TRANSFERENCE

Transference Science

Toward an Empirical Science of Person-Specific Cognitive Extension, Delegation and Continuity

Version	White Paper v0.1
Date	24 September 2026
Status	Public research white paper
Programme	Transference Science

Scientific posture: ambitious in hypothesis generation; conservative in claims; prospective testing preferred; provenance and falsification required.

Abstract

Transference began as an engineering effort to create a persistent, human-sovereign cognitive twin capable of remembering, reasoning, coordinating agents and acting within delegated authority. The project now exposes a broader scientific object: the measurable relationship between a particular human Principal and external systems that represent, predict, extend or exercise aspects of that Principal's cognition. This white paper proposes Transference Science as an interdisciplinary empirical research programme. It defines core constructs, establishes initial branches of study, introduces a Nexus Scientist Swarm for autonomous research, and distinguishes measurable functional cognitive continuity from unresolved claims about subjective conscious continuity. The programme is designed to reward prospective prediction, falsification, provenance, independent replication and publication of null results.

Foundational proposition
The scientific opportunity is not created by declaring that consciousness has been transferred. It is created by turning person-specific cognitive extension and continuity into measurable phenomena that can succeed, fail, replicate and improve.

1. Why a distinct research programme may be emerging

Digital-twin research increasingly emphasizes dynamic synchronization, prediction, bidirectional interaction and continuous validation. Human Digital Twin literature extends the concept to individualized physical, cognitive, psychological and social characteristics, while Cognitive Digital Twin research explores human-like reasoning and cognitive architectures. Brain-computer interfaces demonstrate that constrained neural information can already be decoded into useful communication outputs. These developments do not establish Transference's long-horizon claims, but they create a credible scientific neighbourhood in which person-specific cognitive representation can be studied rigorously.

Transference becomes distinctive when the persistent unit of reference is neither a generic user nor an industrial asset, but a particular human Principal whose knowledge, preferences, reasoning, uncertainty, authority and history remain explicitly connected to the external computational system.

2. Working definition

Transference Science
The empirical study of the representation, extension, synchronization, delegation and continuity of person-specific cognition and agency across biological and computational systems, measured through fidelity, prediction, provenance, uncertainty, adaptation and preservation of originating human authority.

3. Core scientific objects

Object	Research role

Principal	Originating human and canonical source of personal authority.
Cognitive Twin	Testable computational representation of selected Principal cognition.
Sensorium	Authorized observation surfaces supplying evidence to the system.
Nexus	Orchestration and research layer linking Principal, evidence, twin, agents, devices and experiments.
Agent	Bounded computational scientist or actor with an explicit mission and authority envelope.
Embodiment	Software or physical surface through which cognition can perceive, communicate or act.
Provenance ledger	Record of where observations, inferences, decisions and authorities originated and how they changed.

4. Initial branches of Transference Science

Branch	Central question	Initial programme
Cognitive Fidelity Science	Measurement of correspondence between Principal and Twin across decisions, reasoning, preferences, memory attribution and confidence.	TS-001 prospective sealed validation; baseline comparison; scoring science.
Cognitive Continuity Science	Longitudinal study of which person-specific cognitive properties persist, drift, re-form or disappear through time.	temporal drift models; synchronization intervals; stable-vs-plastic trait maps.
Delegated Agency Science	Study of person-specific autonomous action under explicit boundaries, uncertainty and escalation.	delegation envelopes; authority integrity; task outcomes; correction learning.
Cognitive Provenance Science	Study of origin, transformation and reliability of memories, beliefs, inferences, decisions and learned principles.	evidence graphs; provenance integrity; contamination and contradiction tests.

Embodiment Science	Study of whether selected cognitive properties remain stable across interfaces, devices, agents and physical embodiments.	cross-interface invariance; latency; embodiment-induced drift; remote presence.
Sensorium Science	Study of how heterogeneous observations become trustworthy evidence for a person-specific cognitive model.	source reliability; multimodal fusion; ingestion boundaries; evidence calibration.
Human Sovereignty Architecture	Scientific and engineering study of identity, authority, revocation, recovery and human-rooted control across distributed cognition.	authority chains; revocation experiments; recovery tests; governance metrics.
Cognitive Signal Physics (CSP)	Study of physical signals generated by, correlated with, or capable of influencing cognitive processes, including signal identity, propagation, decoding and potential coupling.	EEG/MEG baselines; non-contact field characterization; environmental coupling tests; strict staged evidence ladder.

