SWARP-Agora: A Fractal Wisdom Simulator for Collective Reflection

A Critical Comparison with Contemporary Collective Sensemaking and Deliberation Systems

J. Konstapel Constable Research, Leiden, The Netherlands Leiden, March 28, 2026

Abstract

This paper introduces SWARP-Agora, a fractal adaptive simulation platform designed to address a structural deficit in contemporary democratic systems: the systematic failure to achieve cognitive closure following political expectation failure. Grounded in the author's Political Expectation Failure Theory (PEFT), the platform integrates Paths of Change (PoC), panarchy theory, Active Inference, Markov navigation, Human Design filtering, and a Socratic AI facilitation layer to enable script revision and wisdom ripening at collective scale. A conceptual comparative analysis is conducted across nine established systems and methodologies: SenseMaker and Estuarine Mapping (Snowden), Kumu, the Consilience Project and Game B ecosystem (Schmachtenberger), multi-agent Active Inference simulators (Friston et al.), deliberation platforms (Polis, Loomio, Remesh), the Collective Intelligence Project (CIP), Integral Life and the AQAL framework (Wilber), DebateGraph and related argument-mapping tools, Appreciative Inquiry (Cooperrider), and the Wisdom

Council and Dynamic Facilitation process (Rough). While each system offers substantive contributions to collective intelligence, none integrates expectation-failure mechanics, real-time panarchic rhythm detection (Revolt and Remember functions), and a dedicated wisdom-ripening sub-simulator. SWARP-Agora is positioned as a theoretically distinct architecture: a living laboratory for democratic reflection that directly addresses the structural non-learning identified in PEFT. As no broad empirical validation of SWARP-Agora has yet been conducted, this paper is explicitly a theoretical and conceptual contribution; empirical validation is identified as a critical priority for subsequent research. The paper concludes with seven concrete recommendations for the platform's further development.

Keywords: collective intelligence, expectation failure, wisdom simulation, panarchy, Socratic AI, Active Inference, sensemaking, Human Design, Paths of Change, AQAL, Dynamic Facilitation, argument mapping, Political Expectation Failure Theory

---

1. Introduction

Contemporary democracies exhibit a recurring and structurally undertheorised pathology: the systematic generation, violation, and political exploitation of citizen expectations without achieving the cognitive or institutional resolution that expectation failure should, in principle, trigger. The author terms this condition Political Expectation Failure (PEFT) (Konstapel, 2026). Drawing on Schank and Abelson's (1977) script theory, PEFT argues that expectation failure, when properly processed, initiates a learning sequence of explanation, reminding, generalisation, and script revision. In practice, however, democratic systems

consistently truncate this sequence at the stage of political mobilisation. The result is a cycle of non-resolution driven by what De Vries and Hobolt (2020) identify as political entrepreneurs — actors who structurally profit from exploiting the expectation-performance gap rather than closing it.

This diagnostic convergence from three independent theoretical traditions motivates the development of SWARP-Agora. First, Schank's case-based learning and script theory establishes the cognitive conditions under which genuine belief revision occurs. Second, De Vries's (2018) empirical research on Euroscepticism and political dissatisfaction documents the structural conditions under which such revision is systematically avoided in democratic politics. Third, Gunderson and Holling's (2002) panarchy framework demonstrates that healthy adaptive cycles require both Revolt — bottom-up disruption that destabilises over-consolidated structures — and Remember — top-down wisdom injection that prevents reorganisation from degenerating into mere recycling of existing patterns. Contemporary democratic systems, on this account, are characterised by arrested Revolt and atrophied Remember.

SWARP (Swarm Spatial Adaptive Reasoning Platform) was developed as a practical architectural response to this diagnosis: a fractal simulator that embeds PoC worldviews, panarchic adaptive cycles, and Active Inference dynamics to enable genuine collective learning at scale. Within SWARP, Agora functions as a protected sub-simulator dedicated to wisdom ripening — a deliberate institutional slow space in which experienced individuals engage in Socratic dialogue aimed at achieving the cognitive closure that democratic systems currently structurally prevent.

This paper proceeds as follows. Section 2 describes the methodology of the comparative analysis. Section 3 presents the core architecture of SWARP-Agora. Section 4 provides a systematic critical comparison with nine existing systems. Section 5 offers a synoptic comparative summary. Section 6 assesses the strengths and limitations of SWARP-Agora as a theoretical contribution. Section 7 provides concrete recommendations for future development. Section 8 concludes.

