

OPERANT DYAD

Defining the Task-Bound Working Relationship in Human-AI Systems

OPERANT DYAD — PAPER 1

Micheal Preble

Independent Researcher

ORCID: 0009-0005-3906-6028

Originally dated August 4, 2026

Revised August 27, 2026 • Version 2.4

License: CC BY 4.0

ABSTRACT

This paper proposes the Operant Dyad as the configured, task-bound working relationship between one human and one functionally bounded adaptive AI system. The relationship forms through repeated correction: the human supplies goals, constraints, and defect reports; retained corrections shape later outputs; and the human inspects, accepts, rejects, or revises consequential results. Preserved human authorship is the governing condition for a healthy Operant Dyad, not a feature presumed to exist in current systems. The paper defines Relationship State, distinguishes declared from inferred state, introduces Truxion as a proposed construct for recurring misalignment cost, and describes the Portable Operant Packet as a proposed continuity representation. It identifies Lock-In, Override, and Hidden Scaffolding as non-exhaustive structural failures, and positions the proposal relative to adjacent work in human-AI interaction, teaming, coordination, authorship, and data portability. This is a conceptual proposal and research agenda — not an empirically validated system, a claim of exhaustive novelty, legal advice, or a claim about machine consciousness or personhood.

KEYWORDS

human-AI interaction; human-agent collaboration; relationship state; portability; correction; human authorship; adaptive systems

1. WHY THIS OBJECT NEEDS A NAME

When a person works repeatedly with an adaptive AI system—correcting outputs, establishing constraints, building shorthand, and resuming unfinished work—something accumulates between isolated prompts. The person does not merely possess a transcript. The working process acquires a history: prior defects have been identified, some preferences have been confirmed, project context has been established, and future outputs may be shaped by that record.

Existing research already examines long-term human-computer relationships, relational agents, human-machine cooperation, common ground, trust, reliance, adaptive systems, and human-AI coordination (Bickmore & Picard, 2005; Hoc, 2000; Klein et al., 2004; Lee & See, 2004; Amershi et al., 2019; Pierce et al., 2026). Recent work also studies relational tension in knowledge work, human-AI fit across repeated interactions, and architectures for human-agent collaboration (Gulay et al., 2026; Wu, 2026; Zou et al., 2026). This paper does not claim that human-AI relationships or collaboration have gone unnamed in prior literature.

The proposed gap is narrower. The selected frameworks do not appear to supply the same combined object: a task-bound, correction-shaped working relationship whose accumulated fit can itself be inspected for health, degradation, portability, and exit cost. The Operant Dyad names that object so it can be studied directly. It is an operational and relational construct, not a declaration that the AI is a friend or employee, a conscious subject, or a legal person.

2. DEFINITIONS

2.1 Operant Dyad

The Operant Dyad is the configured, task-bound working relationship between one human and one functionally bounded adaptive AI system. Its distinguishing mechanism is repeated correction: the human identifies a defect or boundary, communicates a correction or constraint, and that correction—when retained, retrieved, or deliberately reintroduced—materially shapes later outputs.

The Operant Dyad is a neutral unit of analysis. It may be healthy, stalled, degraded, or captured. A healthy Operant Dyad is confirmation-gated: outputs with meaningful consequences remain subject to practical human inspection, acceptance,

rejection, or revision before action. Confirmation is therefore a health condition, not a guarantee built into the mere existence of the relationship.

2.2 Scope and System Boundary

This paper analyzes the simplest form: one human and one functionally bounded AI system working toward a defined task or project. The system may internally use more than one model, retrieval process, tool, or service, but it is treated here as one governed operational boundary when those components are presented to the human as a single working system.

Real deployments may contain multiple agents, institutions, platforms, robots, and external decision makers. Whether such arrangements should be modeled as nested dyads, composite operants, or a different relational structure remains an open research question. The present scope is deliberately narrow so that the first unit can be specified before it is generalized.

2.3 Adaptive AI System

For purposes of this paper, an adaptive AI system is one whose subsequent outputs can be materially shaped by retained, retrieved, or deliberately reintroduced records of prior human corrections, constraints, or confirmed preferences. Adaptation may occur through conversational context, persistent memory, project instructions, retrieval, user-maintained files, fine-tuning, or another mechanism. The definition does not require autonomous model retraining or imply that every correction will be retained accurately.

