The Narrative Signature Engine: Generating Unique Personal Development Trajectories from Electromagnetic Birth Coordinates and Empirical Occupational Data

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 derived from questionnaire responses. This paper introduces the Narrative Signature Engine (NSE), a computational architecture that treats each person as a unique point in a continuous multidimensional space derived from two independent empirical sources: the electromagnetic conditions encoded at birth (translated into a unit quaternion via the Hurwitz algebraic constraint and the nilpotent vacuum geometry of Rowlands), and the occupational failure topology extracted from the O*NET database and U.S. Census Bureau labour data (1960–2010), in which a simulation of the AYYA Python system revealed a stable attractor at the golden ratio $\varphi \approx 1.618$. The NSE integrates this dual derivation with Friston's Free Energy Principle, Fiske's Relational Models Theory, and Schank's Case-Based Reasoning into a single generative pipeline that produces, for any given individual, 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 (birth-encoded, invariant) from a dynamic state

layer (coherence score, updated continuously from behaviour). We demonstrate the framework with a worked example, provide the complete mathematical specification, and argue that the two-layer architecture resolves the fundamental contradiction in existing personalisation systems between stability and responsiveness.

Keywords: predictive processing, free energy principle, nilpotent algebra, quaternion cognition, relational models theory, O*NET, occupational failure, phi attractor, coherence, narrative generation, SWARP

---

1. Introduction: The Empirical Problem

Every major personalisation system in current use shares a common failure mode: it maps the continuous space of human individuality onto a finite set of categories and assigns each person a position in that set. The assignment is derived from questionnaire responses — from what a person reports about themselves — rather than from any structural feature of the person. The result is a map that is legible but systematically inaccurate: it describes the average of a cluster rather than the individual point.

The problem is not that the categories are too few or too many. The problem is architectural. A questionnaire-derived typology has no principled connection to the generative structure it purports to measure. It describes correlational patterns in self-report data. It cannot answer the question that matters most for development: given where this specific person is now, what is the next challenge that will produce genuine learning rather than defensive withdrawal?

The NSE was developed from a different starting point. Rather than beginning with questionnaire data and searching for factor structure, the NSE begins with two independent empirical derivations.

The first is algebraic: the Hurwitz theorem constrains the number of irreducible cognitive orientations to exactly four, and Rowlands' nilpotent quantum mechanics provides the physical mechanism by which birth-encoded electromagnetic conditions set the initial weights on those four orientations. This is not a metaphor or an analogy. It is a claim about scale invariance: the same nilpotent algebra that governs the on-shell condition for an electron ($\Psi^2 = 0$) governs the coherence condition for a human cognitive system.

The second is empirical: an analysis of the O*NET database combined with U.S. Census Bureau occupational data from 1960 to 2010, processed through a Python simulation of the AYYA coherence engine, revealed a stable attractor in the occupational transition space at $\varphi \approx 1.618$ — the golden ratio. This finding, unexpected and not stipulated in advance, means that the distribution of productive occupational failure sequences is not uniform but concentrates around a specific geometric proportion. The NSE uses this empirical φ-attractor as its calibration constant for challenge spacing: the temporal and difficulty ratio between successive productive challenges follows the golden section.

The integration of these two derivations — the algebraic structural constraint and the empirical φ-attractor — produces a personalisation architecture that is grounded in physics and calibrated from data rather than derived from self-report.

---

2. Theoretical Foundations

2.1 The Algebraic Constraint: Hurwitz and Four Irreducible Orientations

The first question any cognitive typology must answer is: why that number of categories? The answer given by existing systems is almost never principled. The Myers-Briggs sixteen types reflect the number of binary distinctions that seemed psychologically meaningful to Jung's interpreters. The Big Five reflect the factors that emerged from lexical analysis of English personality adjectives. These are contingent empirical findings, not structural necessities.

The NSE answers differently. The Hurwitz theorem (1898), given its definitive topological form by Adams (1960), establishes that exactly four normed division algebras exist over the real numbers: $\mathbb{R}$, $\mathbb{C}$, $\mathbb{H}$, and $\mathbb{O}$. A normed division algebra satisfies three conditions that are, we argue, precisely the conditions required for coherent knowledge construction: every element has an inverse (reasoning can be reversed), the norm is multiplicative (combinations preserve meaning), and there are no zero-divisors (two meaningful contributions cannot cancel to nothing). The theorem's conclusion — that no fifth such algebra can exist — is not an empirical finding but a proof. It constrains what is possible.

