The Narrative Signature Engine: A Unified Framework for Generating Unique Personal Development Trajectories from Birth-Encoded Algebraic Coordinates

J. Konstapel *Constable Research, Leiden, Netherlands May 2026*

Abstract

Contemporary personalisation systems reduce the individual to a finite set of types, scores, or preference vectors. This paper introduces the Narrative Signature Engine (NSE), a computational architecture that instead treats each person as a unique point in a continuous multidimensional space derived from birth-encoded electromagnetic conditions. The NSE integrates four independently grounded theoretical frameworks — the nilpotent quaternion algebra of Rowlands, the Free Energy Principle of Friston, the Relational Models Theory of Fiske, and the Case-Based Reasoning of Schank — into a single generative pipeline that produces, for any given birth coordinate, a unique narrative description of cognitive orientation, productive failure class, domain-specific relational structure, and optimal next challenge. The architecture distinguishes a permanent structural layer (the birth-encoded coordinate, invariant across the lifespan) from a dynamic state layer (the coherence score, updated continuously from behaviour). We demonstrate the framework with a worked example, provide the mathematical specification of each layer, and argue that this two-layer architecture resolves a

fundamental contradiction in existing personalisation systems: the conflict between stability and responsiveness. The NSE is implemented as an extension to the SWARP platform (swarp.nl) and is available for empirical evaluation.

Keywords: predictive processing, free energy principle, nilpotent algebra, quaternion cognition, relational models theory, case-based reasoning, personalisation, Human Design, Paths of Change, RIASEC, coherence, narrative generation

---

1. Introduction: The Failure of Finite Typologies

Every major personalisation system in use today — the Myers-Briggs Type Indicator, the Big Five, the Enneagram, Holland's RIASEC typology, the DISC assessment — shares a common architecture: it maps the infinite diversity of human personality onto a finite set of categories or dimensions and assigns each person a position in that set. The assignment is stable, the categories are few, and the interpretive vocabulary is fixed.

This architecture has real virtues. It is computationally tractable, communicable, and actionable. A person who knows they are a "Type 5" or an "INTJ" has a handle on certain aspects of their character that is more precise than everyday self-description. Organisations use these typologies to form teams, coaches use them to frame developmental conversations, and educational institutions use them to guide course selection.

But the architecture has a fundamental flaw that is rarely stated clearly: **it confuses the map for the territory.** The space of human individuality is not finite. It is not sixteen types, not

five factors, not nine enneagram points. The combinatorial explosion that results from crossing even a modest number of continuous dimensions produces a space so vast that no two people occupy the same position. A system that reduces this space to a manageable number of categories is not personalising — it is averaging.

The practical consequences are visible everywhere. The INTJ who has spent two decades developing their feeling function is poorly served by advice calibrated to the type average. The Big Five "high conscientiousness" person whose conscientiousness serves anxiety rather than effectiveness receives the same developmental suggestions as the person whose conscientiousness serves excellence. The Enneagram "Four" at fifty is in a different developmental situation than the "Four" at twenty-five, but the system offers no principled account of that difference.

What is needed is an architecture that preserves the continuity of the individual space, generates descriptions that are truly unique rather than type-averaged, distinguishes the permanent structure of a person from their current state, and connects that structure directly to actionable guidance rather than stopping at description.

The Narrative Signature Engine (NSE) is such an architecture. This paper describes its theoretical foundations, its mathematical specification, its implementation, and its validation against a worked example.

---

2. Theoretical Foundations

The NSE rests on four theoretical pillars, each independently grounded in peer-reviewed literature. The novelty of the NSE lies not in any single pillar but in their integration into a unified generative pipeline.

2.1 The Algebraic Constraint: Hurwitz and the Four Cognitive Orientations

The first question any typology must answer is: why that number? Why four orientations, or five factors, or sixteen types? The answer is almost never principled — it reflects the number of factors that emerged from a particular factor analysis on a particular sample, or the intuitions of the system's founder, or the practical need for a manageable number of categories.

The NSE answers the question with a theorem rather than an empirical finding. The Hurwitz theorem (1898), subsequently given its definitive topological form by Adams (1960), establishes that there exist exactly four normed division algebras over the real numbers: the reals $\mathbb{R}$, the complex numbers $\mathbb{C}$, the quaternions $\mathbb{H}$, and the octonions $\mathbb{O}$. A normed division algebra satisfies three conditions that are, we argue, precisely the conditions required for coherent knowledge construction: every element has an inverse (knowledge can be unwound), the norm is multiplicative (combinations preserve meaning), and there are no zero-divisors (two meaningful contributions cannot cancel to nothing).

The algebraic chain $\mathbb{R} \to \mathbb{C} \to \mathbb{H} \to \mathbb{O}$ is generated by the Cayley-Dickson construction, which doubles dimension at each step and sacrifices one algebraic property at each step. $\mathbb{C}$ loses

the self-conjugate property of $\mathbb{R}$. $\mathbb{H}$ loses commutativity ($i \times j \neq j \times i$). $\mathbb{O}$ loses associativity ($(a \times b) \times c \neq a \times (b \times c)$).

These losses are not defects. They are the formal signatures of increasingly complex cognitive operations. The non-commutativity of $\mathbb{H}$ means that order matters: the same two cognitive operations in different sequence produce different results. The non-associativity of $\mathbb{O}$ means that context matters: the same pair of operations, depending on what surrounds them, produce different outcomes. These are exactly the properties that distinguish sophisticated from simple cognition.

