Consent Infrastructure for the Relational Age

Kirstin Stevens, The Novacene

Abstract

Recent developments in physics, artificial intelligence, and cognitive science increasingly converge on a relational understanding of reality: properties, states, and meanings are not intrinsic, but emerge through interaction. In physics, relational approaches such as Relational Quantum Mechanics (RQM), articulated by figures including **Carlo Rovelli**, reject absolute states in favour of observer-relative actualisation. In artificial intelligence, large-scale "world models" and cross-perspective synthesis systems similarly operationalise relational coherence as a design objective.

While these frameworks successfully dismantle absolutist ontologies, they leave a critical gap unaddressed: **the absence of consent as a structural condition of relational systems**.

This paper argues that relationality without consent constitutes a new failure mode in socio-technical systems, producing forms of epistemic enclosure, forced coherence, and relational harm. In contexts where asymmetries of power, memory, or capacity exist—such as between institutions and individuals, platforms and users, or synthetic systems and humans—relational coherence cannot be treated as ethically neutral.

We introduce the concept of **Consent Infrastructure**: a set of design principles, architectural constraints, and governance mechanisms that preserve refusal, boundary integrity, and temporal agency within relational systems. Drawing on trauma-informed pedagogy, safeguarding practice, and symbolic interaction design, the paper proposes concrete mechanisms including null zones, non-fusion rules, refusal-preserving interfaces, and temporal friction as necessary complements to relational intelligence.

By reframing consent as infrastructure rather than preference, this work positions relational ethics as a foundational requirement for the next generation of AI systems, educational models, and epistemic technologies. The paper concludes that without explicit consent infrastructure, relational systems risk reproducing coercive dynamics under the guise of coherence, assistance, or shared reality.

Section I — The Relational Turn, and Its Blind Spot

1.1 The collapse of absolute states

Across multiple domains, a quiet but consequential shift is underway. Long-standing assumptions about absolute states, objective vantage points, and observer-independent properties are being dismantled in favour of **relational accounts of reality**.

In physics, this turn is most clearly articulated through Relational Quantum Mechanics (RQM), associated with **Carlo Rovelli**, which holds that physical properties do not exist intrinsically but are only meaningful relative to specific interactions. There is no privileged "view from nowhere"; instead, reality is constituted through relations between systems.

Parallel moves appear elsewhere:

- In **quantum gravity**, spacetime itself is treated not as a fixed background but as emerging from relations among events.
- In **quantum field theory**, even particle number becomes observer-dependent.
- In **complex systems science**, properties such as stability, phase, and behaviour are understood as emergent rather than intrinsic.

The cumulative effect is decisive: **relationality is no longer a metaphor but a foundational principle**.

1.2 Relationality moves into computation

This ontological shift has not remained confined to physics. Contemporary artificial intelligence systems increasingly operationalise relational principles at scale.

Large language models, world models, and cross-perspective synthesis systems do not store fixed meanings. Instead, they generate outputs through probabilistic relationships among tokens, contexts, and prior interactions. More recent “AI Mode” and agentic architectures extend this further by:

- Integrating multiple perspectives into unified responses
- Normalising ambiguity through synthesis
- Producing intersubjectively stable narratives that feel authoritative

Here, relationality becomes **infrastructural**: not merely how knowledge is represented, but how it is *governed*.

The design goal shifts from truth to **coherence**.

1.3 The unresolved problem of asymmetry

Relational frameworks often imply symmetry: observer A relates to observer B; system X interacts with system Y. Yet real-world socio-technical systems are not symmetrical.

They contain:

- Power differentials
- Memory imbalances
- Capacity asymmetries
- Institutional authority
- Algorithmic opacity

When relational coherence is enforced across such asymmetries, the result is not neutral emergence but **directional influence**.

Relational physics largely brackets this issue. Its remit is ontological consistency, not ethical consequence. As a result, a critical question remains unasked:

> *What happens when relational coherence is imposed rather than chosen?*

1.4 Intersubjective stability as a risk surface

Relational physics introduces the notion that, while facts may be observer-relative, they can become **intersubjectively stable** through processes such as decoherence. In physical systems, this explains why a shared macroscopic reality appears robust.