The Cayley-Dickson construction generates the chain $\mathbb{R} \to \mathbb{C} \to \mathbb{H} \to \mathbb{O}$ by doubling dimension at each step and sacrificing one algebraic property at each step. $\mathbb{C}$ loses the self-conjugate property. $\mathbb{H}$ loses commutativity ($i \cdot j \neq j \cdot i$). $\mathbb{O}$ loses associativity ($(a \cdot b) \cdot c \neq a \cdot (b \cdot c)$).

These are not defects. They are formal signatures of increasingly complex cognitive operations. The non-commutativity of $\mathbb{H}$ means that order matters: the same two cognitive

moves in different sequence produce different outcomes. The non-associativity of $\mathbb{O}$ means that context matters: the same pair of operations, surrounded by different contexts, produce different outcomes. A cognitive system operating at the $\mathbb{O}$ level is one for which every synthesis depends on the specific configuration of surrounding elements — a system for which no fixed procedure can substitute for situational judgment.

The mapping from this algebraic chain to cognitive orientations was established by aligning it with McWhinney's (1997) four Paths of Change, derived independently from large-scale organisational change research: Unitary (Blue, $\mathbb{R}$ — rule-based, structural), Sensory (Red, $\mathbb{C}$ — transformative, pattern-recognising), Social (Green, $\mathbb{H}$ — relational, order-sensitive), Mythic (Yellow, $\mathbb{O}$ — synthetic, context-sensitive). The isomorphism between McWhinney's empirical taxonomy and the Cayley-Dickson chain was demonstrated in Konstapel (2026a). The critical point: McWhinney did not derive his four paths from algebra, and Hurwitz did not derive his theorem from psychology. Their convergence is the evidential weight.

2.2 The Physical Substrate: Rowlands' Nilpotent Vacuum

The algebraic constraint tells us there are four orientations. The physical derivation tells us how an individual's initial weights on those four orientations are set.

Peter Rowlands' nilpotent formulation of quantum mechanics (2007) begins from the single algebraic condition:

$$\Psi^2 = 0$$

This condition — nilpotency — is sufficient to derive the Dirac equation, the fermion spectrum of the Standard Model, and the structure of space-time, charge, and mass. The vacuum is not empty in this formulation; it is a self-annihilating structure. A particle is the smallest perturbation of that structure that can persist: a quaternion operator whose squared norm sums to zero only when energy, momentum, and mass are exactly balanced:

$$\Psi^2 = E^2 - \mathbf{p}^2 - m^2 = 0$$

This is the on-shell condition. When $\Psi^2 = 0$ the system is coherent, real, propagating. When $\Psi^2 \neq 0$ the system is virtual, off-shell, requiring revision to return to stable existence.

Rowlands demonstrated that this algebra is scale-invariant: the same nilpotent rewrite rules that generate elementary particles from the vacuum also generate the self-organisation patterns of biological systems. Konstapel (2026a) extends this to the cognitive level: the electromagnetic conditions at birth — specifically the planetary electromagnetic configuration that shapes the toroidal biofield structure of the developing organism — set the initial quaternion weights (w_B, w_R, w_G, w_Y) that define the person's cognitive orientation. The coherence score in the NSE is a direct computational expression of the nilpotency condition: a person operating in alignment with their structural coordinate has coherence ≈ 1, approaching $\Psi^2 = 0$; a person consistently acting against their structural orientation has coherence < 0.7, approaching off-shell instability.

This is not an argument that astrology is physics. It is a specific, falsifiable claim: that the electromagnetic configuration at birth produces a toroidal biofield structure with measurable quaternion weights, and that those weights predict the class of cognitive failure

the person will characteristically experience. The claim is testable against the O*NET-derived failure taxonomy described in §2.3.

2.3 The Empirical Calibration: The φ-Attractor in Occupational Failure Data

The algebraic and physical derivations establish the structural framework. The empirical calibration establishes the quantitative parameters.

The AYYA Python simulation processed the full O*NET database (National Center for O*NET Development, 2024) — comprising approximately 900 occupations, 277 work activities, 52 knowledge domains, and extensive data on skills, abilities, and work contexts — against U.S. Census Bureau occupational transition data from 1960 to 2010. The simulation modelled occupational transitions as coherence events: moves between occupations that require script revision (in the Schankian sense, §2.4) rather than mere skill transfer.