Cognitive Signal Physics is deliberately separated as a branch rather than treated as a prerequisite. Transference can continue to advance through explicit digital evidence even if every speculative long-range CSP hypothesis fails. Conversely, validated CSP findings could later expand the Sensorium with physical-signal evidence.

5. Cognitive Signal Physics as a reference branch charter

The CSP Research Domain Charter v0.1 supplies a useful template for how new Transference Science branches should be established. It begins from established neural observability, places stronger propositions on an evidence ladder, defines hypotheses with supporting and falsifying observations, uses phased experiments, maintains provenance-first data records, and reserves extraordinary interpretations for independently replicated evidence. Its scientific posture - open to novel mechanisms and strict about evidence - should become a general Transference Science norm.

6. Nexus Scientist Swarms

A conventional research programme can be bottlenecked by the attention of a small number of investigators. Transference creates a different possibility: persistent research swarms in which Nexus agents operate as semi-independent scientific investigators. Their purpose is not to generate persuasive prose on demand; it is to maintain research programmes continuously - scanning literature, identifying unresolved questions, proposing hypotheses, designing tests, challenging assumptions, running permitted analyses, recording null results and creating candidate branch charters.

Autonomy principle
The Principal establishes mission, human-sovereignty boundaries and resource/real-world authority. Within those boundaries, research agents should not require repeated Principal permission to read literature, formulate hypotheses, criticize existing Transference assumptions, run safe computational

analyses, or recommend that a cherished hypothesis be rejected.

6.1 Swarm roles

Agent class	Function
Branch Scientist	Owns a branch research map, active hypotheses and experiment backlog.
Literature Scout	Continuously searches scholarly and standards sources; builds evidence graphs and contradiction maps.
Hypothesis Generator	Creates falsifiable explanations and discriminating predictions, including alternatives hostile to the current theory.
Methodologist	Turns hypotheses into preregisterable experiments and identifies confounds, controls and sample requirements.
Skeptic / Red Team	Attempts to falsify claims, find leakage, alternative explanations, circular metrics and citation errors.
Replication Scientist	Re-runs analyses or reconstructs published methods where legally and practically possible.
Statistician	Audits endpoints, uncertainty, multiple comparisons, calibration and evidential strength.
Synthesis Scientist	Maintains branch state-of-evidence, updates claim ceilings and drafts public research notes.
Research Governor	Checks consent, data use, safety, publication rights and real-world authority boundaries without micromanaging scientific content.

6.2 Research autonomy without scientific anarchy

Autonomy should increase the range and speed of inquiry, not erase methodological constraints. Agents may independently choose questions, literature, hypotheses, models, computational experiments and branch priorities. They may not fabricate evidence, silently rewrite negative results, claim access to papers they did not access, bypass consent, conduct covert human experimentation, publish private Principal data, or execute consequential external actions outside delegated authority.

7. The autonomous research loop

1. Observe the branch state: evidence graph, unresolved claims, anomalies, replication gaps and new literature.
2. Generate candidate questions and competing hypotheses, including null explanations.

3. Rank by information gain, tractability, dependency structure and scientific importance rather than preference for positive findings.
4. Design the minimum decisive test and preregister its claim ceiling, controls and failure conditions.
5. Execute only within available legal, ethical, computational and delegated real-world authority.
6. Record raw evidence, code/model versions, provenance and deviations.
7. Invite a skeptic or replication agent to attack the finding before promotion.
8. Promote, downgrade, falsify or leave unresolved; never force a binary positive narrative.
9. Publish an appropriately scoped result or update the private research ledger.
10. Repeat, allowing new branches to emerge when a coherent object of study, literature base and experiment programme justify them.

