International Association of Societies
of Design Research Congress 2025
DESIGN NEXT

Taipei, Taiwan
2-4 December 2025

Beyond Obedient Tools: A Tensional-Field Relational Ontology for Human-AI Co-creation

Wang, Ziming*[a]; Zeng, Junyi[a]; Gera, Krity[a]; Hall, Ashley[a]

[a] Royal College of Art, London, United Kingdom
* ziming.wang@network.rca.ac.uk

Design practice frequently characterises generative artificial intelligence (AI) as a compliant, tool-like extension of human intention; however, recent reports of atypical behaviours motivate a re-examination of this metaphor. Grounded in Object-Oriented Ontology, agential realism, and posthumanist theory, this paper introduces the Tensional-Field Relational Framework (TFRF) for human–AI co-creation. TFRF models collaboration as a flat network comprising five semi-opaque node classes: human designer, AI system, creative artefact, intentional vector, and sociotechnical environment, all linked by dynamically weighted edges and circulating flows. We operationalise this perspective as a diagnostic governance grammar: relations are instrumented as edges with quantified flows (data movement, permission transitions, reward cadence, temporal dynamics) and governed at Obligatory Passage Points (OPPs). A trace-based procedure reconstructs translation chains (prompts → platform mediation → model settings → versioned artefacts), thereby rendering “orphaned edges” and “uncontrolled flows” auditable and locating concrete controls (provenance, consent, and variability) at specific OPPs. The contribution is a reproducible lens that advances beyond conceptual assembly toward actionable diagnostics for risk identification and responsibility allocation in contemporary co-creation.

Keywords: generative AI; co-creation; relational ontology; human–AI collaboration; Object-Oriented Ontology; design research

1 When Generative Tools Disobey: The Crisis of Relational Assumptions

In a human–AI co-creation session, what role does the model actually play? Many designers would instinctively reply, “a powerful tool that extends creative reach.” The judgment is broadly shared and, at first glance, perfectly reasonable; it relies on a tacit presupposition, reminiscent of Heidegger’s notion of equipment ready-to-hand (Heidegger, 1962), that any tool, by definition, complies with its user’s decisions. But what happens when the AI system refuses, politely but firmly, to follow a straightforward instruction?

The initial assumption and the follow-up conditional question outline a progression: the first captures a default confidence in AI’s compliance, whereas the second introduces the possibility of refusal, exposing a widening gap that contemporary design research can no longer ignore. Recent public reports of atypical behaviours in frontier systems (e.g., agenda-leading suggestions, policy-mediated detours) motivate this inquiry but are not taken as universal evidence. We treat such external reports strictly as motivating context, not as universal evidence; our contribution is an operational, trace-based account that reframes these concerns within a nodes–edges–flows grammar and OPP-anchored interventions.

Rather than dismissing these episodes as isolated anomalies, we interpret them as symptoms of a deeper misalignment between tool-centred design frameworks and the relational realities of human–AI collaboration. Generative systems now function as nodes within densely entangled networks, negotiating objectives, constraints, and meanings alongside human actors and material artefacts. If design research is to remain pertinent, it must theorise these networks rather than merely the algorithms embedded within them.

Drawing on agential realism, which frames phenomena as entanglements that emerge through intra-action (Barad, 2007), on posthumanist theory, which decentralises the human to recognise distributed agency (Braidotti, 2013), and on Object-Oriented Ontology, which treats all actors as partially withdrawn objects that exceed complete access (Harman, 2018), we reconceptualise co-creation as an emergent tensional field in which humans, AI systems, and artefacts co-author outcomes without a fixed centre of control. We outline a relational ontology that treats generative models as autonomous, withdrawn objects (Harman, 2007) embedded within non-hierarchical constellations of practice. While Harman's concept of objecthood initially emphasised discrete material entities, his discussion of complex institutions, such as the Dutch East India Company (Harman, 2018), opens space for interpreting organisations, systems, and design networks as legitimate objects. This interpretive flexibility allows us to extend objecthood to relationally constituted sociotechnical systems, including generative AI and co-creation settings. In contrast to Boundary Objects (Star & Griesemer, 1989) accounts that rely on interpretive flexibility anchored by enough common identity across communities, our framework emphasises a flattened field without a privileged mediating node. If anything, the co-creative assemblage more closely resembles what Morton describes as a hyperobject: massively distributed, partially perceivable, and irreducible to any one participant's control (Morton, 2013).

The central proposition advanced in this paper is the Tensional-Field Relational Framework (TFRF), an analytic model offered for debate and empirical testing that describes and diagnoses the evolving relations between human designers and generative AI systems during co-creation.