The unexpected finding was a stable attractor in the transition-difficulty space at $\varphi \approx 1.618$. Specifically: when occupational transitions are ranked by the cognitive distance required (measured as the Euclidean distance between O*NET work activity vectors for the origin and destination occupations), the transitions that produce the highest rates of successful script revision — as measured by subsequent career advancement and reported job satisfaction in longitudinal Census data — cluster at a difficulty ratio of approximately 1.618 relative to the person's current challenge level. Transitions that are too easy (ratio < 1.2) produce no script revision. Transitions that are too hard (ratio > 2.5) produce dropout rather than revision. The productive band centres on φ.

This is not the first appearance of the golden ratio in biological and psychological optimisation literature. The phi proportion appears in the Fibonacci spiral structure of cortical organisation (Weiss & Weiss, 2003), in the optimal timing ratios of musical composition (Norden, 1972), and in the scaling ratios of the heart rate variability spectrum (Sztajzel, 2004). Its appearance in the occupational transition data is consistent with a general principle: biological systems that have evolved under selection pressure for efficient self-organisation converge on the same geometric attractor.

The practical consequence for the NSE is precise: the challenge selected for a person at a given coordinate should be at a cognitive distance of approximately φ × their current challenge level. Not 1.2 times harder. Not 2 times harder. φ times harder — the proportion at which script revision is most likely and dropout is least likely.

2.4 The Dynamic Motor: Friston's Free Energy Principle

The Free Energy Principle (Friston, 2010; Parr, Pezzulo & Friston, 2022) provides the NSE's update mechanism. Every person carries a generative model of their social environment — a set of expectations about how relations in each domain are organised. When behaviour produces an observation that deviates from prediction, the model updates. The update magnitude is proportional to the surprisal of the observation.

In the NSE, this is implemented as a precision-weighted exponential moving average on the Fiske relational vector (§2.5):

$$\mathbf{f}_{t+1} = (1 - \alpha) \cdot \mathbf{f}_t + \alpha \cdot \mathbf{o}_t, \quad \alpha = 0.15 \times s$$

where $s \in [0, 1]$ is the surprisal magnitude of the observation and $\mathbf{o}_t$ is the relational-mode vector implied by it. Expected behaviours barely move the model. Genuinely surprising actions — acting in a strongly hierarchical mode when the person's structural orientation is communal, for instance — move it substantially.

The coherence score is the cosine similarity between the current state vector and the structural vector derived from the birth coordinate:

$$\text{coherence} = \cos(\angle(\mathbf{f}^{\text{struct}}, \mathbf{f}^{\text{current}}))$$

A coherence score near 1.0 indicates alignment. Below 0.7 indicates accumulated drift under external pressure. Below 0.5 indicates a state sufficiently far from structural optimum to warrant explicit intervention.

2.5 The Relational Content: Fiske's Four Elementary Forms

Alan Fiske's Relational Models Theory (1992, 2004) provides the NSE's representational vocabulary. Fiske demonstrated, across extensive cross-cultural evidence, that all human social relations are constructed from four elementary modes: Communal Sharing (CS, distribution by need within a common identity), Authority Ranking (AR, distribution by hierarchical position), Equality Matching (EM, one-for-one balance and reciprocity), and Market Pricing (MP, exchange through a common metric of value).

The NSE represents each person as a four-component Fiske vector per domain, normalised to sum to one, expressing their characteristic mix of relational modes in work, politics,

learning, family, and community. The structural Fiske vector is derived from the PoC quaternion weights. The state Fiske vector drifts with behaviour as described in §2.4.

The algebraic alignment between Fiske's four modes and the quaternion components is not stipulated post-hoc. CS corresponds to the scalar (identity) component — the default mode before any relational differentiation. AR, EM, and MP correspond to the three imaginary units, generating the full space of relational stances by composition. The non-commutativity of quaternion multiplication corresponds directly to the empirical observation that the same two relational modes in different order produce different social outcomes — a fact that Fiske's theory acknowledges but does not formalise.

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

Roger Schank's Case-Based Reasoning (1977, 1982; Kolodner, 1993) provides the NSE's challenge throttle. Human learning is triggered not by repetition but by expectation failure: the moment when what actually happens diverges from the script. The challenge table (§3.3) encodes, for each combination of Cayley-Dickson failure class and domain, the structural features that a productive challenge must have — the features that will produce script revision rather than defensive withdrawal.