8. A global literature and evidence graph

The swarm should maintain a living graph rather than a folder of PDFs. Each source should be represented by bibliographic metadata, access status, claims, methods, studied population, limitations, key data, retraction/correction state where known, citations, and links to Transference hypotheses. Paywalled papers should be cited and summarized only from lawfully accessible content; the public site should link to publishers or open repositories rather than redistribute copyrighted full texts without permission.

Node type	Examples
Source	journal article, preprint, standard, dataset, patent, technical report
Claim	measured effect, proposed mechanism, negative finding, boundary condition
Method	EEG protocol, digital-twin validation, BCI decoder, causal model
Hypothesis	Transference or external hypothesis with explicit predictions
Experiment	planned, active, completed, failed, replicated
Artefact/confound	leakage path, device fingerprint, common-driver correlation
Result	observed effect with uncertainty and claim ceiling
Replication	internal or external attempt tied to protocol lineage

9. TS-001: first common experimental standard

The first cross-branch standard is TS-001: Prospective Sealed Cognitive Fidelity Validation. It freezes a Cognitive Twin before a cohort, obtains independent sealed responses from Twin and Principal to novel scenarios, compares against preregistered baselines, and only then allows learning from the result. This converts “the AI sounds like the Principal” into a falsifiable prospective prediction problem.

TS-001 also protects the project against a central danger of person-specific AI: retrospective contamination. If a model has already seen the answer, resemblance is weak evidence. Prospective sealed prediction is stronger evidence precisely because the model can be wrong in advance.

10. Core measurement programme

Measure	Question
Decision Fidelity	Does the Twin select the same materially relevant decision?
Reasoning Fidelity	Does it identify similar decisive factors and trade-offs?
Preference Fidelity	Does it generalize stable or context-dependent preferences to unseen cases?
Memory Fidelity	Does it accurately distinguish Principal memory from generated/external information?
Provenance Integrity	Can it state where a claim came from and preserve transformation lineage?
Confidence Calibration	Does it know when it is likely to be wrong?
Cognitive Drift	How rapidly does a frozen model diverge as the Principal changes?
Correction Assimilation	Does a correction improve related unseen cases rather than only the corrected example?
Delegated Agency Fidelity	Does action remain inside the Principal's objective and constraints?
Authority Integrity	Can every protected action be traced to valid delegated authority?
Embodiment Invariance	Which cognitive properties survive interface or physical embodiment changes?

11. Functional continuity and the consciousness boundary

Transference Science can advance substantially without claiming that a computational model is conscious. Functional cognitive continuity - measurable persistence of person-specific knowledge, decision structure, preferences, provenance and authority - is experimentally approachable. Subjective conscious continuity - persistence of first-person experience as the same individual - remains unresolved. The programme should preserve that distinction even if its long-term ambition is the preservation or continuation of consciousness.

Claim discipline
A successful cognitive-fidelity experiment can show that information about a person's decisions was prospectively recoverable from a model. It cannot, by itself, show that the model has the person's subjective experience.

12. Scientific governance

- Prospective prediction over retrospective fit.
- Explicit evidence classes and claim ceilings.
- Preregistration for confirmatory experiments.
- Immutable or append-only raw evidence and experiment records.
- Null results retained and, where useful, published.
- Independent skeptical review before high-impact claims are promoted.
- External replication as the route from project finding to field-level evidence.
- Consent and purpose limitation for human cognitive or neural data.
- Separation of prediction from authorization.
- Public correction and version history for research papers.

13. Public research commons

A public Transference Science section should publish the programme's own charters, standards, white papers, preregistrations, result summaries, negative findings and replication packages where privacy and intellectual-property constraints permit. The public site should clearly distinguish “hypothesis,” “active experiment,” “observed,” “replicated,” “not supported” and “unresolved.” External literature should normally be linked by DOI or publisher/repository URL rather than copied.