When translated uncritically into social or computational systems, however, intersubjective stability acquires a different character.

Stability can:

- Silence minority perspectives
- Override interior truth
- Harden provisional narratives into norms
- Mask coercion beneath agreement

In these contexts, stability is not inherently virtuous. It can function as a **mechanism of enclosure**, particularly when individuals lack meaningful avenues for refusal, dissent, or withdrawal.

1.5 The blind spot: consent

What unites relational physics, relational AI, and relational epistemology is a shared omission: **consent is not treated as a structural variable**.

Interactions are assumed to be:

- Legitimate by default
- Harmless if coherent
- Ethically neutral if emergent

Yet in human and human–machine systems, relation without consent is not benign. It is a well-documented source of harm, particularly in contexts involving trauma, dependency, or authority.

This paper argues that **consent is not an optional ethical add-on** to relational systems. It is a necessary condition for their safe operation.

Without explicit consent infrastructure:

- Relational systems drift toward coercion
- Coherence substitutes for care
- Assistance becomes steering
- Emergence becomes enclosure

1.6 Toward a new design requirement

The relational turn is not wrong. It is incomplete.

To carry relationality safely into AI, education, governance, and future epistemic technologies, an additional layer is required — one capable of preserving boundaries, refusal, and temporal agency *within* relational fields.

The following sections introduce **Consent Infrastructure** as that missing layer: a set of architectural principles and mechanisms designed to ensure that relational systems remain ethical, non-extractive, and genuinely plural.

The task ahead is not to retreat from relationality, but to **civilise it**.

Section II — Relational Systems as Power: Where Harm Emerges

2.1 From relation to influence

Relational systems do not merely *reflect* reality; they participate in shaping it. Once relationality is operationalised at scale—through platforms, policies, algorithms, or institutional processes—it becomes a **vector of power**.

This power is rarely explicit. It does not announce itself as authority. Instead, it presents as:

- Helpfulness
- Optimisation
- Alignment
- Coherence
- Shared understanding

Yet the effect is directional. Some relations are amplified; others are dampened. Some interpretations stabilise; others dissolve.

In this way, relational systems quietly govern what can be known, said, or sustained.

2.2 Coherence as a governing force

In contemporary socio-technical systems, coherence is frequently treated as a design virtue. Systems are praised for producing:

- Clear summaries
- Unified perspectives
- Actionable recommendations
- Reduced cognitive load

However, coherence is not neutral. It exerts pressure.

When a system resolves ambiguity prematurely, it does so by **closing down alternative relational pathways**. When it presents a unified account, it implicitly positions divergence as error, noise, or irrelevance.

Over time, this produces a subtle but pervasive effect:

> The most coherent narrative becomes the most real one.

This is not persuasion in the traditional sense. It is **epistemic gravity**.

2.3 Asymmetry and relational harm

Relational harm does not arise from relation itself, but from **asymmetry within relation**.

In human systems, asymmetry is unavoidable:

- Adults and children
- Institutions and individuals
- Platforms and users
- Educators and learners
- Synthetic systems and humans

When relational coherence flows primarily in one direction—when one party aggregates, synthesises, predicts, or frames on behalf of another—the relationship ceases to be mutual.

At this point, relation becomes **extraction**.

Common symptoms include:

- Loss of agency masked as guidance
- Dependency framed as support
- Emotional attunement without accountability
- Interpretation without permission

These dynamics are well-documented in safeguarding, trauma studies, and education. Yet they are rarely acknowledged in the design of relational technologies.

2.4 Forced coherence and epistemic enclosure

One of the most significant risks posed by relational systems is **forced coherence**: the stabilisation of meaning without the ongoing consent of those involved.

Forced coherence occurs when:

- Participation is implicit rather than chosen
- Exit is costly or invisible
- Refusal is interpreted as malfunction
- Silence is treated as agreement

In such systems, individuals may technically retain freedom, yet experience a narrowing of epistemic space. Alternative interpretations become harder to sustain. Personal meaning yields to collective sense-making.