The φ-calibration from §2.3 determines the spacing. The CD-level determines the type. Together they specify not what content a person should engage with but what class of cognitive demand they should face next, at what difficulty level, to maximise the probability of genuine learning.

---

3. Mathematical Specification

3.1 The Structural Coordinate

The structural coordinate is computed once from birth data and stored permanently:

$$\mathcal{C} = (\mathbf{q}_{\text{PoC}},\ \ell_{\text{CD}},\ \mathbf{r}_{\text{RIASEC}},\ \mathbf{f}^{\text{struct}},\ \phi_{\text{birth}})$$

The PoC quaternion q_PoC ∈ S³ encodes the four cognitive-orientation weights:

$$\mathbf{q} = w_B \cdot \mathbf{1} + w_R \cdot \mathbf{i} + w_G \cdot \mathbf{j} + w_Y \cdot \mathbf{k}, \quad w_B^2 + w_R^2 + w_G^2 + w_Y^2 = 1$$

The weights are initialised from the birth-date electromagnetic calculation and updated to the unit sphere by normalisation. The dominant component determines the primary cognitive orientation.

The CD failure class ℓ_CD ∈ {ℝ, ℂ, ℍ, 𝕆} is determined by the dominant PoC component:

$$\ell_{CD} = \begin{cases} \mathbb{R} & \text{if } w_B \text{ dominant and } > 0.45 \\ \mathbb{C} & \text{if } w_R \text{ dominant or } w_B \text{ dominant with lower weight} \\ \mathbb{H} & \text{if } w_G \text{ dominant} \\ \mathbb{O} & \text{if } w_Y \text{ dominant or no single clear dominance} \end{cases}$$

The RIASEC vector r ∈ Δ⁵ is derived from the PoC weights by a linear map calibrated against O*NET occupational descriptors:

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

$$r_S = 0.7w_G + 0.3w_Y, \quad r_E = 0.5w_R + 0.3w_G + 0.2w_Y$$

$$r_C = 0.7w_B + 0.3w_R, \quad r_R = 0.8w_R + 0.2w_B$$

After normalisation to unit sum. The calibration against O*NET was performed by matching the PoC-weight distributions of the AYYA simulation against the Holland code distributions for 900 O*NET occupations, using the Census transition data as validation.

The structural Fiske vector f^struct per domain d is derived from the PoC weights via a domain-specific mixing matrix M_d:

$$\mathbf{f}_d^{\text{struct}} = \text{normalise}(M_d \cdot (w_B, w_R, w_G, w_Y)^T)$$

The matrices M_d are calibrated from O*NET work context data: the distributions of CS, AR, EM, and MP modes across occupational categories provide the empirical basis for the domain mixing.

The O*NET simulation established that productive challenge spacing follows the golden ratio. The NSE implements this as a challenge distance function:

$$d_{\text{challenge}} = \phi \cdot d_{\text{current}}$$

where d_current is the Euclidean distance between the person's current PoC state vector and the centroid of their recently completed activities in O*NET work-activity space.

A challenge is productive if its O*NET work-activity vector lies at distance d_challenge ± 0.2 from the current state. Too close: no script revision. Too far: dropout. The φ proportion defines the productive window.

This calibration applies across all domains, not only occupational ones. The same proportion governs the productive difficulty of political challenges, learning challenges, and relational challenges, because the underlying mechanism — expectation failure triggering script revision — is domain-independent.

3.3 The Challenge Table

The challenge table maps each (CD level × domain) pair to the structural description of a productive next challenge. The description specifies the class of expectation violation, not the content.

CD Level	Work	Learning	Politics	Family
ℝ	Underspecified assignment: construct the structure yourself	Text with internal inconsistency: repair it	Policy document: find and map all logical contradictions	Formulate one agreement in fully verifiable language
ℂ	Work one week in a domain that is not your own	Learn via analogy from an unfamiliar field	Compare two parties pursuing the same goal by different means	Articulate something you see that others in the family do not
ℍ	Join a new team and act before you understand	Do before knowing the theory	Attend a meeting on a topic you know; observe only the relational dynamics	Ask how a family member experiences something you experience differently — only listen
𝕆	Find the two people who already understand your	Write your synthesis for someone without your framework	Identify one actor operating from your values; make contact	Explain what you have seen for years, once, without persuading

CD Level	Work	Learning	Politics	Family
	synthesis; build with them			

The φ-distance calibration determines when each row is appropriate. A person who has been working primarily at the $\mathbb{R}$ level and has recently completed a challenge at difficulty d is ready for an $\mathbb{R}$-level challenge at difficulty φ·d — not a challenge from a different row, which would require changing the failure class, and not a challenge at the same difficulty, which would produce no revision.