---

2. Methodology

This study employs conceptual comparative analysis as its primary method (Schedler & Santiso, 1998; Gerring, 2012). The approach is appropriate given that SWARP-Agora is at a pre-empirical stage of development and the aim of the comparison is diagnostic and architectural rather than evaluative. Three sources of evidence inform the analysis.

Systematic literature and documentation review. Core theoretical and technical documentation for each compared system was reviewed, drawing on primary sources from the respective developers and institutions: SenseMaker and Estuarine Mapping documentation (Snowden, 2022–2025; Snowden & Boone, 2007), Kumu platform documentation, Consilience Project publications (Schmachtenberger, 2020–2023), Active Inference Institute publications (Friston et al., 2022; Parr et al., 2022), deliberation platform technical papers (Small et al., 2021), CIP whitepapers and annual reports (Siddarth & Huang, 2023; CIP, 2024), Integral Institute publications (Wilber, 2006), DebateGraph documentation, Appreciative Inquiry foundational texts (Cooperrider & Whitney, 2005), and Center for Wise Democracy materials (Rough, 2002).

Functional decomposition. Each system is evaluated across five criteria derived directly from PEFT and panarchy theory: (1) capacity for expectation detection; (2) provision of a mechanism for script revision and cognitive closure; (3) operationalisation of panarchic rhythm detection, specifically Revolt and Remember functions; (4) scalability across fractal organisational levels; and (5) provision of a dedicated wisdom-ripening layer. These criteria are not arbitrary but follow directly from the structural deficits identified in the PEFT diagnosis.

Author's developmental experience. The analysis draws on the author's direct involvement in designing and building SWARP-Agora from 2019 to 2026, including architectural decisions, Human Design integration, Socratic AI specification, and the theoretical integration of McWhinney's Paths of Change with the Grammars of Engagement framework (McWhinney & Konstapel, unpublished manuscript). This proximity to the object of study is acknowledged as a source of potential bias and is addressed through systematic attention to SWARP-Agora's own limitations in Section 6.2.

The analysis is explicitly framed as a theoretical contribution. No empirical user studies of SWARP-Agora at scale have yet been conducted; the comparative judgements offered here concern architectural design and theoretical completeness, not demonstrated outcomes. Empirical validation is identified in Section 7.6 as a critical and urgent next step.

---

3. Core Architecture of SWARP-Agora

3.1 The SWARP Platform

SWARP operates as a large-scale PoC-driven panarchic adaptive cycle, designed to model and simulate collective sensemaking across populations ranging from thousands to millions of participants. Each participant is represented as an autonomous predictive agent whose coupling behaviour, reflection depth, and wisdom potential are filtered through their Human Design profile — specifically, type, authority, and profile line. Navigation across the shared semantic space occurs via Markov chains over a bounded lexicon of approximately 175 interconnected concepts; coupling between participants is measured using the Grammars of Engagement framework developed collaboratively with Will McWhinney (McWhinney & Konstapel, unpublished manuscript).

The theoretical architecture is grounded in McWhinney's (1992) Paths of Change framework, which identifies four fundamental worldview orientations — Unitary (Blue), Sensory (Red), Social (Green), and Mythic (Yellow) — through which individuals and organisations move in any genuine change process. In SWARP, these orientations are treated not as static categories but as coupled oscillators. The introvert/extravert axis determines rotational direction (centripetal versus centrifugal); the judging/perceiving axis provides an orthogonal dimension. Their dynamic interaction generates what the author terms the *anarchy spiral* — the emergent region of highest creative and transformative potential at the centre of the quadrant space, associated in McWhinney's work with deep transformation.

3.2 Agora as the Remember Layer

Agora constitutes the dedicated Remember layer within SWARP's panarchic architecture (following Gunderson & Holling, 2002). It functions as a slow, reflective sub-simulator anchored in the K-phase (Conservation and Deepening) of the adaptive cycle. Within Agora, participants with Human Design profiles indicating high wisdom potential — particularly those with a sixth line (Role Model profile), whose three-phase life cycle of experimentation (ages 0–30), withdrawal and ripening (ages 30–50), and embodied wisdom transmission (from approximately age 50) makes them structurally suited to the Remember function — produce long-form reflective blogs and engage in Socratic dialogue facilitated by a specialised AI agent designated *Socrates*.

