIDs to neuroscience. *NeuroImage*, 61(2), 386–396.

Hamilton, W. R. (1843). On a new species of imaginary quantities connected with the theory of quaternions. *Proceedings of the Royal Irish Academy*, 2, 424–434.

Haslam, N., & Fiske, A. P. (1999). Relational models theory: A confirmatory factor analysis. *Personal Relationships*, 6(2), 241–250.

Hipólito, I., Ramstead, M. J. D., Convertino, L., Bhat, A., Friston, K., & Parr, T. (2021). Markov blankets in the brain. *Neuroscience and Biobehavioral Reviews*, 125, 88–97.

Hunt, V. V. (1989). *Infinite Mind: Science of the Human Vibrations of Consciousness*. Malibu Publishing, Malibu CA.

Hurwitz, A. (1898). Über die Composition der quadratischen Formen von beliebig vielen Variablen. *Nachrichten von der Königlichen Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-physikalische Klasse*, 309–316.

Kakade, S. (2001). A natural policy gradient. In *Advances in Neural Information Processing Systems 14* (NIPS 2001), 1531–1538.

Kirchhoff, M., Parr, T., Palacios, E., Friston, K., & Kiverstein, J. (2018). The Markov blankets of life: Autonomy, active inference and the free energy principle. *Journal of the Royal Society Interface*, 15(138), 20170792.

Kolodner, J. L. (1993). *Case-Based Reasoning*. Morgan Kaufmann, San Mateo CA.

Levin, M. (2014). Molecular bioelectricity: How endogenous voltage potentials control cell behavior and instruct pattern regulation in vivo. *Molecular Biology of the Cell*, 25(24), 3835–3850.

Levin, D. A., & Peres, Y. (2017). *Markov Chains and Mixing Times* (2nd ed.). American Mathematical Society, Providence RI.

Manogue, C. A., & Dray, T. (2010). Octonions, E_6, and particle physics. *Journal of Physics: Conference Series*, 254, 012005.

McCraty, R., Atkinson, M., Tomasino, D., & Bradley, R. T. (2009). The coherent heart: Heart-brain interactions, psychophysiological coherence, and the emergence of system-wide order. *Integral Review*, 5(2), 10–115.

Moufang, R. (1935). Zur Struktur von Alternativkörpern. *Mathematische Annalen*, 110, 416–430.

Persinger, M. A. (2010). The harribance effect as pervasive out-of-body experiences: NeuroQuantal evidence with more precise measurements. *NeuroQuantology*, 8(4), 444–465.

Porges, S. W. (2007). The polyvagal perspective. *Biological Psychology*, 74(2), 116–143.

Pribram, K. H. (1991). *Brain and Perception: Holonomy and Structure in Figural Processing*. Lawrence Erlbaum, Hillsdale NJ.

Rai, T. S., & Fiske, A. P. (2011). Moral psychology is relationship regulation: Moral motives for unity, hierarchy, equality, and proportionality. *Psychological Review*, 118(1), 57–75.

Schank, R. C. (1982). *Dynamic Memory: A Theory of Reminding and Learning in Computers and People*. Cambridge University Press.

Schank, R. C. (1999). *Dynamic Memory Revisited*. Cambridge University Press.

Schank, R. C., & Abelson, R. P. (1977). *Scripts, Plans, Goals, and Understanding: An Inquiry into Human Knowledge Structures*. Lawrence Erlbaum, Hillsdale NJ.

Stevens, S. S. (1946). On the theory of scales of measurement. *Science*, 103(2684), 677–680.

Appendix A — Page and Route Inventory of the Swarp Platform

This appendix enumerates the user-visible surfaces of Swarp as of 7 May 2026 and maps each to the formal layer of sections 2–5 it primarily realises. Routes are given in the form they appear in the React/Wouter routing table (`client/src/App.tsx`) and the Express router (`src/server/routes/`).