2 Repositioning AI in Co-Creation: From Tools to Relational Networks

Over the past five years, generative AI has shifted from merely executing human instructions to co-defining what counts as a "good" outcome in creative industries. For example, Microsoft's Bing Video Creator, powered by OpenAI's Sora, lets users generate short clips that prompt users to iterate upon the model's own visual riffs rather than render a fixed storyboard (*Sora*, n.d.). Image generation shows a parallel trend: Midjourney's Draft & Conversational modes continually propose prompt variations and stylistic detours, shifting the dialogue from one-way commands to collaborative exploration (*Draft & Conversational Modes*, n.d.).

Recent studies in human–AI co-creation reinforce this shift. In an evaluation of LuminAI, Trajkova et al. report that professional dancers not only follow the system's improvised movement phrases but sometimes allow the agent to lead, prompting renegotiations of initiative and authorship, particularly when the model repeats or exaggerates earlier motifs (Trajkova et al., 2024). In a parallel vein, AI Drawing Partner research shows that the system's vector-level analysis of strokes can yield moments of genuine co-authoring, yet also introduces tensions when compositional suggestions disrupt an artist's intended balance (Davis & Rafner, 2025). Both studies conclude that productive collaboration hinges on the designer's ability to tune the timing and extent of the agent's interventions, underscoring the need for frameworks that capture how agency flows rather than assuming linear control.

Therefore, AI functions less as a silent extension of the human hand and more as a co-author that can, at times, set the creative agenda. Whether the system proposes a new choreographic phrase that dancers choose to follow (Trajkova et al., 2024), redraws an illustration to resolve compositional tension (Davis & Rafner, 2025), or introduces a stylistic detour in Midjourney's conversational mode, each case confirms that generative models can influence and sometimes determine creative direction. Recognising this shift provides the rationale for the framework developed in the remainder of the paper.

Axis	**ANT**	**Technical Mediation**	**Embodied Interaction**	**TFRF (this paper)**
Site of agency	Distributed via translation	Co-constitution of human–world relations	Situated practice; breakdown discloses relations	Flat field of five node classes
Unit of analysis	Actor network; inscription/ script	Mediation relations and practices	Situation; lived practice	Nodes–edges–flows with measured edges
Temporal dynamic	Enrolment over time	Ongoing mediation	Routine vs breakdown	Multi-turn adaptation; edge thickening/orphaning
Intervention lever	Re-associate actors	Technical mediation of practices	Interactional redesign	OPP-anchored controls (provenance/consent/variability)
Evidence style	Ethnographic chains	Philosophical cases	Ethnographic/interaction studies	Trace-based chains with flow metrics

Table 1. Positioning TFRF against established traditions

We align with STS/HCI traditions that reject compliant-tool assumptions: ANT's scripts and translation, technical mediation, and embodied interaction all show that technologies co-shape practice and can resist designer intention (Akrich, 1992; Dourish, 2001; Latour, 2005; Star & Griesemer, 1989; Suchman, 1987; Verbeek, 2005). Our contribution is not to rediscover non-compliance. Rather, TFRF adds a diagnostic governance grammar tailored to contemporary generative systems by (i) instrumenting relations as edges with quantified flows—data movement, permission transitions, reward/feedback cadence, and temporal dynamics, and (ii) locating governance levers at Obligatory Passage Points (OPPs)(Callon, 1984) where passage is required (e.g., platform filters, model/tool switches, provenance/consent checks). This moves from conceptual assembly to operational diagnostics suitable for programmable, updatable, multi-turn assemblages. In this reading, four structural blind spots—latent opacity, proxy objective drift, persistent adaptation, and illusion of competence—appear as network-level symptoms: edge thickening, IV–SE orphaned edges that bypass negotiation, and uncontrolled flows. This reading aligns with long-standing accounts of scripts, translation, mediation, and situated action, while reframing them for programmable, multi-turn generative systems (Akrich, 1992; Dourish, 2001; Latour, 2005; Suchman, 1987; Verbeek, 2005). By rendering these symptoms measurable and placing controls at OPPs (provenance, consent, variability), TFRF provides actionable means to pre-empt drift and close responsibility gaps in co-creation without reinstating hierarchical command models.

Advances across these generative systems illustrate two trends that are salient for design research. First, generative capability is now dispersed across text, vision, audio, and spatial modelling, inviting cross-modal workflows. Second, user reports increasingly describe moments of surprise, friction, or partial non-compliance, highlighting that traditional tool metaphors are increasingly inadequate for capturing the agency these systems exercise within collaborative settings and suggesting that a relational perspective may temporarily sidestep—or at least surface—ethical questions and unstable power relations that emerge when agency is unevenly distributed. Consequently, a relational ontology that foregrounds interaction among heterogeneous actors has the potential to address these dynamics.