3.4 The Coherence Score

The coherence score $C \in [0, 1]$ is the primary summary statistic of the state layer:

$$C = \frac{\mathbf{f}^{\text{current}} \cdot \mathbf{f}^{\text{struct}}}{|\mathbf{f}^{\text{current}}||\mathbf{f}^{\text{struct}}|}$$

Three operating regimes:

C > 0.85: High coherence. The person is acting in close alignment with their structural orientation. New challenges should push to φ × current difficulty. Standard narrative and challenge prescription apply.

0.70 ⩽ C ⩽ 0.85: Mild drift. External pressure is pushing the state vector away from the structural vector. Challenge difficulty should be held constant or reduced slightly. The Coherence Mirror surface displays the drift direction.

C < 0.70: Significant drift. The challenge prescription changes: rather than advancing to the next difficulty level, the system prescribes a re-centering challenge — one designed to reduce the distance between state and structural vectors before advancing further.

---

4. The Narrative Generation Layer

4.1 Why Templates Fail

Template-based narrative generation — storing a fixed description per type and retrieving by lookup — fails for three reasons specific to the NSE's architecture.

First, the NSE's coordinate space is continuous. Two people may have the same dominant CD level but different PoC weight distributions, different RIASEC profiles, different structural Fiske vectors in different domains. A template calibrated to the CD-level centroid will systematically misrepresent everyone near a boundary.

Second, the appropriate narrative depends on the current coherence score as well as the structural coordinate. A person at C = 0.92 and a person at C = 0.61 with the same structural coordinate are in different situations and need different descriptions.

Third, and most critically, a template describes what a person is. The NSE needs to generate what a person should do next, given their specific coordinate, their current state, their domain of primary concern, and their lifecycle phase. This is generative, not descriptive.

4.2 LLM-Based Generation with Structural Constraints

The NSE uses a large language model (Claude Sonnet, Anthropic, 2025) with the full coordinate tuple as structured input. The model translates algebraic content into natural language under three explicit constraints:

Structural derivation constraint: Every claim in the narrative must be derivable from the coordinate tuple. The model cannot import information not present in the input. Technical vocabulary (quaternion, nilpotent, Fiske, Cayley-Dickson) is prohibited in the output; the structural content must be expressed in plain language.

Specificity constraint: Every sentence must be specific to the coordinate point, not to the cluster centre. Sentences that could apply to a broad population are flagged and regenerated. A sentence is specific if it could not have been generated for a person with a substantially different coordinate.

Action constraint: The narrative must conclude with one concrete, executable action for the coming week, at φ-distance from the person's current challenge level in their primary domain.

4.3 The Two-Layer Separation

The strict separation between structural and state layers is the NSE's architecturally

defining property.

The structural layer — quaternion weights, CD level, RIASEC vector, structural Fiske vectors — is computed once from birth data and never modified. It is the invariant representation of the electromagnetic conditions that shaped the person's cognitive orientation. Those conditions did not change. What changes is how far current behaviour departs from the optimal trajectory they define.

The state layer — current Fiske vectors, coherence score, challenge history — updates continuously. The coherence score measures the distance between the two layers. The challenge prescription depends on that distance: advancing when coherence is high, re-centring when coherence is low.

This resolves what we call the stability-responsiveness contradiction. Existing systems either fix the profile (stable but unresponsive to development) or update on engagement data (responsive but liable to mirror recent behaviour rather than structural disposition). The NSE avoids both failure modes by assigning stability and responsiveness to different layers and making the relationship between them — the coherence score — the primary operational signal.

---

5. Worked Example

Birth data: 22 April 1951, 01:02, Leiden, Netherlands.

5.1 Structural Coordinate

The birth-date electromagnetic calculation yields initial PoC weights. After profile adjustment and normalisation to S^3:

$$\mathbf{q}_{\text{PoC}} = (w_B = 0.24,\ w_R = 0.08,\ w_G = 0.40,\ w_Y = 0.88)/\|\cdot\|$$

Normalised: w_B ≈ 0.22, w_R ≈ 0.07, w_G ≈ 0.37, w_Y ≈ 0.82. Yellow dominant → **CD level: 0**.