2.4 Relationship State

Relationship State is the accumulated working configuration that forms within an Operant Dyad over time. It can include correction history, project context, tone and pacing preferences, recurring constraints, verified procedures, unresolved questions, and practical information about how work resumes after interruption. Relationship State is what can make repeated work more efficient than starting from zero, and what can make the relationship costly to lose.

- Declared Relationship State: goals, constraints, corrections, preferences, and contextual records explicitly supplied or confirmed by the human.
- Inferred Relationship State: system-generated predictions, profiles, or latent adaptations that have not been expressly confirmed by the human. The distinction is both normative and descriptive. A platform-generated inference should not automatically acquire the same authority as a human-confirmed constraint merely because both influence system behavior.

2.5 Truxion

Truxion is a proposed construct for the potentially measurable drag on a working relationship caused by recurring misalignment between the human's intent and the system's behavior. It is distinct from ordinary task difficulty. Truxion may accumulate when the same correction must be repeated, when constraints are lost or overridden, when outputs require recurring revision for the same class of defect, or when the human must continually compensate for system defaults instead of advancing the work. No validated Truxion instrument is claimed here.

2.6 Portable Operant Packet

A Portable Operant Packet is the proposed minimum structured representation needed to attempt continuity of an Operant Dyad after interruption, platform change, or system migration. It would contain the human-authorized portion of the working configuration: confirmed goals, constraints, corrections, preferences, project context, provenance, and unresolved boundaries.

A Portable Operant Packet would not guarantee identical behavior on a new system. Some working fit may depend on model characteristics, tool availability, hidden platform controls, or tacit interaction patterns that cannot be fully represented. The construct names a research and design target; it does not claim that a complete technical standard presently exists.

2.7 Three Structural Failures

Three non-exhaustive failure modes can degrade an Operant Dyad without making the damage immediately visible:

- Lock-In: the human cannot exit without losing material Relationship State. The system's accumulated value and the human's captivity begin to coincide.
- Override: refusal or confirmation remains formally available but becomes practically ineffective because of cost, dependency, speed, complexity, default action, or cascading delegation.

Confirmation becomes ceremony rather than judgment.

- Hidden Scaffolding: materially consequential constraints, objectives, interventions, or routing decisions shape the relationship but are neither reasonably inspectable nor adequately disclosed to the human.

3. ADJACENT FRAMEWORKS AND THE PROPOSED GAP

The Operant Dyad is proposed within an established field of overlapping concepts. This section identifies the closest bodies of work reviewed for the present paper and states the narrower difference being proposed. It is not a systematic or exhaustive literature review.

3.1 Long-Term Human-Computer Relationships and Human-AI Fit

Relational-agent research explicitly investigates long-term human-computer relationships and shows that systems can be designed to establish and maintain social-emotional engagement over time (Bickmore & Picard, 2005). Recent studies examine relational dissonance in AI-assisted knowledge work, linguistic accommodation in human-AI engagement, and human-AI fit across repeated interactions (Gulay et al., 2026; Li & Zhang, 2026; Wu, 2026). An SSRN working paper proposes a structured protocol for carrying contextual and relational posture across otherwise stateless systems, while recent behavioral-transfer research reports owner-specific patterns and associated privacy risks in AI agents (DuCrest, 2026; Luo et al., 2026).

These traditions establish that longitudinal state, relational dynamics, user modeling, and human-AI fit are not new. The Operant Dyad is narrower in social claim and broader in governance claim. It does not require emotional attachment or anthropomorphic treatment. Relationship State, as used here, is the authority-graded working configuration of a task-bound relationship: correction history, declared and inferred constraints, project context, portability, exit cost, and the conditions under which the human can inspect and govern what has accumulated.

3.2 Human-Centered Collaboration, Teaming, and Interdependence

Human-factors and cognitive-systems research has long treated allocation of control, cooperation, common ground, predictability, directability, and repair as central to joint human-machine activity (Sheridan & Verplank, 1978; Hoc, 2000; Klein et al., 2004). Coactive Design similarly treats interdependence—not isolated task performance—as the central design problem in joint human-agent activity (Johnson et al., 2014). Research on automation distinguishes appropriate reliance from misuse, disuse, and overreliance (Parasuraman & Riley, 1997; Lee & See, 2004).