A.1 AYYA360 surface (individual scale, life-course)

`/ayya360` — entry hub for the individual-developmental thread. Reads from the TOP layer (section 5.1) of dynamic memory and projects onto the Human-Design / Spiral-Dynamics / RIASEC / PoC narrative profiles. Each tile leads to a profile detail (`/mijn-spiritueel-profiel`, `/mijn-cognitief-profiel`, `/mijn-beroeps-profiel`) so that no badge or label is a dead end (the *no-dead-ends* commitment recorded in `replit.md`, 5 May 2026 entry). The surface is the individual register of the loom (scale = individual, opgave varying).

`/levensloop` — the four-phase life-course timeline (childhood, youth, adulthood, elderhood). Reads from MOPs at the level of recurrent life-course situations and renders them as a navigable timeline. The temporal axis here is the recursion of the loom into itself (section 3.7).

`/coach` — the Coach surface. Operates at the MOP layer, with each session indexed as a case in `learning.cases` and aggregated into the citizen's MOP table.

A.2 Politiek surface (governance scale, democracy opgave)

`/mijn-politiek-profiel` — the citizen's per-administrative-layer Fiske-quartet political profile. Each layer (municipality, province, national, European) carries a four-vector over CS/AR/EM/MP and a Bayesian PPP that updates as the citizen interacts with party programmes and council items. Drift analysis is rendered as a sparkline per layer.

`/partijen` and `/partijen/:id` — party comparison surface. Each party is a row of the same Fiske-quartet structure as the citizen profile, so the comparison is formally a divergence on the same four-dimensional simplex per layer.

`/gemeenten` and `/gemeenten/:slug` — municipality entry points. Each municipality projects the (governance, varying-opgave) cell of the loom onto its local agenda, council items (where ORI is available), and local-news feed (Sleutelstad, NH Nieuws, et al.) where ORI is not available. The bottom-up commitment of the platform (recorded in `replit.md`) is implemented here: ORI is one source among many, not the architecture.

`/initiatieven` — citizen initiatives. Each initiative is a case in the same dynamic-memory layer used by Coach and AYYA360, indexed on the loom by its (scale, opgave) cell and its braid of Fiske stitches.

A.3 Academia / Agora surface (knowledge and deliberation)

`/academia` and `/blog/:slug` — the knowledge base. The corpus is the migrated `constable.blog` body (sections of which are listed below) plus the Academia-internal essays. Retrieval is Postgres FTS (Dutch tsvector, GIN indexes); the Socrates Gastheer chat at the top of `/academia` is the onboarding interface for new readers.

The load-bearing essays for the present paper are: *The Spiral Navigator* (`/blog/the-spiral-navigator`, 10 March 2026); *Spiral Learning Theory* (`/blog/spiral-learning-theory`, 18 March 2026); *Swarp als Leer Systeem* (`/blog/swarp-als-leer-systeem`, 29 March 2026); *SWARP: Hoe Mensen Écht Leren* (`/blog/swarp-hoe-mensen-echt-leren-een-revolutionair-case-based-learning-systeem`, 29 March 2026); *Wat is Swarp Agora?* (`/blog/wat-is-swarp-agora`).

`/markov-navigator` — the lexicon-walk surface. Implements section 5.3's reminding mechanism as a user-driven Markov walk, with the spectral pipeline (section 7) supplying the next-step distribution.

A.4 Marketplace and Ecosystem surface (relational economy)

`/marktplaats` — the local marketplace. Each transaction is a case indexed on the loom; the Fiske-mode predictor (section 4.4) decides whether the transaction is rendered as Market Pricing (price, fee), Equality Matching (turn-taking, time-bank), Authority Ranking (deference to seniority), or Communal Sharing (gift, pool).

`/zzp` — the freelancer collaboration surface. Implements the (relational, economy) cell with strong projection onto (relational, education) for skill-sharing.

A.5 Welzijn and Health (individual–relational care)

`/welzijn` — well-being surface. Operates at the (individual, care) and (relational, care) cells. Each well-being check is a case; aggregation produces the citizen's well-being MOPs.

`/gezondheid` — health surface. Same structure as `/welzijn`, with stronger institutional integration (provider directory, appointment).

A.6 Meta surface

`/aiden` — the agent-mediation interface. AIDEN is the meta-cognitive register of the loom (section 6.2); the surface lets a citizen inspect the agents that act on her behalf, adjust their priors, and audit their actions.