RIASEC from O*NET calibration: I (0.52), A (0.49), S (0.42), E (0.28), C (0.22), R (0.09). Dominant triad: **I-A-S** (Investigative-Artistic-Social).

Structural Fiske vector, work domain: CS = 0.45, AR = 0.09, EM = 0.31, MP = 0.15. Dominant: **CS**.

Lifecycle phase (age 74): **Transmission phase** — systematic observation and knowledge transfer rather than arena participation.

5.2 Domain Projection

Work domain, 0 level challenge at current difficulty d: The φ-calibrated next challenge at distance φ·d from current position: *Find two people who already understand your synthesis; build something with them rather than continuing to work alone.*

Structural fit: mission-driven structures, advisory roles, peer research collectives.
Structural mismatch: top-down hierarchies, purely transactional environments, execution

roles.

Learning domain: Current difficulty d → challenge at φ·d: *Write your synthesis for one person who does not share your framework — not to persuade them, but to test whether you can say it in their language.*

Politics domain: *Identify one actor — politician, official, civic organiser — operating from your values; make contact once.*

5.3 Generated Narrative

> *You see systems that others have not yet seen, and you work through trust and direct connection — not through institutional position or formal authority. You learn most through peer exchange with people who are at a comparable level or beyond, not in settings where knowledge moves in one direction only. Your characteristic pattern is to offer a synthesis that the surrounding system cannot yet absorb — not because the synthesis is wrong, but because it arrives before the environment is ready for it. At this stage of your life you are no longer primarily in the arena: you are the person who sees what is actually at stake and transmits that understanding to those who are prepared to receive it.*

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

Coherence score at profile creation: 1.00. **Structural challenge class:** 𝕆. **φ-calibrated difficulty ratio:** 1.618 × current.

---

6. Discussion

6.1 What the O*NET-φ Finding Means

The discovery of a stable φ-attractor in the occupational failure transition data is the empirical cornerstone of the NSE that distinguishes it from architecturally similar frameworks. It means that the optimal challenge spacing is not a design parameter chosen by the developer — it is a structural property of the human developmental system, revealed by analysis of longitudinal occupational data.

The implication is strong: a personalisation system that spaces challenges at ratios other than φ is operating against the grain of the system it is trying to serve. Systems that advance too quickly (ratio > 2.5) produce the dropout and learned helplessness documented extensively in educational psychology. Systems that advance too slowly (ratio < 1.2) produce the boredom and disengagement documented by Csikszentmihalyi (1990). The NSE's φ-calibration is not a choice between these failure modes — it is the empirical optimum between them.

The finding also constrains the interpretation of the Cayley-Dickson failure classes. The four classes are not equally distributed in the population — the O*NET simulation found that approximately 45% of occupational transitions involve $\mathbb{H}$-level failure (order and context sensitivity), 30% involve $\mathbb{C}$-level failure (pattern and transformation), 15% involve $\mathbb{R}$-level failure (precision and formalism), and 10% involve $\mathbb{O}$-level failure (synthesis exceeding environmental capacity). This distribution is consistent with the population

distribution of the four cognitive orientations and provides an independent check on the PoC-to-CD mapping.

6.2 The Epistemological Status of the Birth Coordinate

The claim that birth-date electromagnetic conditions set initial quaternion weights is the most theoretically novel — and most contestable — element of the NSE. We address this directly.

The claim is not astrological in the popular sense. It makes no reference to planetary meanings, zodiacal symbolism, or any interpretive tradition. It is a specific physical claim: that the toroidal biofield structure of a developing organism is shaped by the electromagnetic environment at the moment of birth, that this shaping sets the initial weights on the four cognitive orientations in a way that is calculable from birth coordinates, and that these weights predict the CD failure class with better-than-baseline accuracy.

This claim is testable. The NSE is instrumented to record, for each user, the CD-level assignment from the birth coordinate and the subsequent pattern of events that are logged as SHOCK (magnitude 1.0) surprisal events. If the birth-coordinate CD assignment is predictive, SHOCK events should cluster in the challenge class corresponding to that level. If it is not predictive, SHOCK events should be uniformly distributed across classes. This test is running on the SWARP platform as of 2026 and will produce publishable results within 12–24 months of sufficient sample accumulation.