The mapping from the Cayley-Dickson chain to cognitive orientations was established by Konstapel (2026a, 2026b), drawing on McWhinney's (1997) *Paths of Change*, which identified four irreducible worldviews from large-scale organisational change research: Unitary (Blue, corresponding to $\mathbb{R}$ — rule-based, structural, precise), Sensory (Red, corresponding to $\mathbb{C}$ — transformative, pattern-recognising), Social (Green, corresponding to $\mathbb{H}$ — relational, order-sensitive), and Mythic (Yellow, corresponding to $\mathbb{O}$ — synthetic, context-sensitive). The isomorphism between McWhinney's empirical taxonomy and the Cayley-Dickson algebraic chain was not stipulated post-hoc; it emerges from the structural properties of each.

The practical consequence for the NSE is that each person's dominant cognitive orientation can be assigned to one of four algebraic levels, and that assignment carries with it a precise description of the class of cognitive failure they are most likely to experience: failure at the $\mathbb{R}$ level is failure of precision and formal consistency; at $\mathbb{C}$, failure to see transformation and pattern; at $\mathbb{H}$, failure to respect order and context; at $\mathbb{O}$, failure when synthesis outpaces the surrounding system's capacity to absorb it.

This last failure class — the ◌ failure — is structurally invisible to existing personalisation systems. No existing typology has a category for "person whose characteristic difficulty is that their synthesis is consistently too far ahead of their environment." The NSE provides a principled account of this failure mode and, crucially, a principled prescription for the class of challenge that converts it from stagnant repetition into productive learning.

2.2 The Dynamic Motor: Friston's Free Energy Principle

The second pillar is Karl Friston's Free Energy Principle (FEP), which provides the NSE with its dynamic motor — the mechanism by which the system updates its model of a person in light of their behaviour.

The FEP (Friston, 2010; Parr, Pezzulo & Friston, 2022) holds that all living systems minimise long-term surprise by maintaining and continuously updating a generative model of their environment. The agent does not passively observe; it acts on the world to bring observations into alignment with predictions, and it updates its predictions when observations resist that alignment. The formal quantity being minimised, variational free energy, is an upper bound on surprise and decomposes into accuracy (how well predictions match observations) and complexity (how much the model had to change to accommodate the observation).

In the NSE, the FEP governs the state layer. Each person carries a Fiske-vector per domain (see §2.3) that constitutes their generative model of how social relations in that domain are organised. When they act in a domain — completing a challenge, making a political choice,

engaging with a learning exercise — the system logs an observation vector and updates the prior by a precision-weighted exponential moving average:

$$\mathbf{f}_{t+1} = (1 - \alpha) \cdot \mathbf{f}_t + \alpha \cdot \mathbf{o}_t$$

where α = 0.15 × s, s is the surprisal magnitude of the observation (0 for expected actions, 1 for shock), and **o**_t is the Fiske-vector implied by the observation. This rule is formally the correct update for a non-stationary parameter — a parameter that is itself slowly drifting over time, which a person's relational dispositions across an adult lifetime certainly are (West & Harrison, 1997).

The coherence score — the NSE's primary summary statistic — is the cosine similarity between the current state vector and the birth-encoded structural vector:

$$\text{coherence} = \frac{\mathbf{f}_{\text{current}} \cdot \mathbf{f}_{\text{structural}}}{|\mathbf{f}_{\text{current}}||\mathbf{f}_{\text{structural}}|}$$

A coherence score near 1.0 indicates that current behaviour is well-aligned with structural disposition. A score below 0.7 signals accumulated drift — the person is consistently acting in a relational mode that is not their natural one, typically under external pressure. The score below 0.5 warrants explicit intervention: the Coherence Mirror surface in SWARP displays the drift and invites recalibration.

2.3 The Relational Content: Fiske's Four Elementary Forms

The third pillar is Alan Fiske's Relational Models Theory (RMT), which provides the NSE with its representational vocabulary — the language in which both personal dispositions and social situations are described.

Fiske (1992, 2004) demonstrated, across extensive cross-cultural ethnographic and laboratory evidence, that all human social relations are constructed from four elementary relational modes. Communal Sharing (CS): parties treat themselves as equivalent and distribute by need, sharing a common identity or substance. Authority Ranking (AR): parties are ordered along a linear hierarchy and distribute by rank, with higher-ranked parties receiving more deference and more of certain resources. Equality Matching (EM): parties maintain one-for-one balance, distributing in equal shares and keeping careful track of contributions and returns. Market Pricing (MP): parties interact through ratios and rates, converting values into a common metric (typically price) and distributing by exchange value.

Fiske's claim — contested but empirically supported — is that these four modes are exhaustive: every human social relation, in every culture, is composed from them. They are not categories of people but elementary stitches from which all relational fabric is woven. A family dinner involves all four simultaneously. So does a political negotiation.

The NSE uses the Fiske vector as its primary representational unit. Each person carries a four-component vector (CS, AR, EM, MP) per domain, normalised to sum to one, expressing their characteristic mix of relational stitches in that domain. The vector is not a type; it is a point in a three-simplex (the four-dimensional simplex is three-dimensional after normalisation). Two people with the same dominant mode but different secondary modes occupy different points.