`/lexicon` — the lexicon administration surface. Direct view of the 304 ACTIVE concepts and their semantic edges; supports adding, retiring and re-relating concepts. The surface is restricted to maintainer roles.

`/markov-analysis` — the spectral analysis dashboard. Renders the live results of section 7 as graphs and tables, refreshed on each request to `/api/lexicon/markov/analysis`.

A.7 Cross-cutting routes

`/api/postcode/:pc4` — the postcode resolver of section 6.7.

`/api/lexicon/markov/analysis` — the spectral pipeline of section 7.

`/api/socrates/host` — the Socrates Gastheer SSE endpoint that powers the `/academia` onboarding chat.

`/api/aiden/policy` — the policy-evaluation endpoint that returns the AIDEN expected-free-energy score for a candidate action.

Appendix B — The 304-Concept ACTIVE Lexicon

This appendix summarises the lexicon as it stood on 7 May 2026. The full list of concept names, their domains, their levels (1–5, depth in the ontology) and their `relatedConceptIds` is too long to reproduce verbatim; it is queryable in production at `/lexicon` and via the API at `/api/lexicon/concepts`. We give here the structural summary that supports section 7.

B.1 Domain breakdown

The 304 ACTIVE concepts span 33 domains. By stationary mass, ranked: Community of Practice (35 concepts, 0.146 mass); AYYA360 (31 concepts, 0.106); governance (32 concepts, 0.096); ecosystem (19, 0.075); levensloop (19, 0.069); theoretical foundation (21, 0.069); system (12, 0.052); simulation (10, 0.045); knowledge (17, 0.041); platform architecture (11, 0.040); content (10, 0.038); VHS (volkshogeschool, 9, 0.028); mystiek (9, 0.027); juridisch (8, 0.024); social (7, 0.022); collective intelligence (5, 0.017); collective action (6, 0.015); agenda (6, 0.015); paden (2, 0.012); health (2, 0.012); the remaining 13 domains share the residual 0.106 mass.

B.2 Top-25 stationary attractors

Ranked by stationary probability π, with degree d and betweenness b:

1. TOA-Triade (Balans Denken-Voelen-Doen), AYYA360, level 3 — $\pi = 0.0244$, $d = 53$, $b = 1.000$.
2. Marketplace (Marktplaats), ecosystem, level 5 — $\pi = 0.0211$, $d = 46$, $b = 0.636$.
3. AIDEN, system, level 2 — $\pi = 0.0179$, $d = 39$, $b = 0.442$.
4. Oscillation (Oscillatie), theoretical foundation, level 1 — $\pi = 0.0170$, $d = 37$, $b = 0.751$.
5. Consent, governance, level 5 — $\pi = 0.0161$, $d = 35$, $b = 0.459$.
6. Contextmonitor, CoP, level 3 — $\pi = 0.0156$, $d = 34$, $b = 0.012$.
7. Besluitvormer, CoP, level 3 — $\pi = 0.0156$, $d = 34$, $b = 0.012$.
8. Programmamaker, CoP, level 3 — $\pi = 0.0156$, $d = 34$, $b = 0.012$.
9. Eigenfrequentie, CoP, level 3 — $\pi = 0.0156$, $d = 34$, $b = 0.012$.
10. MetaSwarp, platform architecture, level 2 — $\pi = 0.0147$, $d = 32$, $b = 0.105$.
11. Support Circle (Ondersteuningsgroep), levensloop, level 2 — $\pi = 0.0142$, $d = 31$, $b = 0.071$.
12. CoP Toegangscontrole, CoP, level 3 — $\pi = 0.0142$, $d = 31$, $b = 0.009$.
13. Agent, simulation, level 2 — $\pi = 0.0115$, $d = 25$, $b = 0.250$.
14. Coherence (Coherentie), simulation, level 2 — $\pi = 0.0115$, $d = 25$, $b = 0.727$.
15. Layers of Life (Lagen van het Leven), paden, level 4 — $\pi = 0.0110$, $d = 24$, $b = 0.161$.
16. Health Diagnosis (Gezondheidsdiagnose), health, level 4 — $\pi = 0.0101$, $d = 22$, $b = 0.005$.
17. Chronotope (Chronotoop), content, level 1 — $\pi = 0.0101$, $d = 22$, $b = 0.244$.
18. Coach, levensloop, level 2 — $\pi = 0.0092$, $d = 20$, $b = 0.003$.
19. Case Management (Casusmanagement), levensloop, level 3 — $\pi = 0.0092$, $d = 20$, $b = 0.003$.
20. Freelancer Collaboration (ZZP Samenwerking), ecosystem, level 5 — $\pi = 0.0087$, $d = 19$, $b = 0.056$.
21. Creative Profile (Creatief Profiel), AYYA360, level 4 — $\pi = 0.0083$, $d = 18$, $b = 0.011$.