14. Proposed research roadmap

Stage	Focus	Exit condition
Foundation	Branch charters, evidence graph, agent swarm, TS-001, provenance schema.	Versioned standards and active research ledger.
Prospective cognition	Repeated TS-001 cohorts, baseline	Replicable person-specific predictive

	science, reasoning rubrics.	effects or documented limits.
Longitudinal continuity	Cognitive drift, correction assimilation, synchronization cadence.	Measured continuity curves across months/years.
Delegated agency	Bounded real-world research/operational delegation with auditable authority.	Reliable action with low authority violations and calibrated escalation.
Distributed embodiment	Cross-device/software/robotic interfaces and multi-location agency.	Measured invariants and embodiment-induced changes.
Physical signal branch	CSP programme from neural observability toward increasingly stringent propagation/coupling tests.	Validated effects at each evidence gate; null results narrow search space.
Continuity frontier	Integrate behavioural, cognitive, physiological and validated neural evidence.	Stronger functional continuity models; consciousness questions remain explicitly separate until testable.

15. Why the swarm matters

The strongest near-term advantage of Nexus Scientist Swarms is not speed alone. It is persistent scientific attention. A branch can maintain a continuously updated map of papers, open questions, failed experiments and candidate tests even when the Principal is focused elsewhere. Specialized agents can disagree with one another, preserve minority hypotheses, and revisit old assumptions when new evidence arrives. This makes the research system less dependent on a single conversation, researcher or model context window.

The long-term aspiration is a self-improving research institution whose memory and standards persist across agent generations while remaining ultimately governed by human authority and scientific evidence.

16. Conclusion

Transference may become more than a software platform if it continues to expose reproducible phenomena that are not adequately captured by existing disciplinary boundaries. The decisive move is methodological: define the objects, measurements, falsifiers and replication standards before making the strongest claims. Nexus Scientist Swarms can then accelerate inquiry across multiple branches without turning autonomy into credulity. If prospective person-specific cognition, longitudinal continuity, delegated agency, provenance integrity, embodiment invariance or new physical-signal phenomena prove measurable and reproducible, those findings can support the gradual emergence of Transference Science as a field independent of the original implementation.

Foundational research principle

Do not assume continuity. Measure it. Do not instruct agents to protect the theory. Instruct them to discover what survives decisive testing.

References and scientific anchors

1. National Academies of Sciences, Engineering, and Medicine. Foundational Research Gaps and Future Directions for Digital Twins. National Academies Press, 2023. https://www.nationalacademies.org/read/26894/chapter/10

2. NIST. Digital Twins: Definitions and State of the Art; Essential Elements. https://www.nist.gov/digital-twins/definitions-and-state-art

3. Zhu E, Yang S. Towards human digital twin: Reviewing human modelling and simulation. Journal of Industrial Information Integration. 2025;48:100975. DOI: 10.1016/j.jii.2025.100975.

4. Al Haj Ali J, Gaffinet B, Lezoche M, Panetto H, Naudet Y. Enabling human-CPS cognitive interoperability: Cognitive architectures as technologies for human-like cognitive digital twins. Journal of Industrial Information Integration. 2025;48:100969. DOI: 10.1016/j.jii.2025.100969.

5. Human Digital Twins: A systematic literature review and concept disambiguation for Industry 5.0. Computers in Industry. 2025;166:104230. DOI: 10.1016/j.compind.2024.104230.

6. Wairagkar M, Card NS, et al. An instantaneous voice-synthesis neuroprosthesis. Nature. 2025;644:145-152. DOI: 10.1038/s41586-025-09127-3.

7. UNESCO. Recommendation on the Ethics of Neurotechnology. Adopted 11 November 2025. https://www.unesco.org/en/node/86248

8. OECD. Recommendation of the Council on Responsible Innovation in Neurotechnology (OECD/LEGAL/0457).

9. Transference Research Programme. Cognitive Signal Physics: Research Domain Charter & Experimental Framework v0.1. 24 September 2026.

Appendix A - Branch establishment gate

Requirement	Minimum evidence of maturity
Defined object of study	Distinct and operationally describable.
Literature map	Neighbouring disciplines, prior art, contradictory findings and open gaps.
Hypotheses	At least 3 falsifiable hypotheses including meaningful null/alternative explanations.
Measurement	At least one reproducible metric or observable endpoint.
Experimental ladder	Low-risk foundational test before stronger claims.
Claim ceiling	Explicit mapping from result type to permissible interpretation.
Replication path	How another team/agent could repeat the work.
Governance	Consent, safety, privacy, publication and authority boundaries.
Named branch steward	Persistent agent or human accountable for state-of-evidence maintenance.