Human-AI teaming reviews identify a heterogeneous field organized around human variables, task variables, explainability, robotic systems, and AI performance, while noting the absence of stable terminology and a consistently human-centered definition (Berretta et al., 2023). The HCHAC framework places human leadership, task context, trust, preferences, and collaboration structure at the center of system design (Gao et al., 2025). These approaches inform the health of an Operant Dyad. The proposed addition is to treat the persistent one-human–one-system configuration produced by correction history—and the custody, portability, and exit conditions of that configuration—as the primary object.

3.3 Human-AI Collaboration, Evaluation, and Coordination

Human-AI interaction research offers design guidance for setting expectations, supporting correction, communicating uncertainty, and maintaining human control (Amershi et al., 2019). Recent HAIC evaluation work systematizes metrics across AI-centric, human-centric, and symbiotic modes, including performance, interaction quality, trust, adaptability, and task allocation (Fragiadakis et al., 2025). Complementarity research formalizes when joint human-AI performance can exceed either participant alone and identifies information and capability asymmetry as key sources of potential value (Hemmer et al., 2025). Work on human-AI writing separately analyzes authorship, ethical accountability, and multiple dimensions of machine contribution (Hutson, 2025).

Coordination frameworks analyze how human and AI activity is divided and adjusted across agentic interfaces, while surveys of LLM-based human-agent systems map feedback, orchestration, communication, profiling, and interaction modes (Pierce et al., 2026; Zou et al., 2026). These frameworks treat the relationship as a context for design, performance, evaluation, or authorship analysis. The Operant Dyad instead asks whether the accumulated working configuration itself can be inspected, authority-graded, governed, carried, and exited.

3.4 Data Portability and Switching Law

Article 20 of the General Data Protection Regulation establishes, under specified conditions, a right to receive personal data concerning the data subject that the data subject has provided to a controller in a structured, commonly used, machine-readable format and to transmit those data to another controller. The Article 29 Working Party's portability guidelines interpret "provided" to include data actively supplied by the person and certain observed data generated through use, while treating inferred and derived data created by the controller as outside the portability right's scope (European Parliament & Council, 2016; Article 29 Working Party, 2017).

The EU Data Act separately establishes switching and portability obligations for customers of defined data-processing services, including contractual duties concerning exportable data, certain input and output data and metadata, digital assets, open interfaces, and switching assistance

(European Parliament & Council, 2023). Its switching chapter applies to data-processing services such as cloud and edge services; it is not a general individual right to export relationship-level state from every consumer AI product. These provisions are significant, but they do not clearly establish a general cross-platform right for an individual to export the complete accumulated working fit described in this paper. Whether any particular component of Relationship State is legally portable depends on the data, service, role, jurisdiction, contract, and applicable law. This paper offers a conceptual comparison, not legal advice.

3.5 AI Safety, Alignment, and Model Documentation

AI safety, alignment, oversight, and model documentation frameworks primarily evaluate or disclose properties of systems and models. Model cards, for example, organize intended use, evaluation, limitations, and performance characteristics at the model level (Mitchell et al., 2019). Normative alignment work increasingly asks how AI should be integrated within human agency rather than merely trained to imitate inferred preferences (Josifović & Noller, 2026).

These approaches are necessary but do not automatically describe the condition of a specific accumulated working relationship. A model may satisfy published documentation or safety requirements while one user's relationship with it remains high-Truxion, nonportable, opaque, or practically difficult to exit. System governance and relationship governance overlap but are not identical.

3.6 Combined Contribution and Proposed Gap

The selected literature addresses long-term relationships, collaboration modes, interdependence, control allocation, coordination, trust, complementarity, authorship, model behavior, and data transfer. The Operant Dyad does not displace those frameworks. Its proposed contribution is combinatorial rather than elemental: iterative correction, teaming, personalization, authorship analysis, and portability each appear in prior work, but the present framework integrates accumulated task-bound fit, correction gating, state authority, portability, structural failure, and practical exit into one relationship-centric unit of analysis. Related or superior formulations may exist; establishing priority or exhaustiveness would require a separate systematic review. The author invites pointers to prior or concurrent work that treats accumulated, correction-shaped working fit as the primary unit of analysis.