The Socrates AI is designed to embody the Socratic method in its authentic form: sustained questioning, elenctic probing of assumptions, and maieutic facilitation of the participant's own insight, as opposed to the provision of information or direction. Output from Agora is injected into lower-scale α-phases (Reorganisation) of the wider SWARP simulator, enabling Revolt signals at the periphery to produce genuine script revision rather than merely triggering further cycles of failure exploitation.

3.3 Real-Time Rhythm Detection

Detection of Revolt and Remember functions is implemented through three core metrics, computed at 15-minute intervals across all fractal organisational layers:

- **Markov Path Similarity**: the degree of convergence in the conceptual trajectories that users follow through the shared lexicon, indicating either cognitive entrainment or

productive divergence.

- **Coupling Score**: the frequency and depth of co-navigation, blog response, and co-creation, operationalised via the Grammars of Engagement framework.
- **Free Energy Delta**: the magnitude of predictive surprise per user or group, serving as a real-time proxy for the intensity of expectation violation in the Active Inference sense.

A Revolt event is flagged when a sudden Ω-phase transition at a smaller organisational scale cascades upward within a 24–48 hour window. A Remember event is flagged when Agora output is measurably absorbed into α-phase activity — operationalised as a decline in Free Energy Delta and the initiation of new Markov paths originating within the Agora sub-simulator.

---

4. Comparative Analysis

4.1 SenseMaker and Estuarine Mapping (Snowden / The Cynefin Co.)

SenseMaker is the most mature distributed ethnographic sensemaking platform currently in operation. It collects micro-narratives from large populations and enables participants to self-interpret their own stories through triadic or dyadic signifier sets, producing vector fields that reveal emergent cultural and cognitive patterns without imposing researcher categories. Recent development has concentrated on the Estuarine Mapping framework (Snowden, 2022–2025), which maps the constraints and constructors of a system's

affordance landscape using the estuary as a guiding metaphor — a dynamic mixing zone between fixed and fluid conditions. Estuarine Mapping supports what Snowden describes as maieutic facilitation: assisting groups in perceiving what is already emergent rather than imposing predetermined solutions. MassSense applications extend this capacity to real-time, large-scale narrative sensing.

Strengths: Proven scalability, participant-led interpretation that reduces researcher imposition, and sophisticated grounding in complexity science and Cynefin epistemology.

Limitations relative to SWARP-Agora: SenseMaker excels at detection and pattern recognition but provides no built-in mechanism for script revision or cognitive closure in Schank's sense. There is no dedicated wisdom-ripening sub-system, no individual oscillator model, and no explicit operationalisation of panarchic Revolt and Remember as real-time dashboard metrics. Estuarine Mapping remains a diagnostic and facilitation framework; it does not simulate the adaptive cycle or actively inject wisdom products into reorganisation phases. Kumu — a related web-based visual mapping tool focused on relationship networks and causal loop diagrams — shares these limitations while offering strong structural mapping for existing systems. SWARP-Agora goes beyond both by constituting a generative simulator that actively closes the expectation-failure loop through Socratic AI dialogue, detects panarchic rhythms in real time, and applies Human Design as an individual oscillator filter.

4.2 The Consilience Project and Game B (Schmachtenberger et al.)

The Consilience Project and Game B ecosystem engage with meta-rationality, civilisational design, and the structural avoidance of multipolar coordination failures.

Schmachtenberger's systems diagnosis substantially converges with PEFT: existing institutional arrangements are structurally incapable of producing the collective learning and cognitive closure they require for long-term stability (Schmachtenberger, 2020–2023). Game B communities attempt to prototype alternative social forms grounded in genuine mutual sensemaking and collaborative epistemology.

Strengths: Profound systemic diagnosis operating at civilisational scale, and communities of practice attempting to embody alternative patterns in practice.

Limitations relative to SWARP-Agora: The Consilience Project and Game B ecosystem remain primarily discursive and community-based. No operational fractal simulator exists with real-time Markov navigation, panarchic rhythm monitoring, or individual oscillator filtering. There is no equivalent to Agora's Socratic AI layer or to Human Design as a precision filter for wisdom potential. These projects constitute valuable cultural and philosophical groundwork but do not, at present, constitute an executable architecture for collective wisdom simulation.