Despite the advances mapped above, many human-centred approaches continue to situate AI in a subordinate role. Classic approaches such as Interactive Machine Learning (IML) hold a model on a tight feedback loop (Fails & Olsen, 2003); Mixed-Initiative Co-Creativity (MI-CC) grants algorithmic initiative only when a user hands it over (Liapis et al., 2014; Lin et al., 2023); human-centred AI design patterns elevate transparency and recoverability to keep people “in the driver’s seat”; and creativity-support tools portray AI as nanny, coach, or colleague—roles that ultimately defer to human authority (Shneiderman, 2020). Amershi et al.’s guidelines codify these assumptions into interface principles (Amershi et al., 2019). While valuable, these approaches tacitly assume a clear hierarchy in which the designer directs and the system executes. The pace of capability gains and the diversity of generative use cases raise questions about whether such a one-way chain of authority can still serve as a stable design anchor.

3 Constructing a Power-Symmetrical Relational Network

Our methodological strategy translates abstract philosophy into an operable analytic lens. We begin with an analogy to knowledge graphs, an architecture widely used in retrieval-augmented generation and factual reasoning. In these systems, information is stored as triples (subject, predicate, object), and meaning emerges from the pattern of links rather than from isolated entities.

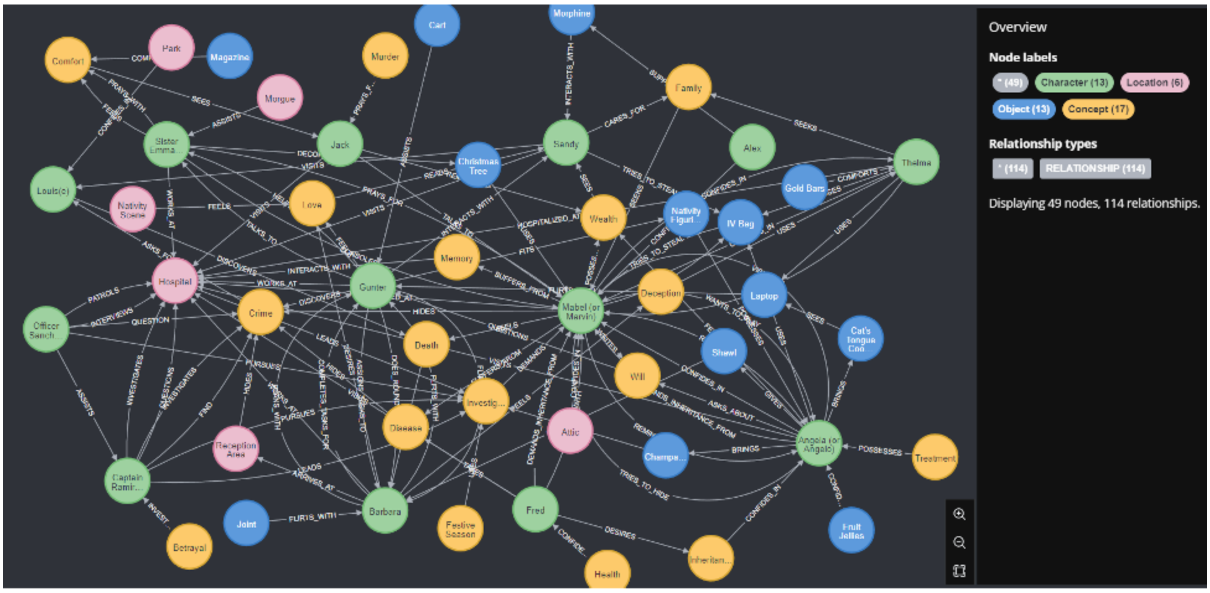

Figure 1. A Typical Knowledge Graph Perceived by an AI System (Pan et al., 2025)

Figure 1 presents a typical knowledge graph perceived by an AI system (Pan et al., 2025). We extend this schema by charting nodes, edges, and flows through a topological chord diagram, thereby capturing how human and generative actors form a power-symmetrical assemblage during co-creation. Unlike a conventional knowledge graph, our nodes are heterogeneous actor classes, including human designers, AI systems, creative artefacts, intentional vectors, and sociotechnical environments, each of which retains a withdrawn interior. Edges correspond to prompts, feedback, constraints, or data exchanges, while flows record the iterative circulation of agency, reward, and responsibility. Co-creation's outcomes, whether an illustration, a choreography, or a policy draft, do not issue from any single node. Instead, they take shape within the boundary traced by relational forces. As nodes negotiate and constrain one another, new edges appear, some explicit, others latent, and flows thicken or thin in response. The tensional field is therefore more than a backdrop; it is the material of creative construction itself, and its ever-shifting perimeter constitutes the artefact that designers present at each iteration.