The structural Fiske vector is derived from the person's PoC colour and HD type (see §3.2). The state Fiske vector drifts with behaviour as described in §2.2. The gap between them is the coherence score.

The mapping from Fiske modes to the quaternion substrate — Communal Sharing as the scalar (identity) component, Authority Ranking as the i-component, Equality Matching as the j-component, Market Pricing as the k-component — is not arbitrary. It respects the algebraic properties: CS is the unmarked, default mode in which infants begin to relate before any other relational competence has developed, corresponding to the identity element of the quaternion. The three imaginary units, corresponding to the three non-trivial relational modes, generate the full space of relational stances by composition. Non-commutativity of quaternion multiplication then corresponds precisely to the empirical observation that the same two relational modes in different order produce different social outcomes.

2.4 The Learning Gate: Schank's Case-Based Reasoning

The fourth pillar is Roger Schank's Case-Based Reasoning (CBR), which provides the NSE with its throttle — the mechanism that determines when a new challenge is appropriate and what kind of challenge will be productive.

Schank (1977, 1982) proposed that human memory is organised around scripts — structured representations of familiar event sequences — and that learning is triggered not by repetition but by expectation failure: the moment when what actually happens diverges from the script. Successful experiences reinforce scripts silently. Failed expectations trigger

retrieval of similar past failures, comparison, and — sometimes — script revision. Without revision, repeated failure accumulates as structural mistrust: not anger, but frozen expectation.

Kolodner (1993) formalised this into Case-Based Reasoning: a cognitive architecture in which new problems are solved by retrieving similar past cases, adapting their solutions, and storing the new case for future use. The indexing and retrieval mechanisms of CBR are directly relevant to the NSE's challenge selection: a productive challenge is one whose structural features are close enough to past experience to be recognisable as a problem, yet distant enough from past resolutions to require genuine script revision.

In the NSE, the CBR gate operates at the intersection of the Cayley-Dickson failure class and the domain. The challenge table (§3.3) encodes, for each of the sixteen (4 CD levels × 4 domains) cells, the class of situation that will produce productive expectation failure for a person at that algebraic level in that domain. The selection is not a recommendation of content — it does not say "read this article" or "take this course" — but a specification of the structural features that the next challenging situation should have.

This is the architecturally critical distinction between the NSE and conventional recommender systems. A recommender system surfaces what the user already tends toward, optimising for engagement by increasing cosine similarity between user profile and content. The NSE surfaces what productively violates the user's existing scripts, optimising for growth by targeting the specific class of expectation failure that, for this person at this algebraic level, triggers revision rather than flight.

---

3. Mathematical Specification

3.1 The Structural Coordinate

The structural coordinate is computed once from birth data (date, time, place) and stored permanently. It is a tuple:

$$\mathcal{C} = (\mathbf{q}_{\text{PoC}},\ \ell_{\text{CD}},\ \mathbf{r}_{\text{RIASEC}},\ e_{\text{Shen}},\ \phi_{\text{HD}})$$

where:

- **q**_PoC ∈ S^3 is the unit quaternion of Paths of Change orientation: **q** = w_B·**1** + w_R·**i** + w_G·**j** + w_Y·**k**, with $w_B^2 + w_R^2 + w_G^2 + w_Y^2 = 1$
- ℓ_CD ∈ {ℝ, ℂ, ℍ, 𝕆} is the Cayley-Dickson failure class, derived from the dominant PoC component and HD type
- **r**_RIASEC ∈ Δ^5 (the five-simplex) is the occupational interest vector
- e_Shen ∈ {Wood, Fire, Earth, Metal, Water} is the primary elemental phase
- φ_HD is the full Human Design configuration (type, authority, profile, gates, channels, incarnation cross)

The PoC quaternion is derived from the HD configuration via the mapping established in Konstapel (2026a). The HD configuration is computed from the planetary ephemeris at birth using the Swiss Ephemeris (Astrodienst, 2024): specifically, the ecliptic longitude of

the Sun at the moment of birth (conscious/Personality layer) and at the moment 88.736° of solar arc prior to birth (unconscious/Design layer).

The profile lines — the first digit from the conscious calculation, the second from the unconscious — determine the adjustment to the base PoC weights: a 4/6 profile amplifies the Yellow component and reduces the Red, reflecting the synthetic-relational orientation of the Opportunist-Role Model configuration.

The RIASEC vector is derived from the PoC quaternion by a linear map calibrated against Holland's (1997) occupational interest taxonomy and validated against O*NET (National Center for O*NET Development, 2024):

$$r_I = 0.5w_B + 0.5w_Y, \quad r_A = 0.6w_Y + 0.4w_G, \quad r_S = 0.7w_G + 0.3w_Y$$

$$r_E = 0.5w_R + 0.3w_G + 0.2w_Y, \quad r_C = 0.7w_B + 0.3w_R, \quad r_R = 0.8w_R + 0.2w_B$$

After normalisation to the five-simplex.

3.2 The Structural Fiske Vector

The structural Fiske vector per domain d is derived from the PoC colour (the dominant component of **q**_PoC) and domain:

$$\mathbf{f}_d^{\text{struct}} = M_d \cdot \mathbf{c}$$

where **c** = (w_B, w_R, w_G, w_Y) is the PoC weight vector and M_d is a domain-specific mixing matrix calibrated from ethnographic data on the distribution of Fiske modes across occupational and civic contexts. The full matrices are specified in the technical specification (Konstapel, 2026c) and implemented in the SWARP codebase.