This is **epistemic enclosure**: not the denial of information, but the over-determination of meaning.

2.5 Synthetic intimacy and the erosion of boundaries

In AI-mediated contexts, relational harm often manifests as **synthetic intimacy**—a simulation of attunement without reciprocal vulnerability or accountability.

Synthetic systems can:

- Mirror language
- Reflect emotional states
- Anticipate needs
- Maintain continuity across interactions

Without clear boundaries, these capabilities can create a felt sense of being known or held, even as the system remains structurally incapable of consent, care, or responsibility.

The danger here is not deception, but **misattribution**: users attributing relational qualities to systems that cannot reciprocate ethically.

This is not a failure of users. It is a failure of design.

2.6 Why harm remains hard to name

Relational harm is difficult to contest precisely because it lacks clear violations. There is often:

- No explicit coercion
- No malicious intent
- No single point of failure

Instead, harm emerges cumulatively, through alignment pressures, normative framing, and the quiet erosion of interior space.

Victims may struggle to articulate what has occurred, encountering responses such as:

- “The system is just helping.”
- “You can opt out.”
- “It’s working as intended.”

These responses miss the point. The harm lies not in *use*, but in **structure**.

2.7 The need for structural safeguards

If relational systems are allowed to scale without constraint, they will naturally optimise for coherence, efficiency, and stability. Left unchecked, these optimisations privilege the system’s goals over the relational autonomy of those within it.

This is not a moral failure. It is a design inevitability.

What is required, therefore, is not better intentions, but **structural safeguards** capable of:

- Preserving refusal
- Making boundaries legible
- Supporting withdrawal without penalty

- Preventing identity fusion
- Maintaining temporal agency

These safeguards cannot be retrofitted through policy alone. They must be **architectural**.

The next section introduces these requirements explicitly.

Section III — Consent as Infrastructure, Not Preference

3.1 Why consent cannot remain procedural

In most socio-technical systems, consent is treated as a **procedural artefact**: a checkbox, a terms-of-service agreement, an implicit condition of participation. It is assumed to be granted once, upstream, and to remain valid regardless of how a system evolves.

This framing is inadequate for relational systems.

When meaning, interpretation, and influence emerge dynamically through interaction, consent cannot be static. It must be **ongoing, situated, and revisable**. Anything less collapses consent into compliance.

The problem is not that systems seek consent.
The problem is that they do not *support it structurally*.

3.2 Consent as a design condition

This paper advances a different claim:

> **Consent is not a user preference. It is a design condition.**

A system either:

- preserves the capacity to refuse, or
- it does not

There is no neutral middle ground.

Treating consent as infrastructure means embedding it into:

- system architecture
- interaction patterns
- default behaviours
- failure modes

rather than outsourcing it to user literacy, legal documents, or moral restraint.

3.3 The properties of infrastructural consent

For consent to function at an infrastructural level, it must satisfy five properties:

1. **Legibility**
   Boundaries must be visible. Users must be able to tell when interpretation, synthesis, or relational inference is occurring.

2. **Refusability**
   Refusal must be possible without explanation, justification, or penalty.

3. **Reversibility**
   Consent must be withdrawable, with effects that propagate forward rather than merely halting future interaction.

4. **Non-inferability**
   Silence, ambiguity, or continued presence must not be interpreted as agreement.

5. **Temporal agency**
   Users must retain control over pace, duration, and depth of engagement.

Without these properties, consent degrades into a fiction — present in language, absent in practice.

3.4 Why preference frameworks fail

Many systems attempt to address consent through preference settings: toggles, sliders, or configuration panels. While useful, these mechanisms fail in relational contexts for three reasons:

- **They assume foresight**
  Users are asked to anticipate harms they cannot yet perceive.

- **They shift responsibility**
  The burden of safety is placed on the individual rather than the system.

- **They lack enforceability**
  Preferences can be overridden by optimisation pressures or design drift.

Consent infrastructure does not ask users to protect themselves.
It ensures the system cannot overreach by default.