4.3 Multi-Agent Active Inference Simulators (Friston / Active Inference Institute)

Multi-agent Active Inference systems model collectives as ensembles of predictive agents that minimise variational free energy through action and perception (Friston et al., 2022; Parr et al., 2022). This framework is directly aligned with SWARP's own theoretical backbone. The Active Inference Institute has extended this to social cognition and collective dynamics, producing simulation models of how groups form and update shared generative models and how expectation violation propagates through networks of coupled agents.

Strengths: Unmatched mathematical rigour, direct integration with neuroscientific models of prediction and adaptive learning, and increasing applicability to multi-agent social dynamics.

Limitations relative to SWARP-Agora: Most implementations remain either purely computational or experimentally small-scale. They lack a protected wisdom-ripening space analogous to Agora, do not incorporate PoC worldviews or Human Design as individual oscillator filters, and do not operationalise Revolt and Remember as real-time participatory metrics. SWARP-Agora may be understood as an applied, participatory instantiation of multi-agent Active Inference — sharing the theoretical foundation while addressing the implementation and human-engagement challenges that pure simulation architectures have not yet resolved.

4.4 Deliberation Platforms: Polis, Loomio, and Remesh

Polis, Loomio, and Remesh represent the current state of the art in large-scale democratic deliberation technology. Polis applies machine learning clustering to identify areas of consensus and disagreement across large populations in real time (Small et al., 2021). Loomio provides structured asynchronous deliberation with explicit decision-making protocols. Remesh facilitates live large-group dialogue with AI-assisted theme detection.

Strengths: Proven accessibility, deployment in real political and organisational contexts, and genuine advances beyond simple preference aggregation toward structured deliberation.

Limitations relative to SWARP-Agora: All three platforms operate at the level of preference and opinion rather than at the level of the generative scripts and expectation architectures that produce those preferences. None addresses the upstream cognitive structures that PEFT identifies as the locus of democratic pathology. There is no integration of panarchic theory, coupling metrics, wisdom filtering, or script revision mechanisms. In Schank's terms, these platforms are sophisticated venues for expressing opinions following script failure, but not instruments for the revision of the scripts themselves.

4.5 The Collective Intelligence Project (CIP)

The Collective Intelligence Project, co-founded by Divya Siddarth and Saffron Huang, is an R&D nonprofit focused on directing transformative technology toward collective benefit through structured large-scale deliberation (Siddarth & Huang, 2023). Its Global Dialogues initiative combines stratified sampling, AI-enabled facilitated deliberation, and bridging-based aggregation to capture both majority signals and minority concerns across more than seventy countries. In 2024–2025, it engaged over ten thousand participants in structured processes designed to surface shared priorities through collective reflection (CIP Annual Report, 2024). CIP also develops Community Models — AI systems aligned with locally defined normative constitutions — and is building an Epistemic Evaluation Suite to assess truthfulness, impartiality, and good-faith reasoning in AI systems.

Strengths: CIP represents the most sophisticated current attempt to bring structured global deliberation into AI governance. Its combination of demographic rigour, deliberative process design, and AI integration is substantively impressive. Its Grandmothers' Collective

project — encoding wisdom from elders worldwide — directly parallels Agora's focus on long-life-experience as a primary wisdom source and epistemic resource.

Limitations relative to SWARP-Agora: CIP operates primarily at the level of preference and values elicitation; it provides no mechanism for deep script revision in Schank's sense. It lacks a panarchic rhythm detection layer, a continuous adaptive cycle architecture, and a Socratic AI designed to drive participants toward cognitive closure rather than normative consensus. CIP does not filter participants by reflective wisdom potential: all voices are treated as equally valid inputs regardless of depth of experience or reflective capacity. The Grandmothers' Collective is conceptually proximate to Agora but lacks the Markov navigational infrastructure and Free Energy metrics that make SWARP a generative simulator rather than a documentation initiative.