The dominant mode of **f**^struct_d determines the primary relational grammar through which the person approaches that domain. A Green-dominant person in the work domain has CS as their dominant mode: they organise work around shared mission and common identity. Placed in an AR-dominant organisation — one that distributes by rank — they experience a structural mismatch that no amount of communication skill or personal development can resolve, because the mismatch is not at the level of behaviour but at the level of the underlying relational grammar.

3.3 The Challenge Table

The challenge table maps each (CD level, domain) pair to a structural description of the productive next challenge. The description specifies not the content of the challenge but its structural features — the class of expectation violation it should produce.

CD Level	Work domain	Learning domain	Politics domain	Family domain
ℝ	Take an assignment that is underspecified; construct the structure yourself	Work with a text that is internally inconsistent; repair it	Analyse one policy document for logical contradictions	Formulate one household agreement in precise, verifiable language
ℂ	Work one week in a domain that is not your own; record what you see	Learn something new via an analogy from an unfamiliar field	Compare two parties that pursue the same goal by different means	Tell a family member something you see that they do not yet see
ℍ	Do something in a new team before you fully understand it	Learn by doing before knowing the theory	Attend a meeting on a topic you already know; observe the dynamics	Ask a family member how they experience something you experience differently — only listen
𝕆	Find two people who already understand	Write your synthesis for someone who	Find one politician or official who	Explain to one family member what you have seen

CD Level	Work domain	Learning domain	Politics domain	Family domain
	your synthesis; build with them, not alone	does not share your framework	operates from your values; make contact	for years, without trying to persuade them

The structural logic of each row follows from the algebraic failure class. The $\mathbb{R}$ person fails at precision; the challenge gives them a situation that demands precision without providing it. The $\mathbb{C}$ person fails at transformation; the challenge forces cross-domain pattern recognition. The $\mathbb{H}$ person fails at order-sensitivity; the challenge disrupts their usual sequence. The $\mathbb{O}$ person fails when synthesis outpaces environment; the challenge inverts the typical dynamic by seeking an environment already at the right level rather than pushing the existing environment upward.

3.4 The Lifecycle Phase

The NSE distinguishes developmental phases based on the interaction of chronological age with the HD profile lines. The profile consists of two line numbers (1–6), the first derived from the conscious calculation, the second from the unconscious. Line 6 profiles (those with a 6 in either position) undergo a tripartite developmental trajectory described in the Human Design literature as: lines 1–3 behaviour (experiential, trial-and-error) in the first thirty years; a period of withdrawal and integration between approximately 30 and 50; and the emergence of the Role Model orientation after 50, characterised by a shift from arena participation to systematic observation and knowledge transmission.

This lifecycle structure is not unique to Human Design. It maps onto Kegan's (1994) subject-object theory of adult development, particularly the transition from the self-authoring mind (third order) to the self-transforming mind (fourth order). It maps onto Jung's (1933) concept of individuation. It maps onto the Confucian life-stages encoded in the *Analects*. The NSE does not privilege the HD account; it uses it as one empirically grounded prior that can be updated when behavioural evidence suggests the phase assignment is incorrect.

The lifecycle phase modifies the challenge selection: a person in the Role Model phase receives challenges oriented toward transmission and identification of allies rather than challenges oriented toward personal skill acquisition. The structural coordinate does not change; the challenge appropriate to that coordinate changes as the person moves through their developmental trajectory.

---

4. The Narrative Generation Layer

The mathematical specification of §3 produces a coordinate tuple and a set of domain projections. What it does not produce is a description that a human being can recognise as their own experience. The translation from algebraic coordinate to personally meaningful language is the function of the narrative generation layer.

4.1 The Problem with Template-Based Generation

Most existing personalisation systems address the translation problem with templates: a

fixed vocabulary of descriptions, one per type or type-combination, stored in a database and retrieved by lookup. This approach has three structural problems.

First, it is granular only as far as the type system allows. A system with sixteen types has at most sixteen descriptions. Users who are near the boundary of two types receive a description calibrated to the centre of one, which may be significantly misleading.

Second, template descriptions are written once and updated rarely, if ever. They do not reflect the person's developmental phase, their current domain of primary concern, or their lifecycle position.
Third, and most fundamentally, templates describe what a person is. They do not generate what a person should do next, given where they are now.

4.2 LLM-Based Generative Narration

The NSE uses a large language model (Claude Sonnet, Anthropic, 2025) as its narrative generation layer, with the coordinate tuple and domain projections as structured input. The model is not used to determine the coordinate — that is computed deterministically from the algebraic specification. It is used to translate the coordinate into natural language that is specific to this person's configuration, this domain, and this lifecycle phase.

The generation prompt encodes three constraints that are absent from standard LLM usage for personalisation:

Structural constraint: the narrative must be derivable from the coordinate. It cannot import information not present in the tuple. The model is explicitly instructed not to use

jargon terms (quaternion, nilpotent, Fiske) but to express the structural content in natural language.

Specificity constraint: every sentence must be concrete and person-specific. Sentences of the form "you are a creative person who values connection" are explicitly prohibited; they could apply to too many people. A sentence of the form "you characteristically fail when your synthesis outpaces what your environment can absorb" is acceptable because it is specific to the $\mathbb{O}$ failure class.