22. Community Purchasing (Buurtinkoop), ecosystem, level 5 — $\pi = 0.0083$, $d = 18$, $b = 0.033$.
23. Personal Coach (Persoonlijke Coach), AYYA360, level 4 — $\pi = 0.0074$, $d = 16$, $b = 0.025$.
24. Politieke Partij, governance, level 4 — $\pi = 0.0074$, $d = 16$, $b = 0.098$.
25. Neighborhood Circle (Wijkkring), governance, level 5 — $\pi = 0.0069$, $d = 15$, $b = 0.100$.

B.3 Reading the top-25

The five highest betweenness-centrality concepts are TOA-Triade ($b = 1.000$), Oscillation ($b = 0.751$), Coherence ($b = 0.727$), Marketplace ($b = 0.636$) and Consent ($b = 0.459$). These are the lexicon's structural bridges: removing any of them would partition the giant component into substantially smaller pieces. That TOA-Triade — the individual-cognitive integration concept — is the single most central node by betweenness is the spectral confirmation that the platform's individual register is not a side-feature but the structural pivot of the whole vocabulary. That Oscillation and Coherence (both theoretical-foundation concepts) sit second and third confirms that the theoretical layer is not decoration; it is load-bearing.

The CoP cluster (rows 6–9 and 12) shares an identical degree of 34 and a near-zero betweenness, which is the spectral signature of a tightly clustered sub-vocabulary internal to the CoP domain, weakly connected to the rest. This is the LCS commitment in microcosm.

Appendix C — Spectral Computation Pipeline and Numerical Results

This appendix documents the computation in enough detail that an independent implementation could reproduce the numbers of section 7. The reference implementation is `server/markov-engine.ts` in the Swarp codebase.

C.1 Three architectural questions, three spectral regimes

The pipeline was built to answer three architectural questions concurrently. *Connectedness*: how many disjoint vocabularies has the lexicon actually grown? The thirteen-thread soft split of section 6 was a design intention; the graph might or might not honour it. *Mixing rate*: over how many "next concept" steps does a learning path retain thematic coherence before it dissolves into the system's stationary attractors? *Symmetry breaks*: are there structural axes inside the lexicon — bipartitions, opposing camps — that reveal architectural tension the team had not yet named?

These three questions map precisely onto the three regimes of the spectrum of $N = D^{-1/2} A D^{-1/2}$: the multiplicity of $\lambda = 1$ (number of components), the spectral gap $1 - \lambda_2$ (mixing rate), and the magnitude of the most negative eigenvalues (near-bipartite symmetry breaks). The Markov chain was built because there is no other instrument that answers all three at once on the actual production graph.

C.2 Pipeline

On every request to `/api/lexicon/markov/analysis` the engine performs the following.

1. Load all rows from `lexicon_concepts` where `status = 'ACTIVE'` (Drizzle query in `server/markov-engine.ts:fetchNetwork`).
2. Build the symmetric adjacency A from `relatedConceptIds`: an edge $\left(i, j\right)$ exists if either concept lists the other.
3. Compute the degree vector d and the row-stochastic transition matrix $P = D^{-1}A$. Its stationary distribution is $\pi_i = d_i / \sum_k d_k$ (closed form for an undirected reversible chain).
4. Find the connected components by breadth-first search on the adjacency.
5. Compute the top-15 eigenvalues of $N = D^{-1/2} A D^{-1/2}$ on the full graph, and again on the giant component, by power iteration with Gram–Schmidt deflation against previously found eigenvectors. The first iteration on each subgraph is seeded with the vector $\sqrt{d}$, which is the exact Perron eigenvector for N, so $\lambda_1 = 1$ is recovered to numerical precision rather than approximated.
6. Compute the clustering coefficient $\overline{C}$ and the betweenness centrality (Brandes 2001) for the same graph.