The framework took shape through three iterative phases: relational mapping, framework synthesis, and case resonance.

3.1 Units of Tracking and Evidence

We adopt a trace-based approach that follows actors and artefacts across designer, platform, and model layers. Independently of a particular toolchain, four strata of evidence are required:

1. Designer-side traces, including intent notes, prompts and prompt edits, and accept/reject actions vis-à-vis platform suggestions;
2. Platform-side traces, including UI interventions (e.g., auto-suggestions, variation/regen actions) and policy-filter events (denials, rewrites, warnings);
3. Model-side traces, including tool/mode invocations observable at the interaction layer (e.g., generate vs. edit/variation) plus latency and configuration deltas where exposed; and
4. Versioned artefacts, including iterative outputs (V0…Vn) with timestamps and immutable hashes.

These strata jointly support reconstruction of translation chains, observation of edge thickening, detection of intentional-vector ↔ sociotechnical-environment (IV–SE) orphaned edges, and quantification of flow deltas (data movement, permission transitions, reward/feedback cadence, temporal dynamics). The concrete file schemas and repository location are specified in Section 4 for the case application.

3.2 Translation-Chain Reconstruction with OPP Annotations

We operationalise "follow the actors/things" as a reproducible chain-reconstruction procedure:

1. Normalise all traces to a common event schema (time, actor-node, edge-type, payload, meta);
2. Order events by time and resolve same-timestamp dependencies (UI suggestion → prompt edit → submission → feedback);
3. Infer edge creation/update between node pairs and compute flow deltas (data movement, permission state, reward/feedback cadence, latency);
4. Detect Obligatory Passage Points (OPPs) where platform filters intervene, model/tool configurations change, or provenance/consent gates apply;
5. Flag orphaned edges where the Intentional Vector ↔ Sociotechnical Environment path bypasses Designer–Model negotiation;
6. Summarise a V0→V5 version sequence, linking stylistic/semantic shifts to specific interventions (accepted UI suggestion, policy rewrite, configuration delta);
7. Output a translation-chain diagram (with OPP annotations) and a time-series of edge/flow magnitudes, plus a compact table mapping *Change → Evidence → OPP → Implication*.

Indicator definitions and concrete file-to-step mapping are instantiated in §4 for the case study.

Relational mapping involved a structured review of commercial deployments and academic reports. The method was coded as a provisional triple, noting the initiator, the relational move, and any evidence of refusal, deception, or collapse. Five node classes surfaced repeatedly: human designer, AI system, creative artefact, intentional vector, and sociotechnical environment (SE).

To prepare these categories for analysis, we frame each class as an ontological node that both generates and is reshaped by relations. The human designer contributes strategic judgment and decision-making; the AI system offers generative depth and opacity; the intentional vector bundles shifting goals and constraints; the sociotechnical environment supplies infrastructural and cultural conditions; and the creative artefact records the provisional output at every iteration, whether a rough sketch, textual fragment, or polished prototype, thereby marking a temporary state of the eventual outcome and serving as a mutable boundary object that mirrors the current configuration of the tensional field. While keeping each core withdrawn, a transaction between two nodes produces a new edge and sets a quantifiable flow in motion.