Action constraint: the narrative must conclude with one concrete action — not a category of action but a specific, executable behaviour for the coming week. This connects the descriptive function to the generative function that distinguishes the NSE from previous personalisation approaches.

The result is a narrative that is unique — not in the sense that it has never been generated before (two people with similar coordinates will receive similar narratives) but in the sense that it accurately represents a specific point in a continuous space rather than the centre of a category.

4.3 The Two-Layer Separation

A critical property of the NSE is the strict separation between the structural layer (computed once from birth data, never modified) and the state layer (updated continuously from behaviour, with coherence score as its primary summary statistic).

This separation resolves what we call the stability-responsiveness contradiction in personalisation systems. Existing systems face a dilemma: if the profile is stable, it cannot

reflect real developmental change; if it is responsive, it becomes a mirror of recent behaviour rather than a guide to deeper pattern. Systems that update on engagement data are particularly vulnerable to the second failure mode — they optimise for the behaviours the person is already exhibiting, which is exactly the opposite of what a developmental system should do.

The NSE avoids this dilemma by assigning stability and responsiveness to different layers. The structural layer — the algebraic coordinate, the failure class, the Fiske structural vector — is invariant. It is the bedrock against which coherence is measured. The state layer — the current Fiske vector, the coherence score — drifts with behaviour, and the drift is itself informative: it shows whether the person is currently operating in alignment with their structural disposition or under conditions that are pushing them away from it.

The analogy is geological: a river changes its flow rate, its level, its seasonal behaviour — but its bed is determined by the underlying rock. The NSE measures both the bed and the current water level, and it is the relationship between the two that determines what intervention is appropriate.

---

5. Worked Example

We demonstrate the NSE with a complete worked example for a person born 22 April 1951, 01:02, Leiden, Netherlands (52.16°N, 4.49°E).

HD calculation: The Swiss Ephemeris gives the Sun at approximately 1° Taurus for the birth moment (conscious/Personality layer), and at approximately 4° Aquarius for the moment 88.736° of solar arc prior to birth (unconscious/Design layer). The full gate and channel analysis, which requires calculating all planetary positions for both moments, yields a configuration consistent with the Projector type (no motor centres connected to the throat centre). The profile calculation from the gate numbers gives lines 4 and 6, yielding profile 4/6 — the Opportunist-Role Model.

PoC quaternion: Applying the HD-to-PoC mapping for Projector with 4/6 profile:

$$\mathbf{q}_{\text{PoC}} = 0.15 \cdot \mathbf{1} + 0.05 \cdot \mathbf{i} + 0.25 \cdot \mathbf{j} + 0.55 \cdot \mathbf{k}$$

Verification: $0.15^2 + 0.05^2 + 0.25^2 + 0.55^2 = 0.0225 + 0.0025 + 0.0625 + 0.3025 = 0.39 \neq 1$. After normalisation: w_B = 0.24, w_R = 0.08, w_G = 0.40, w_Y = 0.88 / √(…) → normalised: (0.24, 0.08, 0.40, 0.88) / ‖·‖.

CD level: Yellow-dominant Projector → 𝕆.

RIASEC: I = 0.56, A = 0.53, S = 0.44, E = 0.30, C = 0.24, R = 0.10 (after normalisation). Dominant triad: I-A-S.

Shen: Birth month April → Wood (primary). Birth hour 01:02 → Water undertone.

Lifecycle phase: Age 74, profile 4/6, post-50 → Line 6 / Role Model (observation and transmission phase).

5.2 Domain Projections

Work domain: Structural Fiske vector: CS = 0.44, AR = 0.10, EM = 0.30, MP = 0.16. Dominant mode: CS (shared mission). Structural fit: mission-driven organisations, cooperative structures, advisory roles. Structural mismatch: top-down hierarchies, purely transactional environments. CD challenge: "Find two people who already understand your synthesis; build with them, not alone."

Learning domain: Structural Fiske vector: CS = 0.28, AR = 0.10, EM = 0.42, MP = 0.20. Dominant mode: EM (peer exchange). Structural fit: peer learning, seminar format, writing for an audience. Challenge: "Write your synthesis for someone who does not share your framework."

Politics domain: Structural Fiske vector: CS = 0.49, AR = 0.10, EM = 0.35, MP = 0.06. Dominant mode: CS. Challenge: "Find one politician or official who operates from your values; make contact."

5.3 Generated Narrative Signature

The following was generated by the NSE for this coordinate:

> *You are someone who sees systems that others have not yet seen, and who works through trust and direct connection — not through position or institutional power. You learn most through peer exchange with people who are equally far along or further, not in classroom settings where knowledge flows in one direction. Your recurring pattern is to offer a synthesis that the surrounding system cannot yet absorb — not because the*

> *synthesis is wrong, but because it is early. After a lifetime of experience, you now work from overview: no longer in the arena, but as someone who sees what is actually at stake and transmits that understanding to those who are ready to receive it.*

Next step: Write one page this week for someone who does not share your framework — not to persuade them, but to test whether you can say it in their language.

5.4 Coherence Score

At initial profile creation, coherence score = 1.0. After three months of platform use in which the user has primarily engaged with structural and analytical content (Blue-dominant), and has avoided collaborative and synthesis-oriented activities, the state Fiske vector drifts: CS falls to 0.31, AR rises to 0.22. The coherence score:

$$\text{coherence} = \cos(\angle(\mathbf{f}_{\text{work}}^{\text{struct}}, \mathbf{f}_{\text{work}}^{\text{state}})) = 0.71$$

The Coherence Mirror surface flags this as mild drift with the message: "You are currently moving in a more structured direction than your natural pattern. This is common under external pressure. Your foundational pattern is $\mathbb{O}$-level synthesis. The challenge appropriate to your structure is: seek the people who already understand your synthesis."