4.6 Integral Life and AQAL (Wilber / Integral Institute)

Ken Wilber's AQAL framework (All Quadrants, All Levels, All Lines, All States, All Types) constitutes the most comprehensive available metatheory for integrating personal development, collective evolution, and systemic change (Wilber, 2006). The Integral Life platform operationalises this framework through curated content, community, and contemplative practice. AQAL's four quadrants — Interior Individual (I), Exterior Individual (It), Interior Collective (We), Exterior Collective (Its) — bear a structural resemblance to PoC's four worldview quadrants, though the mapping is non-trivial: AQAL emphasises developmental levels and maturity lines that PoC does not systematically theorise, while PoC foregrounds dynamic change paths and coupling dynamics that AQAL leaves largely implicit.

Strengths: The AQAL framework provides the most thorough available integration of wisdom traditions, developmental psychology, systems theory, and contemplative practice. Its recognition that interior individual change and exterior collective design must proceed in tandem is deeply convergent with SWARP's PoC architecture.

Limitations relative to SWARP-Agora: Integral Life is primarily a content and developmental community platform, not an adaptive simulator. It does not implement panarchic rhythm detection, Markov navigation, Free Energy metrics, or real-time coupling scores. Human Design is not integrated; individual developmental profiles are assessed through conventional stage instruments and coaching. Most critically, there is no Agora-equivalent: no protected, slow wisdom-ripening space with Socratic AI facilitation specifically designed to produce script closure and inject wisdom products into ongoing adaptive cycles.

4.7 DebateGraph and Argument Mapping Tools

DebateGraph (Baldwin & Price) combines argument visualisation with collaborative wiki editing to make the strongest arguments on all sides of complex public debates freely available and continuously subject to challenge and refinement (DebateGraph.org). It has been deployed by the BBC, the European Commission, the UK Prime Minister's Office, and CNN. Related tools include Debatemap, the MIT Deliberatorium, and Argunet.

Strengths: These tools make the logical structure of complex debates explicit, navigable, and publicly available. They support the analysis of wicked problems by visualising the interconnectedness of positions and concerns. DebateGraph's wiki-based openness enables continuous collaborative improvement of shared argumentative resources.

Limitations relative to SWARP-Agora: Argument mapping tools operate at the level of explicit propositional structure — what positions exist and how they relate logically — rather than at the level of the implicit expectation scripts that generate those positions. They provide no mechanism for script revision, no panarchic rhythm detection, no wisdom-ripening space, and no individual oscillator model. Crucially, argument maps can proliferate indefinitely without producing cognitive closure: the map grows in complexity but the underlying scripts remain unrevised. SWARP-Agora's Markov navigation over a bounded lexicon addresses precisely this pathology by closing the relevant conceptual space rather than expanding it without limit.

4.8 Appreciative Inquiry (Cooperrider)

Appreciative Inquiry (AI), developed by David Cooperrider in the mid-1980s, is a strengths-based approach to organisational change that orients collective inquiry toward identifying and amplifying existing peak experiences and capacities as a basis for envisioning and designing preferred futures (Cooperrider & Whitney, 2005). Its 4-D cycle — Discover, Dream, Design, Deliver/Destiny — guides collectives from strength identification through possibility-envisioning to implementation. AI has been applied at substantial scale across organisations including the United Nations and numerous national contexts.

Strengths: AI is one of the few change methodologies that explicitly foregrounds the experiential and narrative quality of collective wisdom rather than beginning from problem diagnosis. Its emphasis on generative questions, grounded positive imagery, and social-constructivist co-creation is genuinely convergent with the spirit of Agora. Appreciative

Inquiry Summits can engage hundreds or thousands of participants in co-creative processes that build genuine collective ownership.

Limitations relative to SWARP-Agora: AI's defining strength — its sustained positive focus — is simultaneously its primary limitation from a PEFT perspective. By systematically avoiding problem and failure analysis, it cannot complete the script revision cycle that expectation failure requires. Schank's model requires that failures be examined, explained, and generalised before new scripts can be formed; AI bypasses this requirement by moving directly to the positive core, which risks constructing new expectations on a foundation of unrevised scripts. There is no panarchic rhythm detection, no Markov navigation, no individual oscillator model, and no sustained Socratic challenge to ingrained assumptions.

4.9 The Wisdom Council and Dynamic Facilitation (Rough / Center for Wise Democracy)

Jim Rough developed Dynamic Facilitation in the early 1980s as a method for enabling creative resolution of seemingly intractable workplace conflicts. The Wisdom Council Process extends this to large social systems: every three to four months, eight to twelve randomly selected members of an organisation, city, or nation meet for two days with a trained dynamic facilitator, produce a unanimous creative message, and present it to the wider system for further dialogue (Rough, 2002). The process has been embedded in the constitutions of two Austrian federal states — Vorarlberg and Salzburg — and has been applied in corporate and governmental contexts.