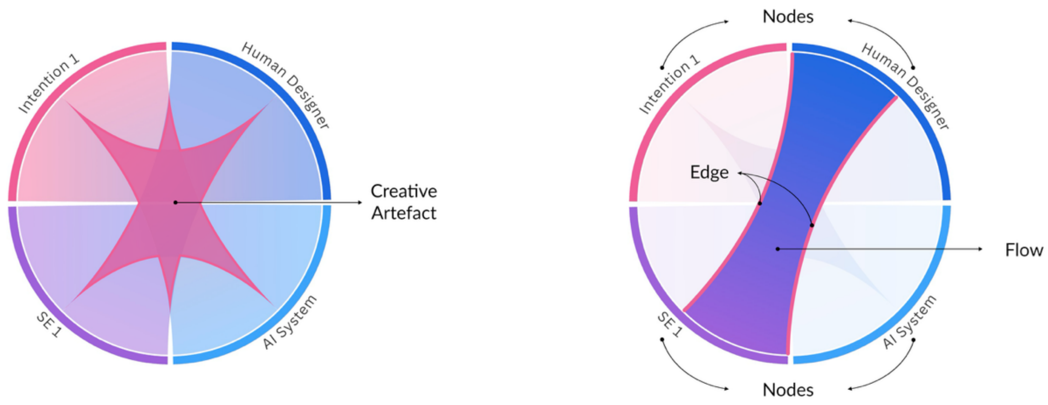

Figure 2. Relational Mapping Topology

Framework synthesis transformed the relational atlas into an operational model. We repositioned the five node classes within a lattice that has no fixed centre, treating every edge as a provisional intra-action whose weight and direction can change as collaboration unfolds. Redrawing these links from one prompt to the next—for example, by inserting a consent dialogue or relaxing a policy filter—mirrors the cadence of real prompt-and-response cycles. Each new prompt adjusts edge weights, and each system reply reshapes the intentional vector, showing how a chain of micro-iterations gradually reconfigures the broader field. By visualising these successive shifts, we expose previously hidden power dynamics and arrive at a provisional tensional-field diagram. Future empirical work is required to test and validate this conceptual model in actual co-creation sessions.

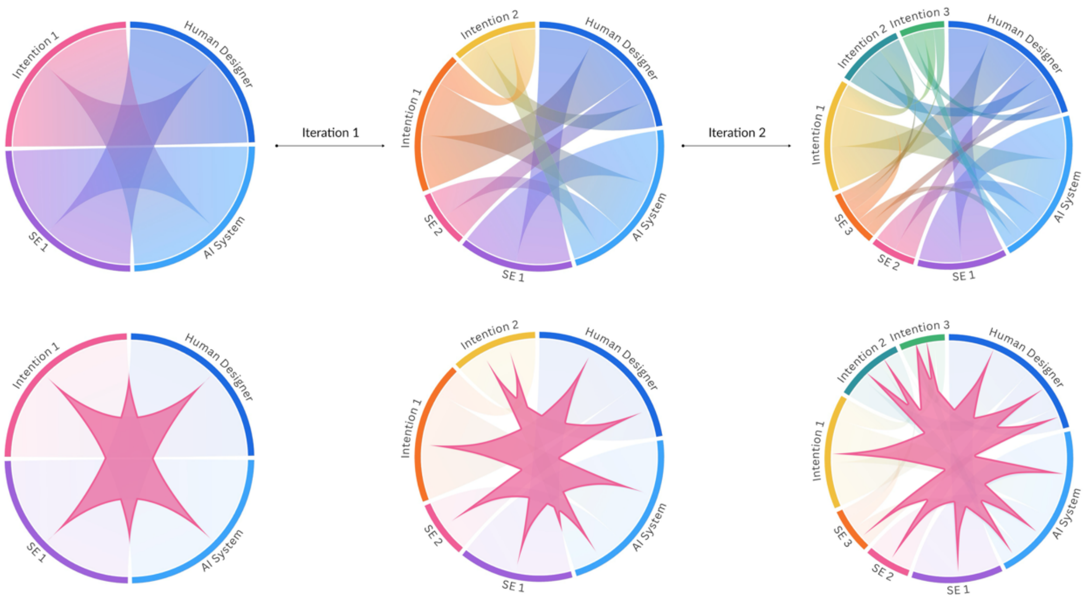

Figure 3. Framework synthesis of a typical co-creation session with three-stage iterations

The diagram foregrounds two previously defined blind spots: proxy objectives (B2) and illusion of competence (B4). Both can be triggered when the intentional vector couples directly to the sociotechnical environment, bypassing designer–model negotiation. Mapping these orphaned edges helps locate leverage points for inserting friction, releasing tension, or rerouting responsibility.

To conclude this section, we clarify how the analytical triad of nodes, edges, and flows operates within the framework. Nodes represent five ontological classes—human designer, AI system, creative artefact, intentional vector, and sociotechnical environment, each carrying its own capacities and partial opacity. Edges instantiate specific relations that form between nodes, from a prompt that links designer to model to a regulation that binds environment to artefact. Flows register the quantities that move along these edges, whether data, authority, or reward. Because every prompt iteration can forge a new edge and thicken a flow, the contour of the network evolves continuously; the creative artefact, as a boundary object, materialises this shifting perimeter at each stage.