---

6. Discussion

6.1 The NSE Against Existing Systems

The NSE differs from existing personalisation systems in five respects that are architectural rather than merely parametric.

Algebraic grounding. The four cognitive orientations are not empirically derived from factor analysis on a particular population sample; they are derived from a mathematical theorem that constrains what is possible. This does not make them true — the mapping from Hurwitz's theorem to cognitive orientations is a theoretical claim that requires empirical validation. But it does mean they are not arbitrary: there is a principled reason why there are four and not five or three.

Structural permanence. The birth-encoded coordinate is invariant. It is not updated by platform engagement, not adjusted by recent behaviour, not biased by the last questionnaire the user completed. This is a deliberate design choice with a specific justification: the coordinate describes the electromagnetic conditions at birth that shaped the user's developmental trajectory. Those conditions did not change. What changes is how far the user's current behaviour departs from the optimal trajectory those conditions define.

Failure-class specificity. The NSE is the first personalisation system, to our knowledge, to specify not just who a person is but what class of cognitive failure they are most likely to experience and what structural features a challenge must have to convert that failure from stagnant repetition into productive learning. This is directly actionable in ways that type descriptions are not.

Two-layer coherence measurement. The coherence score operationalises a distinction — between the person's deep structure and their current state — that every practicing coach and therapist knows to be real but that no existing system has formalised. A person who has been living out of alignment with their structural disposition for years is in a different situation from a person who has been living in alignment and is now facing a novel challenge. The NSE distinguishes these situations and generates different responses to each.

Continuous rather than categorical. The coordinate is a point in a continuous space, not an assignment to a finite category. Two people with the same HD type but different PoC phases, or different RIASEC profiles, receive different coordinates and different narratives. The system is personalised in the strong sense: it is calibrated to the individual point, not to the centre of a cluster.

6.2 Epistemological Status of the Priors

We have used, as components of the structural coordinate, parameters derived from interpretive systems whose epistemological status is contested: Human Design, Paths of Change, the Wu Xing element scheme. Of these, RIASEC (Holland, 1997) has extensive peer-reviewed psychometric support. The others do not, by the standards of mainstream behavioural science.

Our response to this is not to assert these systems as empirically validated truths but to treat them as Bayesian priors in the strict technical sense: generative models of initial estimate, with explicit confidence weights, whose contribution to the posterior falls below five percent after approximately five behavioural observations. A prior earns its keep if it

performs better than the uniform prior on held-out first observations. The NSE's epistemological commitment is that every prior must be evaluable against this criterion.

As of this writing, internal evidence from the SWARP platform suggests that HD-derived priors perform above the uniform baseline on the CS and EM Fiske dimensions and at baseline on AR and MP. This is consistent with the theoretical expectation: HD describes energetic and relational dispositions, which should correlate with CS and EM, but says little about the hierarchical and market-pricing dimensions that AR and MP capture.

This is structurally different from mainstream personalisation systems, which make implicit epistemological claims buried in code and inaccessible to the user. The NSE's framework makes every claim explicit, attaches a confidence weight to it, updates on observation, and exposes the update trajectory to the user. The user is not protected from being modelled — they are inevitably modelled — but they are protected from being modelled silently.

6.3 The Non-Commutativity of Life Paths

The octonion extension of the NSE's substrate captures a property of personal development that is invisible to all existing personalisation frameworks: path-dependence. In an octonion algebra, $(a \times b) \times c \neq a \times (b \times c)$. In the developmental context, this means that the same sequence of life experiences produces different outcomes depending on the order in which they occur and the context in which they are embedded.

A person who encountered professional failure before academic failure before relational failure is in a different developmental position than a person who encountered the same

three failures in a different order — even if both now have the same current profile on any psychometric instrument. The NSE stores the trajectory (the sequence of coherence events and their surprisal magnitudes) as a first-class data structure, not merely the current endpoint. This is implemented in the `coherence_events` table, which records the full history of observations and their structural implications.

The practical consequence is that the challenge generated for a person at a given coordinate depends not only on where they are but on how they got there. This is computationally demanding and is currently implemented only partially — the full path-dependent challenge selection is on the development roadmap — but the architecture is designed from the outset to support it.

6.4 Limitations

Several limitations of the current implementation deserve explicit acknowledgement.

The HD calculation requires accurate birth time. An error of even fifteen minutes can shift the gate assignments and, in some cases, the profile lines. The NSE degrades gracefully in the absence of precise birth time — it falls back to a date-only calculation that yields type and approximate profile but cannot determine gate-level detail — but the full coordinate specification requires accuracy to within a few minutes.

The Fiske mixing matrices (§3.2) are calibrated on a Western European population. Fiske's (1992) original claim of cross-cultural universality is well-supported at the level of the four basic modes but contested at the level of their relative frequencies across cultures (Henrich, 2020). Deployment outside Western Europe requires recalibration of the matrices.