Dynamic Facilitation employs a four-poster technique — statements, solutions, concerns, and data — consistently oriented toward what Rough calls *choice-creating*: a relational

mode in which diverse participants address real problems with authentic voice and emotional honesty, seeking breakthrough rather than negotiated compromise.

Strengths: The Wisdom Council Process is the closest existing analogue to Agora among all compared systems. Its commitment to authentic voice, creative breakthrough, and the unanimous message as a genuine voice-of-the-whole resonates strongly with SWARP's Grammars of Engagement and the Socratic spirit of Agora. Its embedding in Austrian constitutional structures demonstrates the institutional viability of wisdom-generating processes at governmental scale. Dynamic Facilitation's choice-creating mode aligns explicitly with the anarchy-spiral dynamic in SWARP's PoC architecture.

Limitations relative to SWARP-Agora: Wisdom Councils are periodic, offline, small-group events, typically involving eight to twelve participants every four months. They lack a continuous adaptive cycle architecture, a real-time rhythm-detection layer, a Markov-navigable knowledge infrastructure, and an individual oscillator model for participant selection by wisdom potential. Facilitation depends on trained human dynamic facilitators, which constrains scalability. Most critically, there is no persistent Remember layer: Wisdom Council outputs are presented to the larger system but are not systematically injected into ongoing adaptive cycles through a mechanism analogous to SWARP's Free Energy and Markov dynamics.

---

5. Comparative Summary

System	Expectation Detection	Script Revision / Closure	Revolt / Remember Detection	Individual Oscillator Model	Wisdom-Ripening Layer	Scalabilit
SenseMaker + Estuarine	Strong	None	Conceptual only	None	None	High
Kumu	Structural	None	None	None	None	Medium
Consilience / Game B	Diagnostic	None	None	None	None	Small grou
Active Inference Simulators	Mathematical	Implicit	None	None	None	Experime
Polis / Loomio / Remesh	Preference-level	None	None	None	None	High
CIP Global Dialogues	Strong	None	None	None	Partial (Grandmothers)	High

System	Expectation Detection	Script Revision / Closure	Revolt / Remember Detection	Individual Oscillator Model	Wisdom-Ripening Layer	Scalabilit
Integral Life / AQAL	Developmental	Implicit	None	Partial (stage)	None	Small–Medium
DebateGraph / Argunet	Propositional	None	None	None	None	Medium
Appreciative Inquiry	Positive	Partial	None	None	Partial	High
Wisdom Council	Authentic	Partial	None	None	Partial	Small grou
SWARP-Agora	**Active Inference**	**Socratic AI**	**Real-time**	**Human Design**	**Agora (full)**	**Fractal / Large**

6. Critical Assessment of SWARP-Agora

6.1 Theoretical Strengths

SWARP-Agora's primary theoretical distinction lies in its integration of four levels that existing systems address only partially and in isolation: the cognitive level (Schank's script theory, expectation failure, and closure mechanisms); the empirical-political level (De Vries's research on the structural conditions of democratic dissatisfaction); the systemic level (panarchy, Revolt/Remember, and fractal adaptive cycles); and the individual oscillator level (Human Design as a precision filter for wisdom potential and coupling style). No existing system in the comparative field combines all four.

The Socratic AI layer provides a mechanism — authentic Socratic elenchus and maieutic questioning — that is absent from all compared systems. The real-time Revolt/Remember detection dashboard, operating across fractal organisational scales via three continuously computed metrics, operationalises panarchic dynamics in a manner no compared system achieves. The closest analogues — CIP's Grandmothers' Collective and the Wisdom Council process — share the wisdom-orientation but lack the continuous adaptive architecture and individual oscillator precision that SWARP-Agora provides.

6.2 Limitations and Boundary Conditions

Four limitations require honest acknowledgement.

Absence of empirical validation. This paper is a theoretical and architectural contribution. No controlled empirical study of SWARP-Agora's effect on cognitive closure rates, script

revision quality, or expectation recalibration has yet been conducted. All comparative claims concerning SWARP-Agora's superiority to existing systems are therefore claims about architectural design and theoretical completeness, not about demonstrated outcomes. This is the most significant current limitation.