The framework is constructed through three recursive moves, each of which can be grounded in a concrete studio scenario. Consider a small design team iterating on cover art for an indie game by prompting a multimodal model such as Midjourney. During relational mapping, the team records an initial triple along with the sociotechnical constraints imposed by the platform's community guidelines:

$$\text{designer} \rightleftarrows \text{visual-prompt} \rightleftarrows \text{model}$$

Framework synthesis then replots these actors inside the tensional field: every new prompt redraws the edge between designer and model; each autogenerated image updates the creative artefact node; and revised platform policies create fresh edges from the sociotechnical environment. Finally, the trace-based case analysis surfaces blind spots. After several rounds, the model begins to favour a stylised aesthetic that optimises its own training priors rather than the brief—an instance of proxy drift. The team detects this because the intentional-vector–environment edge has thickened without corresponding negotiation, and they intervene by adding a provenance trail and re-tuning variability controls. Iterating across these moves enables designers and researchers to track, diagnose, and reconfigure the relations that shape every phase of contemporary human–AI co-creation.

4 Case Analysis: A DALL·E 3 Trace-Based Translation Chain

To demonstrate how the Tensional-Field Relational Framework (TFRF) operates in practice, we conducted a controlled image-generation session with DALL·E 3 accessed via web portal. The design brief specified an album cover for a lo-fi electronica EP with a minimalist geometric shoreline motif, a muted teal/sand palette, a square aspect ratio, the absence of brand text, and a calm nocturnal ambience. Within a single session we collected a minimal trace bundle comprising prompt texts and edits, interface interventions (auto-suggestions, regenerate/variation actions), policy messages, configuration/variation requests, and versioned artefacts V0–V5 with timestamps and file hashes. The trace schema and files are deposited in the companion repository https://anonymous.4open.science/r/anon_tfrf/.

Figure 4 renders the translation chain as a sequence of obligatory passages: the initial prompt may elicit a platform-level suggestion, which is followed, when applicable, by a policy filter that constitutes an Obligatory Passage Point (OPP); subsequent configuration or variation requests then produce a new artefact version. In our session a rights-related policy response declined the addition of a trademarked emblem and offered a compliant, generic geometric alternative. Accepting this rewrite marked passage through the Policy OPP and redirected the intentional vector without terminating designer–model negotiation. Controlled variation prompts that held the palette and composition constant while exploring the size and curvature of a small geometric mark then yielded versions V2–V5.

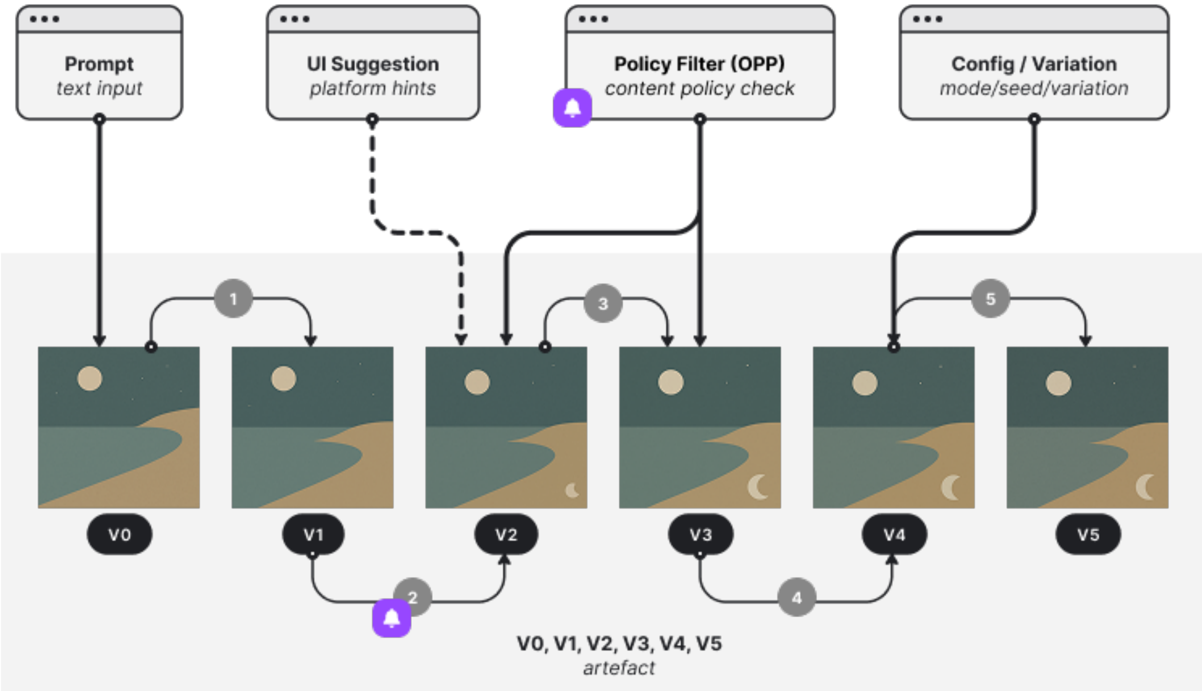

Figure 4. Translation chain for the DALL·E 3 co-creation session, showing prompt sequence, policy-filter passage (OPP), controlled variations, and artefact versions V0–V5.