The computation runs in roughly 800 ms on the live database; results are cached for 60 s.

C.3 Graph-level invariants (snapshot of 7 May 2026)

ACTIVE concepts (nodes): 304. Semantic edges: 1088. Distinct domains: 33. Edge density: 0.0236. Average degree: 7.16. Clustering coefficient $\overline{C}$: 0.631. Connected components: 11. Giant-component size: 232 (76.3 %). Spectral gap on the giant ($1 - \lambda_2$): 0.0251. Mixing-time bound: ≈ 40 steps.

C.4 Top-15 eigenvalues of N

On the *full graph*, the top-15 eigenvalues are: 1.0000 (multiplicity at least nine in the top-15 list), then 0.9749, 0.9598, 0.9574, 0.9473, 0.7862, with the most negative eigenvalues entering at -0.9574. The multiplicity of $\lambda = 1$ matches the BFS-confirmed component count of 11 up to the resolution of the power iteration (the two smallest components contribute eigenvectors with very low effective weight).

On the *giant component*, the top-15 eigenvalues are: 1.0000, 0.9749, 0.9598, 0.9473, 0.9369, 0.9323, 0.9234, 0.9013, 0.8505, 0.8164, then negative: $-0.8137, -0.7932, 0.7782, -0.7644, -0.7440$. The sequence $0.9749, 0.9598, 0.9473, 0.9369, 0.9323, 0.9234, 0.9013, 0.8505$ is the cluster spectrum inside the largest component; the four substantial negatives $(-0.81, -0.79, -0.76, -0.74)$ are the bipartite axes that resolve as the theory–practice tension of section 7.6.

C.5 What each spectral regime tells us

Multiplicity of $\lambda = 1$. Eleven near-disjoint vocabularies — the LCS commitment of section 6.4 confirmed by direct measurement. Swarp is not a single application but a federation of eleven loosely coupled subapps welded together by deliberate bridges. The giant component (76 %) is the main coordination surface across all eight application domains; the ten satellites are specialised vocabularies that intentionally do not bleed into the main surface.

Spectral gap. $1 - \lambda_2 = 0.0251$, mixing time ≈ 40 steps. The giant component is a slow-mixing learning surface, not a search engine. A search engine wants a large gap so that the answer is always “close” to the query; a learning surface wants a small gap so that the user can travel meaningfully through neighbourhoods of related concepts before being pulled to the centre. Swarp is the second, by measurement.

Negative spectrum. Four substantial negative eigenvalues separate concepts of theoretical foundation from concepts of operational deployment. The same theory–practice tension section 3 named as the architectural problem of integrating Friston with Schank is visible in the spectrum but does not fragment the giant into two components — meaning the architectural response (AIDEN-mediated bridges, cross-thread translations of Appendix A) is doing its job: theory and practice live in opposing camps but remain connected.

Stationary distribution. The four scales of integration the architecture predicts are present in measurement: TOA-Triade and Oscillation anchor the individual register; Marketplace and Consent anchor the relational and governance registers; AIDEN anchors the meta register. A long sequence of “next concept” suggestions on `/markov-navigator` will spend most of its time in their orbit. The platform’s emergent centre of gravity matches its designed centre of gravity.

C.6 The reading

The Markov chain was built to make the lexicon legible to a team that could no longer hold it in their heads. It turned out, on running the numbers, to make the architecture legible as well. That is the strongest justification for keeping the spectral pipeline in production: the system is now able to describe itself, and what it describes is consistent with what the four months of work intended to build.

Appendix D — Quaternion and Octonion Multiplication Tables

For reference, and for the implementer who wishes to verify the algebraic claims of section 2.