Adoption and accessibility barriers. Current adoption remains limited. The platform's simultaneous integration of PoC, panarchy, Active Inference, Human Design, Markov navigation, Grammars of Engagement, and Socratic AI creates a substantial accessibility barrier that risks confining the system to expert users.

Scientific status of Human Design. The Human Design filter introduces an individual selection mechanism whose scientific validity remains contested and for which peer-reviewed empirical evidence is limited. Its use in SWARP is pragmatic and architectural rather than grounded in validated psychometric science.

Author proximity. The analysis draws significantly on the author's own developmental experience, introducing a potential confirmation bias. Comparative assessments of SWARP-Agora's strengths relative to other systems should be read with this proximity in mind.

---

7. Recommendations for Future Development

The following seven recommendations are derived directly from the comparative analysis and from the gaps identified in existing systems.

7.1 Simplify the entry experience. The primary adoption barrier is complexity. A dedicated onboarding path designed for Agora's target participants — individuals with substantial life experience, reflective orientation, and comfort with depth — should reduce the cognitive load of first entry to a single guided sequence: provide birth data → receive Human Design filter profile → write first blog → meet Socrates. All further complexity can be introduced progressively.

7.2 Develop a minimal Wisdom Council integration. The Wisdom Council process is SWARP's closest structural analogue among existing systems and has demonstrated institutionalisable results at governmental scale. SWARP should build a structured Wisdom Council mode: each quarter, a small randomly selected group of Agora participants convenes virtually for a two-day Socratic session, produces a unanimous message, and this message is formally injected into the wider SWARP simulator as a Remember signal. This would bridge offline wisdom-council practice with SWARP's continuous online adaptive architecture.

7.3 Integrate an Appreciative Inquiry discovery phase into Agora onboarding. Appreciative Inquiry's Discover phase — identifying peak experiences and existing strengths — is well suited to Agora's initial participant engagement. The Socrates AI can facilitate an AI-style entry dialogue as the first Agora interaction, grounding the participant's wisdom in their positive life core before the elenctic challenge begins. This approach preserves AI's strength in authentic engagement while ensuring that script failure analysis, which AI methodology tends to bypass, is subsequently engaged rather than circumvented.

7.4 Build a SWARP-CIP bridge for global value elicitation. CIP's Global Dialogues methodology provides a proven infrastructure for reaching demographically representative samples across more than seventy countries in eight or more languages. Agora could benefit from CIP's global reach for identifying wisdom-potential participants, while CIP's dialogues could benefit from SWARP's Agora as a deeper reflective follow-up layer for participants demonstrating high reflective depth. A formal partnership or data bridge between the two platforms would extend the scope of both.

7.5 Operationalise a formal Remember injection protocol. The pathway from Agora's wisdom output to the wider SWARP simulator is architecturally specified but not yet formally protocolised. A Remember injection protocol should define: which Agora outputs (blogs reaching a threshold coupling score, Socratic dialogues achieving demonstrable closure, Wisdom Council messages) trigger a Remember event; how these are tagged in the Markov lexicon; and how they are surfaced to users in α-phase activity. Without this protocol, the Remember function risks remaining theoretical rather than operational.

7.6 Commission an independent empirical validation study. The single most important step for SWARP-Agora's scientific credibility and institutional adoption is a peer-reviewed empirical study comparing cognitive closure rates, script revision quality, and expectation recalibration between SWARP-Agora participants and control groups using comparable platforms (SenseMaker, Polis, or a deliberation platform). A controlled comparison with 50–100 participants would be sufficient to generate publishable findings and provide the evidence base that adoption decisions in public institutions require.

7.7 Develop a SWARP-DebateGraph integration for propositional scaffolding. DebateGraph's explicit argument-structure visualisation could serve as a valuable scaffold

for participants who find unstructured Socratic dialogue cognitively demanding. A lightweight integration — allowing Agora participants to attach a DebateGraph argument map to their blog as optional propositional scaffolding — would bridge the explicit-propositional and experiential-reflective modes of inquiry, expanding accessibility without compromising reflective depth.