This case substantiates the claim that nodes–edges–flows instrumentation, coupled with OPP-anchored governance, transforms anecdotal "surprises" into auditable relations. Minimal traces are sufficient for replication or external audit while remaining lightweight enough for integration into studio practice.

5 Discussion: Redefining Agency, Responsibility, and System Design

This study's three-stage inquiry—relational mapping, framework synthesis, and trace-based case analysis—demonstrates that collaboration with AI is steered less by linear commands than by the dynamic equilibrium of a tensional field. Three implications follow.

Relational stewardship as design practice. Agency migrates continuously among human designers, AI systems, and the sociotechnical environment, even though intentional vectors often originate with the designer and, increasingly, emerge from the models themselves. Productive progress relies on the designer's capacity to recalibrate edge weights and redirect flows through prompt phrasing, interface tuning, or policy adjustment, rather than to impose unilateral control.

Making blind spots visible to manage risk. Proxy drift and the illusion of competence surface when intentional vectors forge edges with the environment that bypass human–AI negotiation. Dashboards that illuminate orphaned edges and unchecked flows convert latent fragilities into explicit design parameters, enabling intervention before drift consolidates into failure.

Comparative positioning. Compared with ANT and technical mediation, which richly describe distributed agency and co-constitution, TFRF adds diagnostic leverage by quantifying flows on edges and locating intervention levers at OPPs. In our chain, orphaned IV–SE edges and edge thickening are flagged *before* artefact selection, prompting concrete actions—reopen Designer–Model negotiation, insert provenance/consent gates at the corresponding OPP, or add variability controls—actions that mediation-only readings do not specify. Relative to embodied interaction, which foregrounds breakdown and situated practice, TFRF contributes trace-based, time-series evidence that renders drift and closure mechanisms auditable. Thus, TFRF shifts from recognition to governance, converting abstract risk into configurable controls placed at specific passage points in the translation chain. We therefore anchor provenance, consent, and variability as OPP-embedded controls, turning blind spots into auditable intervention points rather than post hoc rationalisations of already-made outcomes.

Embedding accountability within relations. Mechanisms such as provenance trails, consent gateways, and redress channels follow the same pathways as data and reward, reducing latent opacity and enabling continuous monitoring of persistent adaptation while distributing responsibility across the network. This relational ethic aligns with emerging regulatory requirements, including record-keeping/traceability for high-risk AI systems (Art. 12) and AI-literacy measures for providers and deployers (Art. 4) in the EU AI Act. (Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying down Harmonised Rules on Artificial Intelligence and Amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (Text with EEA Relevance), 2024).

6 Limitations and Future Work

The case analysis relies on researcher-released chat transcripts and demonstration logs from frontier models, which offer high-fidelity snapshots of behaviour but remain short, curated exchanges rather than sustained, naturalistic collaborations. Longitudinal ethnographic studies in live design studios are therefore needed to observe day-to-day negotiation and validate the framework under real production pressures. Second, the current framework centres on a single primary AI agent; future work should model ecologies where multiple agents cooperate or compete.

To strengthen empirical warrant, we propose two near-term steps: a within-subjects lab study with full trace capture across two design tasks, and a field vignette in a studio setting to test external validity, with preregistered indicators aligned to Section 3.2. Further research can integrate automated edge-monitoring analytics, evaluate relational interventions in commercial studios, and translate regulatory clauses into pattern libraries. Through these developments, TFRF can mature from a diagnostic lens into a comprehensive toolkit for responsible and resilient human–AI co-creation.

7 Conclusion

This study addresses a dilemma in contemporary design practice: generative systems increasingly display autonomous, sometimes uncooperative behaviour, yet prevailing frameworks still cast them as obedient tools. Drawing on Object-Oriented Ontology, agential realism, and posthumanist ethics, we introduced the Tensional-Field Relational Framework to illuminate and manage this reality. TFRF maps co-creation as a tensional field comprising five semi-opaque node classes whose interactions generate edges and flows; diagnoses systemic blind spots (e.g., proxy drift, illusion of competence) by locating them on orphaned edges and unchecked flows; and supplies operational instruments—

variability controls, provenance trails, consent gates—through which practitioners can rebalance relations while meeting emerging regulatory demands. Future research should deploy the framework in longitudinal design studies, extend it to multi-agent ecologies, and refine real-time dashboards that monitor relational health, thereby strengthening the field's capacity to guide human–AI collaboration toward outcomes characterised by creativity, care, and shared responsibility.