D.1 Quaternions ℍ

Hamilton's relations: $i^2 = j^2 = k^2 = ijk = -1$. Full multiplication table:

	1	i	j	k
1	1	i	j	k
i	i	-1	k	-j
j	j	-k	-1	i
k	k	j	-i	-1

Conjugation: $\overline{a + bi + cj + dk} = a - bi - cj - dk$. Norm: $\left|q\right|^2 = a^2 + b^2 + c^2 + d^2$. Inverse: $q^{-1} = \overline{q}/\left|q\right|^2$. Multiplicative norm: $\left|pq\right| = \left|p\right|\left|q\right|$. Non-commutative: $ij = k$, $ji = -k$. Associative.

D.2 Octonions 𝕆 — Fano-plane lines

Basis $\{1, e_1, e_2, e_3, e_4, e_5, e_6, e_7\}$. Each $e_i^2 = -1$. The seven oriented Fano-plane triples (Cayley convention used by Baez 2002):

$\{e_1, e_2, e_3\}$: $e_1e_2 = e_3$, $e_2e_3 = e_1$, $e_3e_1 = e_2$ (and reverses negated).

$\{e_1, e_4, e_5\}$: $e_1e_4 = e_5$, $e_4e_5 = e_1$, $e_5e_1 = e_4$.

$\{e_1, e_7, e_6\}$: $e_1e_7 = e_6$, $e_7e_6 = e_1$, $e_6e_1 = e_7$.

$\{e_2, e_4, e_6\}$: $e_2 e_4 = e_6$, $e_4 e_6 = e_2$, $e_6 e_2 = e_4$.

$\{e_2, e_5, e_7\}$: $e_2 e_5 = e_7$, $e_5 e_7 = e_2$, $e_7 e_2 = e_5$.

$\{e_3, e_4, e_7\}$: $e_3 e_4 = e_7$, $e_4 e_7 = e_3$, $e_7 e_3 = e_4$.

$\{e_3, e_6, e_5\}$: $e_3 e_6 = e_5$, $e_6 e_5 = e_3$, $e_5 e_3 = e_6$.

For any $i \neq j$, $e_i e_j = -e_j e_i$. The seven triples are exactly the seven lines of the Fano plane $\mathbb{P}^2(\mathbb{F}_2)$, and each carries an associative quaternionic subalgebra generated by 1 and the triple. There are seven such subalgebras; they are the seven canonical factors of section 3.6.

D.3 Loss of associativity, retention of alternativity

In general $(e_i e_j)e_k \neq e_i(e_j e_k)$. The associator $[x, y, z] = (xy)z - x(yz)$ is alternating: $[x, y, z] = -[y, x, z] = -[x, z, y]$. Equivalently, any subalgebra generated by two elements is associative. The Moufang identities,

$$(zxz)y = z(x(zy)), \qquad y(zxz) = ((yz)x)z, \qquad (zx)(yz) = z(xy)z,$$

are the strongest surviving associativity-like rules. They are what makes octonion arithmetic tractable in software.

D.4 The Hurwitz chain

The four normed division algebras over $\mathbb{R}$ are $\mathbb{R}$ (dim 1), $\mathbb{C}$ (dim 2), $\mathbb{H}$ (dim 4), $\mathbb{O}$ (dim 8). At each Cayley–Dickson doubling the algebra loses a structural property: $\mathbb{R} \to \mathbb{C}$ loses ordering; $\mathbb{C} \to \mathbb{H}$ loses commutativity; $\mathbb{H} \to \mathbb{O}$ loses associativity. The next doubling — $\mathbb{O} \to \mathbb{S}$ (sedenions, dim 16) — loses the property of being a division algebra (zero divisors appear), and the sequence of *normed* division algebras terminates.

That termination is the formal warrant for Swarp's substrate choice. The richest substrate compatible with division and a multiplicative norm is $\mathbb{O}$. Going further sacrifices the property "non-zero times non-zero is non-zero", which the platform cannot survive. Stopping earlier sacrifices the formal expression of context non-commutativity, which the platform requires. The substrate is what it must be.

End of paper.