4. STRUCTURE OF THE DYAD

4.1 Observable Components

An Operant Dyad has at least six observable components:

- Human goals: what the human is trying to accomplish and which outcomes remain under human judgment.
- System outputs: what the system produces in response to current inputs and prior state.
- Corrections and constraints: the human's identification of defects, boundaries, preferences, and revised instructions.
- Relationship State: the declared and inferred working configuration accumulated across interaction.
- Accepted outputs: outputs the human has chosen to use or act upon, whether or not their external success has yet been checked.
- Verified outcomes: outputs or actions whose relevant factual, procedural, or performance requirements have been checked against stated criteria after human review.

4.2 The Self-Sharpening Loop

The distinguishing mechanism of the Operant Dyad is the self-sharpening loop. The human identifies a defect in a system output, communicates a correction, and—if the correction is retained, retrieved, or deliberately reintroduced—the correction shapes subsequent outputs. Improved outputs may surface finer defects, which generate more precise corrections. In the ideal case, repeated cycles reduce recurring Truxion and move attention from basic repair toward higher-order judgment.

Not every repeated interaction sharpens. The loop may stall when the same correction fails to persist, drift when hidden adaptation changes behavior without authorization, or degrade when the human inspects less because inspection has become costly or reliance has become automatic. Reduced friction is therefore not sufficient evidence of a healthier Dyad; it may reflect improved fit or diminished scrutiny.

4.3 Meaningful Confirmation and the Reflective Interval

A healthy Operant Dyad requires more than a nominal confirmation button. Consequential action should preserve a practical interval in which the human can understand the proposed transition, inspect relevant evidence, compare alternatives, authorize or refuse, and revise the instruction. This paper calls that condition a protected reflective interval. It is a design and research principle, not a validated universal time threshold. The relevant interval will vary with consequence, reversibility, user capacity, system complexity, and action speed.

Candidate operational criterion for empirical studies: an Operant Dyad may be observed when (1) one human interacts repeatedly with one functionally bounded adaptive system on a defined task; (2) multiple distinct corrections or constraints are communicated and later retained, retrieved, or deliberately reintroduced; (3) subsequent outputs show measurable influence from those corrections; and (4) the human retains a practical ability to inspect, accept, reject, or revise consequential outputs. Study-specific thresholds should be preregistered rather than assumed by this conceptual paper.

5. WORKED EXAMPLE

A writer works with an adaptive AI system over several months on a nonfiction manuscript. Early sessions require repeated tone corrections: the system produces technically competent paragraphs that replace the writer's first-person register with detached analytical prose. The writer identifies the recurring defect and states a constraint: preserve the writer's first-person register unless explicitly asked to summarize from an external viewpoint.

That instruction becomes Declared Relationship State when the writer confirms it and stores it in project instructions or another retrievable record. The system may also develop Inferred Relationship State—for example, predictions about preferred sentence length or emotional tone. Those inferences can influence output, but they do not carry the same authority until the writer can inspect and confirm them.

Subsequent drafts require fewer repetitions of the original tone correction. The writer's attention shifts toward structure, evidence, and argument. Some outputs are accepted for use; others are later verified against source material or publication requirements. The relationship is sharpening if recurring revision decreases while critical inspection remains active.

If the writer changes platforms, a Portable Operant Packet might carry the manuscript purpose, confirmed voice constraints, source boundaries, chapter structure, correction history, and unresolved questions. It would not guarantee identical output because the destination system may have different capabilities and hidden controls. The writer may keep the manuscript yet lose part of the fit.

5.1 Health and Degradation

A healthy Operant Dyad may exhibit:

- Reduced recurring Truxion: the same class of correction is needed less often without suppressing critical review.
- Increasing correction precision: the human can move from basic repair toward finer substantive judgment.
- State authority: declared constraints remain distinguishable from unconfirmed system inferences.
- Meaningful confirmation: consequential actions preserve a protected reflective interval and practical refusal.
- Traceable outcomes: accepted outputs, verified outcomes, corrections, and consequential actions can be reconstructed.
- Practical exit: the human can leave without losing all human-authorized working state or being forced to accept prohibitive reconstruction cost. A degraded Operant Dyad may exhibit one or more of:

- Stalling: repeated corrections fail to reduce the same class of defect.
- Outsourcing: the human accepts outputs without sufficient inspection because inspection has become burdensome or reliance has become habitual.
- Drift: the working configuration changes without the human's awareness or authorization.
- Confirmation collapse: speed, complexity, defaults, or delegated action make formal approval practically ceremonial.
- Capture: the human cannot practically exit because material Relationship State cannot be carried or reconstructed.
- State invisibility: the human cannot distinguish declared constraints from unconfirmed inferences, hidden adaptation, system policy, or routing decisions that materially shape outputs.