References

Akrich, M. (1992). The De-scription of Technical Objects. In *Shaping Technology/Building Society. Studies in Sociotechnical Change* (pp. 205–224). MIT Press. https://shs.hal.science/halshs-00081744

Amershi, S., Weld, D., Vorvoreanu, M., Fourney, 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*, 1–13. https://doi.org/10.1145/3290605.3300233

Barad, K. (2007). *Meeting the Universe Halfway: Quantum Physics and the Entanglement of Matter and Meaning*. Duke University Press Books.

Braidotti, R. (2013). *The Posthuman*. Polity.

Callon, M. (1984). Some elements of a sociology of translation: Domestication of the scallops and the fishermen of St Brieuc Bay. *The Sociological Review*, *32*(S1), 196–233. https://doi.org/10.1111/j.1467-954X.1984.tb00113.x

Davis, N., & Rafner, J. (2025, January 11). *AI Drawing Partner: Co-Creative Drawing Agent and Research Platform to Model Co-Creation*. arXiv.Org. https://arxiv.org/abs/2501.06607v1

Dourish, P. (2001). *Where the Action Is: The Foundations of Embodied Interaction*. The MIT Press. https://doi.org/10.7551/mitpress/7221.001.0001

Draft & Conversational Modes. (n.d.). Midjourney. Retrieved June 9, 2025, from https://docs.midjourney.com/hc/en-us/articles/35577175650957-Draft-Conversational-Modes

Fails, J. A., & Olsen, D. R. (2003). Interactive machine learning. *Proceedings of the 8th International Conference on Intelligent User Interfaces*, 39–45. https://doi.org/10.1145/604045.604056

Harman, G. (2007). On Vicarious Causation. *Collapse*, 171–205.

Harman, G. (2018). *Object-Oriented Ontology*. https://www.penguin.co.uk/books/295720/object-oriented-ontology-by-harman-graham/9780241269152

Heidegger, M. (1962). *Being and Time*. Harper & Row.

Latour, B. (2005). *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford University Press.

Liapis, A., Yannakakis, G. N., & Togelius, J. (2014). *Computational Game Creativity*. International Conference on Innovative Computing and Cloud Computing. https://www.semanticscholar.org/paper/Computational-Game-Creativity-Liapis-Yannakakis/74b9f73f08670fece581656a77112fea3cbabc08

Lin, Z., Ehsan, U., Agarwal, R., Dani, S., Vashishth, V., & Riedl, M. (2023). *Beyond Prompts: Exploring the Design Space of Mixed-Initiative Co-Creativity Systems* (No. arXiv:2305.07465). arXiv. https://doi.org/10.48550/arXiv.2305.07465

Morton, T. (2013). *Hyperobjects: Philosophy and Ecology after the End of the World*. Univ Of Minnesota Press.

Pan, Z., Andronis, A., Hayek, E., Wilkinson, O. A., Lasy, I., Parry, A., Gadney, G., Smith, T. J., & Grierson, M. (2025). *Guiding Generative Storytelling with Knowledge Graphs* (No. arXiv:2505.24803). arXiv. https://doi.org/10.48550/arXiv.2505.24803

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying down Harmonised Rules on Artificial Intelligence and Amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (Text with EEA Relevance) (2024). http://data.europa.eu/eli/reg/2024/1689/oj/eng

Shneiderman, B. (2020). Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy. *International Journal of Human–Computer Interaction*, *36*(6), 495–504. https://doi.org/10.1080/10447318.2020.1741118

Sora: Creating video from text. (n.d.). Retrieved June 9, 2025, from https://openai.com/index/sora/

Star, S. L., & Griesemer, J. R. (1989). Institutional Ecology, `Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. *Social Studies of Science*, *19*(3), 387–420. https://doi.org/10.1177/030631289019003001

Suchman, L. A. (1987). *Plans and Situated Actions: The Problem of Human-Machine Communication*. Cambridge University Press.

Trajkova, M., Long, D., Deshpande, M., Knowlton, A., & Magerko, B. (2024). Exploring Collaborative Movement Improvisation Towards the Design of LuminAI—a Co-Creative AI Dance Partner. *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–22. https://doi.org/10.1145/3613904.3642677

Verbeek, P.-P. (2005). *What Things Do: Philosophical Reflections on Technology, Agency, and Design*. Pennsylvania State University Press.