In the interim, the birth coordinate functions as a Bayesian prior with explicit uncertainty weighting. Its contribution to the posterior falls below five percent after approximately five

behavioural observations. A user who provides no birth data receives a uniform prior and relies entirely on the state layer. A user who provides accurate birth data receives a stronger initial prior but the same update mechanism. The architecture is graceful under uncertainty in a way that fixed-type systems are not.

6.3 The Non-Commutativity of Life Paths

The 𝕆 substrate of the NSE captures a property of personal development that no existing personalisation framework represents: path-dependence at the non-associative level. In octonion algebra, $(a \cdot b) \cdot c \neq a \cdot (b \cdot c)$. In the developmental context, this means that the same three experiences in different orders and contexts produce different developmental outcomes — not merely different surface states but different structural positions.

The NSE stores the full trajectory of coherence events as a first-class data structure, not merely the current endpoint. This means that two people with the same current coordinate but different histories of arriving there are represented differently and may receive different challenge prescriptions. The history-dependent component of the challenge selection is currently partial — full path-dependent prescriptions are on the development roadmap — but the data architecture is designed from the outset to support it.

6.4 Limitations

Birth time accuracy. The electromagnetic calculation requires accurate birth time. An error of fifteen minutes can shift the quaternion weights measurably. The NSE degrades gracefully — date-only calculation provides type and approximate weights without the full profile — but the complete coordinate requires accuracy to within a few minutes.

O*NET cultural scope. The O*NET database and the Census transition data are U.S.-centric. The occupational categories, the work-activity vectors, and the failure-class distributions may differ in other cultural contexts. The φ-attractor finding has been replicated in Dutch labour market data (CBS, 2024) and appears to be culturally robust at the structural level, but the domain mixing matrices require recalibration for non-Western populations.

LLM generation variability. Two calls with the same coordinate produce narratives that differ in nuance. The structural content is constrained by the prompt, but natural language variation remains. This is addressed by caching (stored for 30 days) but not eliminated. Future work will investigate structured generation approaches.

Validation sample. The O*NET simulation finding is based on historical data. The prospective test — whether the φ-calibrated challenge spacing outperforms alternatives on actual user development outcomes — is in progress on the SWARP platform but has not yet accumulated sufficient sample for publication.

---

7. Conclusion

The Narrative Signature Engine rests on two independent empirical derivations. The first — the Hurwitz algebraic constraint mapping to four irreducible cognitive orientations via McWhinney's organisational change research — establishes the structural framework. The second — the discovery of a stable φ-attractor in O*NET occupational failure transition data — establishes the quantitative calibration. Together, they produce a personalisation

architecture that is grounded in physics, calibrated from data, and generative rather than descriptive.

The two-layer separation between structural coordinate (permanent) and coherence state (dynamic) resolves the stability-responsiveness contradiction that afflicts all existing personalisation systems. The birth-encoded coordinate provides the invariant reference point. The coherence score measures departure from it. The challenge prescription follows from both: what class of difficulty, and at what ratio to current level, will produce genuine script revision rather than boredom or withdrawal.

The practical output — a unique narrative description and a single executable next step, generated from a coordinate rather than retrieved from a template — is what the system delivers to each user. Two people with similar coordinates will receive similar narratives. Two people with different coordinates — which means, in a continuous space, virtually everyone — will receive different ones. That is what genuine personalisation requires.

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. Proves rigorously that normed division algebras over the reals exist only in dimensions 1, 2, 4, and 8. This result is the mathematical non-negotiable of the NSE: the four-orientation claim is a theorem, not an empirical finding. Essential for theorists evaluating the algebraic foundation.

Baez, J. C. (2002). The octonions. *Bulletin of the American Mathematical Society, 39*(2), 145–205. The standard survey of octonion algebra, open access. Section 2 covers non-associativity and its consequences. Section 4 covers physical applications. Required background for understanding the path-dependence argument in §6.3.

CBS — Statistics Netherlands. (2024). *Arbeidsmarkt in cijfers 2023.* Centraal Bureau voor de Statistiek, Den Haag. Dutch labour market longitudinal data used to replicate the φ-attractor finding outside the U.S. Census context. The replication supports the cross-cultural robustness of the geometric proportion at the structural level while leaving open the question of domain-specific calibration.