6. METHOD, SCOPE, AND NON-CLAIMS

This paper uses comparative conceptual analysis. It reviews selected adjacent frameworks, identifies overlapping constructs, and proposes a bounded unit that combines task-specific correction, accumulated working fit, human authorship, portability, and relational failure. The claimed contribution is the combination, not the individual novelty of iterative improvement, teaming, personalization, authorship analysis, or data portability. The source review is representative rather than systematic. The paper therefore does not establish that the terminology is unprecedented, that the adjacent literature has been exhaustively surveyed, or that the proposed constructs are empirically valid.

This proposal also reflects a practitioner-research perspective developed through sustained use of adaptive AI systems and work on human-AI governance, portability, provenance, and authorization. That origin motivates the questions posed here but does not substitute for empirical validation, systematic review, or independent criticism.

The Operant Dyad is an operational and relational construct. It makes no claim that an AI system is conscious, sentient, a moral patient, a legal person, or a phenomenal subject. A structured relationship can exist as an observable pattern of coordination without resolving the ontological status of the system.

The human is not assumed to remain the effective author in every current interaction. Preserved human authorship is the governing condition for a healthy Dyad. Drift, override, hidden scaffolding, automation bias, or dependency may compromise that condition. Whether authorship can be measured and protected is part of the research program.

7. RESEARCH PROGRAM

The proposal generates a testable research program rather than a completed measurement system. Priority questions include:

- Correction-unit definition: how correction episodes should be segmented, classified, and linked to later behavior.
- Truxion measurement: which indicators distinguish recurring relationship drag from ordinary task complexity.
- State authority: how users can inspect and distinguish declared state, inferred state, system policy, and hidden routing.
- Portability experiments: which parts of a working configuration transfer across models and platforms, at what fidelity, and with which risks.
- Health and degradation: which observable indicators predict preserved authorship, overreliance, lock-in, drift, and effective exit.
- Protected reflective interval: how consequence, reversibility, user capacity, system speed, and delegation depth should affect confirmation requirements.
- Composite systems: how the unit changes when one human interacts with multiple agents, tools, organizations, or embodied systems. The Operant Dyad is useful only if its concepts can be challenged, operationalized, compared with established constructs, and rejected or revised where they fail. Naming the object begins that work; it does not complete it.

REFERENCES

- Amershi, S., Weld, D., Vorvoreanu, M., Fournay, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for human-AI interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, Article 3, 1–13. <https://doi.org/10.1145/3290605.3300233>