---

8. Conclusion

SWARP-Agora represents a theoretically distinct advance beyond existing sensemaking, deliberation, and collective intelligence systems. Where SenseMaker and Kumu detect and visualise system patterns; where the Consilience Project and Game B theorise civilisational redesign; where Active Inference simulators provide mathematical rigour; where CIP elicits global values; where Integral Life supports individual developmental mapping; where DebateGraph makes argument structure explicit; where Appreciative Inquiry builds on collective strength; and where Wisdom Councils generate authentic unanimous creative messages — SWARP-Agora is specifically designed to simulate the conditions under which collective wisdom can ripen and be systematically injected back into adaptive cycles that are otherwise structurally condemned to repeat their failures.

The central theoretical claim is architectural and diagnostic. The political expectation failure that characterises contemporary democracies is not primarily a problem of information access, preference aggregation, or institutional design at the surface level. It is a problem of cognitive closure: the systematic incapacity of collective systems to complete

the learning cycle that script failure, properly processed, makes possible. SWARP-Agora is designed to provide that missing mechanism.

The most promising pathway to adoption combines SWARP's architectural depth with the human-scale accessibility of Wisdom Councils, the positive relational entry of Appreciative Inquiry, and the global demographic reach of CIP's deliberative infrastructure. The seven development recommendations offered in Section 7 provide a concrete roadmap for this integration.

Whether wisdom, once ripened in Agora, can genuinely seed the Remember function in the larger social simulator — and whether this architecture can achieve the adoption it requires to demonstrate its effects empirically — remains the central open question. It is also, in the author's view, the central practical question for the future of democratic self-governance under conditions of complexity.

Declarations

Conflict of interest: The author is the principal developer of the SWARP-Agora platform described in this paper. This proximity is acknowledged as a source of potential bias; the comparative analysis has been conducted with systematic attention to SWARP-Agora's own limitations.

Funding: No external funding was received for this study.

Data availability: No empirical datasets are associated with this theoretical contribution.

References

Cooperrider, D. L., & Whitney, D. (2005). *Appreciative inquiry: A positive revolution in change*. Berrett-Koehler.

De Vries, C. E. (2018). *Euroscepticism and the future of European integration*. Oxford University Press.

De Vries, C. E., & Hobolt, S. B. (2020). *Political entrepreneurs: The rise of challenger parties in Europe*. Princeton University Press.

Friston, K. J., et al. (2022). Active inference and learning. *Neural Computation*, *34*(1), 1–49.

Gerring, J. (2012). *Social science methodology: A unified framework* (2nd ed.). Cambridge University Press.

Gunderson, L. H., & Holling, C. S. (Eds.). (2002). *Panarchy: Understanding transformations in human and natural systems*. Island Press.

Konstapel, J. (2026). *Political expectation failure: Toward a unified theory of democratic instability*. Unpublished manuscript, Leiden.

McWhinney, W. (1992). *Paths of change: Strategic choices for organizations and society*. Sage.

McWhinney, W., & Konstapel, J. (unpublished). *Grammars of engagement*. Unpublished manuscript.

Parr, T., Pezzulo, G., & Friston, K. J. (2022). *Active inference: The free energy principle in mind, brain, and behavior*. MIT Press.

Rough, J. (2002). *Society's breakthrough! Releasing essential wisdom and virtue in all the people*. 1st Books.

Schank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals and understanding*. Lawrence Erlbaum.

Schank, R. C. (1999). *Dynamic memory revisited*. Cambridge University Press.

Schedler, A., & Santiso, J. (1998). Democracy and time: An invitation. *International Political Science Review*, *19*(1), 5–18.

Schmachtenberger, D. (2020–2023). Various writings and lectures. Consilience Project and Game B. Retrieved from https://consilienceproject.org

Siddarth, D., & Huang, S. (2023). *Collective Intelligence Project whitepaper*. CIP.

The Collective Intelligence Project. (2024). *CIP annual report 2024*. CIP.

Small, C. T., et al. (2021). Polis: Scaling deliberation through computational social choice. *arXiv preprint*. arXiv:2102.03970.

Snowden, D. J., & Boone, M. E. (2007). A leader's framework for decision making. *Harvard Business Review*, *85*(11), 68–76.

Snowden, D. J. (2022–2025). *Estuarine mapping framework documentation*. The Cynefin Co.

Wilber, K. (2006). *Integral spirituality: A startling new role for religion in the modern and postmodern world*. Integral Books.