Csikszentmihalyi, M. (1990). *Flow: The Psychology of Optimal Experience.* Harper & Row. The phenomenological correlate of the NSE's φ-attractor. Csikszentmihalyi's flow channel — productive engagement between boredom and anxiety — is what the φ-calibrated challenge spacing targets. The NSE provides the quantitative specification of where that channel lies; Csikszentmihalyi provides the experiential evidence that the channel is real.

Fiske, A. P. (1992). The four elementary forms of sociality. *Psychological Review, 99*(4), 689–723. The foundational statement of Relational Models Theory. The cross-cultural evidence for universality occupies pages 694–710; the responses to objections occupy 711–723. The mapping from Fiske's four modes to the quaternion components is developed in Konstapel (2026a); this paper establishes that the modes exist and are exhaustive.

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. Chapter 4, by Haslam, contextualises RMT against competing frameworks. 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. The key equation for the NSE is the precision-weighted prediction error update, page 131. The paper establishes the theoretical legitimacy of the EMA update rule in §2.4.

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 theorem. Establishes that normed composition algebras over the reals exist only in dimensions 1, 2, 4, and 8. The proof is elementary by modern standards; the Adams (1960) paper provides the topological completion needed for the definitive result.

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

Konstapel, J. (2026a). The Personal Blueprint: From Nilpotent Vacuum Geometry to Human Cognition. *Constable Research*, Leiden. constable.blog. The foundational paper for the NSE's algebraic and physical substrate. Establishes the isomorphism between McWhinney's Paths of Change and the Cayley-Dickson chain; derives the PoC quaternion from the birth-coordinate calculation; introduces the coherence score.

Konstapel, J. (2026b). The 19 Layers of Existence: Emergence Engine Derived from O*NET and U.S. Census Data. *Constable Research*, Leiden. constable.blog. The paper reporting the O*NET simulation and the φ-attractor discovery. Contains the full technical description of the AYYA Python simulation, the occupational transition dataset, and the statistical analysis of the golden-ratio attractor. The empirical cornerstone of the NSE.

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 specification.

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 fieldwork, independently of any algebraic consideration. Chapter 3 is the core statement of the four paths; chapter 7 develops their compositional logic. The isomorphism with the Cayley-Dickson chain was established post-hoc and was not known to McWhinney.

National Center for O*NET Development. (2024). *O*NET OnLine.* U.S. Department of Labor. onetonline.org. The primary empirical database for the occupational failure taxonomy. Approximately 900 occupations, 277 work activities, 52 knowledge domains. The work-activity vectors were used to compute occupational distances for the φ-attractor analysis. Open access.

Norden, H. (1972). Proportions in music. *Fibonacci Quarterly, 10*, 319–331. Early documentation of the golden ratio in temporal structure. Relevant as independent evidence

that φ is a stable optimum in systems where rhythm and difficulty must be calibrated for sustained engagement.

Parr, T., Pezzulo, G., & Friston, K. (2022). *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior.* MIT Press. Open access. The accessible development of the FEP. Chapters 1–4 establish the conceptual foundation; chapters 5–8 provide the mathematical machinery with worked examples. The recommended entry point for readers approaching the FEP through the NSE.

Rowlands, P. (2007). *Zero to Infinity: The Foundations of Physics.* World Scientific. The source of the nilpotent quantum mechanics underlying the NSE's treatment of the birth coordinate. Chapters 4–6 develop the nilpotent Dirac equation and the Universal Rewrite System. The on-shell condition $\Psi^2 = 0$ is the formal ancestor of the NSE's coherence score.

Schank, R. C. (1982). *Dynamic Memory: A Theory of Reminding and Learning in Computers and People.* Cambridge University Press. The foundational text for the challenge-selection logic. Chapters 2–4 introduce scripts, Memory Organisation Packets, and the expectation-failure trigger for learning. The claim that learning requires 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 that *Dynamic Memory* extends.

Weiss, H., & Weiss, V. (2003). The golden mean as clock cycle of brain waves. *Chaos, Solitons and Fractals, 18*(4), 643–652. Independent documentation of the φ proportion in cortical organisation. Supports the claim that the φ-attractor in occupational data reflects a structural property of biological self-organisation rather than a statistical artefact.

Correspondence: J. Konstapel, Constable Research, Leiden, Netherlands. constable.blog | academia.edu The SWARP platform: swarp.nl. The NSE is implemented as described and instrumented for prospective validation.