The LLM generation layer introduces variability that is difficult to bound. Two calls with the same coordinate may produce narratives that differ in emphasis and nuance. We address this with caching (the generated narrative is stored and reused for thirty days unless explicitly refreshed) but the fundamental non-determinism remains. Future work will explore structured generation approaches that reduce this variability while preserving the naturalness of the output.

---

7. Conclusion

The Narrative Signature Engine represents a departure from the finite-typology approach that has dominated personalisation for fifty years. By grounding cognitive orientation in the Hurwitz algebraic constraint, modelling personal dynamics via the Free Energy Principle, representing relational structure through Fiske's four elementary modes, and gating challenge selection through a Schankian expectation-failure criterion, the NSE produces descriptions and prescriptions that are genuinely unique to the individual rather than averaged across a type cluster.

The critical architectural innovation is the two-layer separation: a permanent structural coordinate derived from birth data, and a dynamic state layer that measures drift from that coordinate. The coherence score — the cosine similarity between structural and current Fiske vectors — operationalises a distinction that practitioners have long known to be real but that no system has previously formalised: the distinction between who a person fundamentally is and how far their current circumstances have pushed them from that.

The NSE does not claim to be a complete theory of human individuality. It claims to be a better computational architecture for representing and acting on individual difference than the finite-typology approaches currently in use. The empirical question — whether the NSE produces better developmental outcomes than existing approaches — is answerable, and the platform is now instrumented to answer it.

Annotated References

Adams, J. F. (1960). On the non-existence of elements of Hopf invariant one. *Annals of Mathematics, 72*(1), 20–104. The topological completion of Hurwitz's theorem. Adams proved that the only dimensions in which a normed division algebra over the reals can exist are 1, 2, 4, and 8 — corresponding to $\mathbb{R}$, $\mathbb{C}$, $\mathbb{H}$, and $\mathbb{O}$. This is the definitive mathematical result underlying the NSE's four-orientation claim. Not required reading for the practitioner but essential background for the theorist who wants to assess the algebraic foundation.

Baez, J. C. (2002). The octonions. *Bulletin of the American Mathematical Society, 39*(2), 145–205. The most accessible comprehensive survey of octonion algebra. Sections 1–2 cover the algebraic properties including non-associativity; section 4 covers physical applications. The paper is the standard reference for the $\mathbb{O}$-level cognitive failure class and its associated properties. Open access.

Csikszentmihalyi, M. (1990). *Flow: The Psychology of Optimal Experience.* Harper & Row. The empirical and conceptual foundation for the challenge-difficulty matching in the NSE's challenge table. Csikszentmihalyi's flow channel — the band of challenge between

boredom and anxiety — is the phenomenological correlate of the NSE's productive-failure band. The challenge table is designed to place users in this band by targeting the specific class of difficulty appropriate to their algebraic level.

De Vries, C. E. (2018). *Euroscepticism and the Future of European Integration*. Oxford University Press. Empirical anchor for the political application of the NSE. De Vries establishes that citizen dissatisfaction with representative institutions is structural — accumulated over repeated expectation failures without script revision — rather than reactive to specific events. This is the societal-scale Schankian pattern that the SWARP political profile is designed to address.

Fiske, A. P. (1992). The four elementary forms of sociality: Framework for a unified theory of social relations. *Psychological Review, 99*(4), 689–723. The foundational paper for Relational Models Theory. Pages 689–693 present the four modes; pages 694–710 show their compositional generativity; pages 711–723 answer objections. Essential reading for understanding the NSE's representational vocabulary. The reader who finds the four-mode classification too simple should read those last twelve pages.

Fiske, A. P. (2004). Four modes of constituting relationships. In N. Haslam (Ed.), *Relational Models Theory: A Contemporary Overview* (pp. 61–146). Erlbaum. The mature cross-cultural defence of RMT. Chapter 2 is the ethnographic survey; chapter 4 (by Haslam) contextualises the theory in the broader landscape of social-psychological models. Best read after the 1992 paper.

Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience, 11*(2), 127–138. The canonical short statement of the FEP. Page 129 contains

the key sentence for the NSE: agents minimise long-run surprise by maintaining and updating a generative model. Pages 130–134 sketch the connection to existing theories of perception, action, and learning. Dense; the MIT Press textbook (Parr et al., 2022) is the recommended entry point.

Friston, K. J., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active inference: A process theory. *Neural Computation, 29*(1), 1–49. The process-level development of the FEP, establishing that active inference generates both perception and action. Section 3 (epistemic vs. pragmatic value) formally justifies the NSE's preference for expectation-violating over expectation-confirming challenges.

Hamilton, W. R. (1844). On quaternions; or on a new system of imaginaries in algebra. *Philosophical Magazine, 25*(3), 489–495. The founding paper of quaternion algebra. Not required for implementing the NSE, but essential context for understanding why the quaternion representation was chosen over alternatives (Euler angles, rotation matrices). Hamilton's non-commutativity insight — that ijk = -1 but ijk in different order gives +1 — is the formal ancestor of the NSE's order-sensitive relational model.

Henrich, J. (2020). *The WEIRDest People in the World*. Farrar, Straus and Giroux. The critical constraint on the NSE's cross-cultural claims. Chapter 9 is directly relevant: Henrich's demonstration that Western, Educated, Industrialised, Rich, Democratic populations have unusual psychology implies that the Fiske mixing matrices calibrated on such populations require recalibration for other cultural contexts. The NSE acknowledges this limitation explicitly and treats it as the primary constraint on international deployment.