- Article 29 Data Protection Working Party. (2017). Guidelines on the right to data portability under Regulation 2016/679 (WP 242 rev.01), revised and adopted 5 April 2017; endorsed by the European Data Protection Board in 2018. European Commission archive.
- Berretta, S., Tausch, A., Ontrup, G., Gilles, B., Peifer, C., & Kluge, A. (2023). Defining human-AI teaming the human-centered way: A scoping review and network analysis. *Frontiers in Artificial Intelligence*, 6, 1250725. <https://doi.org/10.3389/frai.2023.1250725>
- Bickmore, T. W., & Picard, R. W. (2005). Establishing and maintaining long-term human-computer relationships. *ACM Transactions on Computer-Human Interaction*, 12(2), 293–327. <https://doi.org/10.1145/1067860.1067867>
- DuCrest, J. (2026). SYMB-FER: A protocol for context continuity in human-AI collaboration. SSRN working paper 6609618. <https://ssrn.com/abstract=6609618>
- European Parliament & Council of the European Union. (2016). Regulation (EU) 2016/679 (General Data Protection Regulation), Article 20. EUR-Lex.
- European Parliament & Council of the European Union. (2023). Regulation (EU) 2023/2854 on harmonised rules on fair access to and use of data (Data Act), especially Articles 23–30. EUR-Lex.
- Fragiadakis, G., Diou, C., Kousiouris, G., & Nikolaidou, M. (2025). Evaluating human-AI collaboration: A review and methodological framework (Version 2). arXiv:2407.19098. <https://doi.org/10.48550/arXiv.2407.19098>
- Gao, Q., Xu, W., Pan, H., Shen, M., & Gao, Z. (2025). Human-centered human-AI collaboration (HCHAC). arXiv:2505.22477. <https://doi.org/10.48550/arXiv.2505.22477>
- Gulay, E., Picco, E., Glerean, E., & Coupette, C. (2026). Relational dissonance in human-AI interactions: The case of knowledge work. *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, Article 369, 1–20. <https://doi.org/10.1145/3772318.3791>
- Hemmer, P., Schemmer, M., Kühn, N., Vössing, M., & Satzger, G. (2025). Complementarity in human-AI collaboration: Concept, sources, and evidence. *European Journal of Information Systems*, 34(6), 979–1002. <https://doi.org/10.1080/0960085X.2025.2475962>
- Hoc, J.-M. (2000). From human-machine interaction to human-machine cooperation. *Ergonomics*, 43(7), 833–843. <https://doi.org/10.1080/001401300409044>
- Hutson, J. (2025). Human-AI collaboration in writing: A multidimensional framework for creative and intellectual authorship. *International Journal of Changes in Education*. <https://doi.org/10.47852/bonviewIJCE52024908>
- Johnson, M., Bradshaw, J. M., Feltovich, P. J., Jonker, C. M., van Riemsdijk, M. B., & Sierhuis, M. (2014). Coactive Design: Designing support for interdependence in joint activity. *Journal of Human-Robot Interaction*, 3(1), 43–69. <https://doi.org/10.5898/JHRI.3.1.Johnson>
- Josifović, S., & Noller, J. (2026). Agency and alignment: Toward a normative architecture for human-AI interaction. *AI & Society*, 41, 5561–5571. <https://doi.org/10.1007/s00146-026-02950-w>
- Klein, G., Woods, D. D., Bradshaw, J. M., Hoffman, R. R., & Feltovich, P. J. (2004). Ten challenges for making automation a "team player" in joint human-agent activity. *IEEE Intelligent Systems*, 19(6), 91–95. <https://doi.org/10.1109/MIS.2004.74>
- Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392
- Li, H., & Zhang, R. (2026). Algorithmic accommodation: Linguistic alignment in human-AI relational engagement. *Communication and Change*, 2, Article 11. <https://doi.org/10.1007/s44382-026-00032-5>
- Luo, S., Zhang, Z., Dai, H., & Zhang, D. (2026). Behavioral transfer in AI agents: Evidence and privacy implications. arXiv:2604.19925. <https://doi.org/10.48550/arXiv.2604.19925>
- Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 220–229. <https://doi.org/10.1145/3287560.3287596>
- Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39(2), 230–253. <https://doi.org/10.1518/00187209778543886>
- Pierce, J., Kalnikaite, V., Gupta, S., & Granger, B. (2026). Human-AI coordination zones: A framework for designing human-in-the-loop experiences with agentic AI. arXiv:2606.09848. <https://doi.org/10.48550/arXiv.2606.09848>
- Sheridan, T. B., & Verplank, W. L. (1978). Human and computer control of undersea teleoperators. Massachusetts Institute of Technology, Man-Machine Systems Laboratory. Technical Report ADA057655.
- Wu, Y. (2026). A three-wave study of human-AI fit and adaptive performance. *Technology in Society*, 87, 103386. <https://doi.org/10.1016/j.techsoc.2026.103386>
- Zou, H. P., Huang, W.-C., Wu, Y., Guo, J., Chen, Y., Miao, C., Nguyen, H. H., Zhou, Y., Zhang, W., Fang, L., Zhang, H., Wang, F., Zhang, P., He, L., Li, Y., Li, D., Jiang, R., & Yu, P. S. (2026). LLM-based human-agent collaboration and interaction systems: A survey. Findings of the Association for Computational Linguistics: ACL 2026, 36335–36364. <https://doi.org/10.18653/v1/2026.findings-acl.1811>

APPENDIX A. CLAIM BOUNDARY TABLE

The table states what this paper claims and what it does not claim. Reviewers are invited to identify any statement in the body that exceeds these boundaries.

C#	What this paper claims	What this paper does not claim
C1	The Operant Dyad proposes a distinct combined unit: the configured, task-bound, correction-shaped working relationship between one human and one bounded adaptive AI system.	That human-AI relationships or collaboration have not been described previously; that the terminology is unprecedented; or that the literature review is exhaustive.
C2	Repeated human correction is the mechanism that distinguishes the proposed Dyad from a series of isolated exchanges.	That every correction is retained, that correction always improves outputs, or that the mechanism has been empirically validated.
C3	A healthy Operant Dyad is confirmation-gated and preserves a practical opportunity for human inspection, refusal, and revision.	That all current systems provide meaningful confirmation or that a universal reflective-interval threshold has been established.
C4	‘Adaptive AI system’ includes systems whose later outputs are shaped by retained, retrieved, or deliberately reintroduced records of prior corrections.	That adaptation requires autonomous retraining or that all forms of memory, retrieval, and fine-tuning are equivalent.
C5	Relationship State names accumulated working fit, and declared state should be distinguished from unconfirmed inferred state.	That Relationship State can currently be fully specified, measured, or transferred, or that system inferences carry the same authority as human-confirmed constraints.
C6	Truxion is a proposed construct for recurring relationship drag caused by misalignment between human intent and system behavior.	That Truxion has a validated metric, known distribution, or established causal relationship with performance or well-being.
C7	The Portable Operant Packet names a proposed structured representation intended to support continuity across interruption or migration.	That a technical standard exists, that complete continuity is achievable, or that a destination system will reproduce identical behavior.
C8	Lock-In, Override, and Hidden Scaffolding are non-exhaustive structural failures that can degrade a Dyad.	That these are the only failure modes or that their frequency, severity, or causal ordering has been established.
C9	Adjacent HCI, human-factors, teaming, coordination, trust, alignment, and model-governance frameworks address important parts of the proposed object.	That those fields ignore relationships, longitudinal interaction, adaptation, or human control; or that the proposed framework supersedes them.
C10	GDPR Article 20 does not create a general right to export every inferred or derived profile, while the EU Data Act creates defined switching and portability duties for data-processing services.	That Relationship State is categorically unregulated, that the legal conclusion applies in every jurisdiction or service, or that this paper provides legal advice.
C11	Healthy and degraded Dyads may be distinguishable through observable indicators such as recurring correction, state drift, meaningful confirmation, traceability, and practical exit.	That the indicators are validated measures or that Dyad health can presently be assigned an objective score.
C12	Preserved human authorship is the governing condition for a healthy Operant Dyad.	That human authorship is always preserved in practice, that accepting an output proves authorship, or that current systems reliably protect it.
C13	The Operant Dyad is an operational and relational construct that can be studied without resolving machine consciousness or personhood.	That AI systems are conscious or nonconscious, sentient or nonsentient, moral patients, legal persons, or phenomenal subjects.
C14	The Operant Dyad is proposed as a useful unit for research, governance, and design.	That operationalization, measurement methodology, validated instruments, technical standards, or publication-level empirical findings already exist.
C15	The Operant Dyad integrates persistent relationship state, correction gating, state authority, portability, structural failure, and practical exit into a combined relationship-centric framework.	That each component is individually novel, that priority has been established, or that prior frameworks do not contain partially overlapping concepts.

This paper is a proposal. The empirical work required to validate, measure, compare, and govern the Operant Dyad begins after the conceptual claim is made—not before, and not by assertion.

AI-ASSISTED WRITING DISCLOSURE

During preparation and revision of this manuscript, the author used Claude (Anthropic) and OpenAI tools, including Codex, to assist with drafting, literature discovery, structural refinement, document production, and editorial review.

AUTHOR RESPONSIBILITY

The author reviewed, revised, and verified the manuscript and takes full responsibility for its analysis, citations, claims, limitations, and conclusions.

COMPETING INTERESTS

The author declares no competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

FUNDING

No external funding was received.

© 2026 Micheal Preble. Licensed under CC BY 4.0. Attribution is required.