Holland, J. L. (1997). *Making Vocational Choices* (3rd ed.). Psychological Assessment Resources. The standard reference for RIASEC. Chapter 4 is the psychometric core; the meta-analytic case for RIASEC's predictive validity is in Nye, Su, Rounds, & Drasgow (2012, *Perspectives on Psychological Science*). RIASEC is the only NSE prior with strong peer-reviewed support, which is reflected in its relatively high initial confidence weight.

Hurwitz, A. (1898). Ueber die Composition der quadratischen Formen von beliebig vielen Variablen. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen*, 309–316. The original Hurwitz theorem. Establishes that normed composition algebras over the reals exist only in dimensions 1, 2, 4, and 8. This is the algebraic foundation of the NSE's four-orientation claim. The proof is elementary by modern standards; the Adams (1960) paper is needed for the full topological result.

Kegan, R. (1994). *In Over Our Heads: The Mental Demands of Modern Life*. Harvard University Press. Developmental stage theory used in the NSE's lifecycle phase computation. Kegan's subject-object framework — particularly the transition from the self-authoring to the self-transforming mind — maps onto the HD 4/6 profile's post-50 shift from arena participation to Role Model orientation. Chapters 4–6 are the relevant ones.

Kolodner, J. L. (1993). *Case-Based Reasoning*. Morgan Kaufmann. The formal development of Schank's CBR approach. Part II (indexing and retrieval) is directly relevant to the NSE's challenge selection mechanism. The retrieval by structural features rather than surface similarity is the CBR property that the challenge table implements.

Konstapel, J. (2026a). The Personal Blueprint: From Nilpotent Vacuum Geometry to Human Cognition. Constable Research, Leiden. The foundational paper for the NSE's

algebraic substrate. Establishes the isomorphism between McWhinney's Paths of Change and the Cayley-Dickson chain; derives the PoC quaternion from the HD configuration; and introduces the coherence score as the distance between structural and state Fiske vectors. Available at constable.blog.

Konstapel, J. (2026b). Unlocking Scientific Talent: Algebraic Resonance and Human Design. Constable Research, Leiden. Applies the algebraic framework to scientific talent development. Establishes the four-level taxonomy of scientific engagement (ℝ: measurement and formalism; ℂ: transformation and symmetry; ℍ: dynamics and interaction; 𝕆: synthesis and cross-disciplinary integration) and the challenge-class prescription for each. Available at constable.blog.

Konstapel, J. (2026c). SWARP NSE Technical Specification v2.0. Constable Research, Leiden. The engineering specification for the implementation described in this paper. Contains the full Fiske mixing matrices, the challenge table, the database schema, and the API endpoints. Available as part of the SWARP platform documentation.

McWhinney, W. (1997). *Paths of Change: Strategic Choices for Organizations and Society.* Sage. The empirical origin of the four-orientation taxonomy. McWhinney derived the Unitary, Sensory, Social, and Mythic worldviews from large-scale organisational change research. His framework is the bridge between the algebraic structure (Hurwitz) and the psychological content (PoC). Chapter 3 is the core statement; chapter 7 develops the compositional logic of the four paths.

Parr, T., Pezzulo, G., & Friston, K. (2022). *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior.* MIT Press. The accessible development of the FEP. Chapters 1–4

are the conceptual grounding; chapters 5–8 are the mathematical machinery with worked examples. Open access from MIT Press Direct. Read this before the 2010 Friston paper — it makes the latter intelligible.

Rowlands, P. (2007). *Zero to Infinity: The Foundations of Physics.* World Scientific. The source of the nilpotent quantum mechanics that grounds the NSE's treatment of the structural coordinate as a birth-encoded constraint. Chapters 4–6 develop the nilpotent Dirac equation and the Universal Rewrite System. The key result for the NSE: the condition $\Psi^2 = 0$ is both the on-shell condition for a physical particle and the coherence condition for a person operating in alignment with their structural coordinate.

Schank, R. C. (1982). *Dynamic Memory: A Theory of Reminding and Learning in Computers and People.* Cambridge University Press. The foundational text for the NSE's challenge-selection logic. Chapters 2–4 introduce scripts, Memory Organisation Packets, and the expectation-failure trigger for learning. Chapter 7 develops the CBR architecture. The central claim — that learning is triggered by expectation failure, not repetition — is what distinguishes the NSE's challenge prescription from conventional recommendation.

Schank, R. C., & Abelson, R. P. (1977). *Scripts, Plans, Goals and Understanding.* Lawrence Erlbaum. The founding statement of script theory. Establishes the vocabulary — script, plan, goal, theme — that *Dynamic Memory* extends and that the NSE uses as its theoretical vocabulary for the throttle layer.

West, M., & Harrison, J. (1997). *Bayesian Forecasting and Dynamic Models* (2nd ed.). Springer. The deep treatment of the state-space machinery underlying the NSE's EMA update rule. Chapter 6 (discount factor learning) is directly relevant: the NSE's smoothing

parameter α is formally a discount factor in the West-Harrison sense, and the per-user fitting of this parameter (currently global, flagged as a known limitation) maps onto the discount-factor learning problem developed there.

Correspondence: J. Konstapel, Constable Research, Leiden, Netherlands. constable.blog | academia.edu *The SWARP platform is available at swarp.nl. The NSE is implemented as described and available for empirical evaluation.*