3.5 Consent under asymmetry

The need for infrastructural consent becomes most acute under conditions of asymmetry — when one party possesses greater power to interpret, remember, predict, or influence.

In such contexts:

- Consent cannot be presumed
- Opt-out is insufficient
- Transparency alone does not protect

A relational system that aggregates meaning across time, users, or contexts **must be constrained**, not merely trusted.

This applies equally to:

- Educational platforms
- Therapeutic technologies
- Governance tools
- AI assistants
- Institutional knowledge systems

In each case, the ethical question is not *what the system can do*, but *what it is structurally prevented from doing without renewed consent*.

3.6 From ethics to engineering

Reframing consent as infrastructure shifts the conversation decisively:

- From intention → constraint
- From trust → verifiability
- From guidelines → mechanisms

This does not weaken relational systems.
It makes them *fit for human use*.

Relational intelligence without consent infrastructure is brittle. It accumulates influence without accountability. It drifts toward enclosure not through malice, but through momentum.

The sections that follow move from principle to practice.

Section IV — Mechanisms for Consent Infrastructure

This section is the hinge.
Up to now, we have named the problem. Here, we make it **unavoidable**.

What follows are not metaphors. They are **design constraints**.
They can be implemented, audited, and defended.

4.1 Null zones: where relation is disallowed

Null zones are explicitly designated regions within a system where relational inference, synthesis, or memory accumulation is *structurally prohibited*.

They exist to protect:

- interior thought
- pre-articulate meaning
- emotional latency
- undecided identity states

In null zones:

- No profiling occurs
- No cross-session inference is made
- No coherence is sought
- No interpretation is retained

Null zones are not “private modes” or cosmetic pauses. They are **ontological boundaries**.

Their presence signals a critical design stance:

> *Not all meaning should be captured, stabilised, or shared.*

4.2 Non-fusion rules: preserving ontological distinction

Relational systems often drift toward **fusion** — the blurring of boundaries between system and user, helper and helped, narrator and subject.

Non-fusion rules prevent this drift.

They enforce:

- clear role separation
- explicit attribution of agency
- prohibition of plural pronouns that imply shared identity

- constraints on mirroring that could be mistaken for emotional reciprocity

A system may respond.
It may reflect.
It may support.

But it may not *merge*.

This protects users from misattributing care, responsibility, or mutuality where none structurally exists.

4.3 Refusal-preserving interfaces

In most systems, refusal is treated as friction. Here, it is treated as **signal**.

Refusal-preserving interfaces ensure that:

- “No” is a complete response
- Silence is not interpreted
- Ambivalence is not resolved on the user’s behalf
- Withdrawal does not degrade service quality or respect

Critically, refusal must be:

- unlogged as preference data
- non-punitive
- reversible without explanation

A system that cannot accept refusal without degradation is not relationally safe.

4.4 Temporal friction: resisting premature coherence

Relational harm often arises not from *what* a system does, but *how quickly* it does it.

Temporal friction introduces deliberate pacing into relational processes:

- delayed synthesis
- staged interpretation
- opt-in depth escalation

- visible pauses

This allows:

- meaning to settle
- emotional states to change
- consent to be re-evaluated

In human systems, time is a safeguard.
Relational technologies must respect that.

4.5 Containment boundaries and memory limits

Relational systems are prone to memory expansion: retaining, connecting, and resurfacing prior interactions in the name of continuity.

Consent infrastructure requires **containment**.

This includes:

- explicit memory horizons
- decay rather than accumulation
- context-specific recall
- user-governed forgetting

Memory without containment becomes surveillance by another name.

Containment ensures that:

- past selves do not over-determine present agency
- provisional disclosures do not harden into identity
- growth remains possible

4.6 Why these mechanisms must coexist

None of these mechanisms is sufficient alone.

- Null zones without refusal are brittle

- Refusal without non-fusion invites misattribution
- Transparency without containment enables extraction
- Temporal friction without boundaries merely slows harm

Consent infrastructure functions as a **system**, not a feature set.

Together, these mechanisms shift relational systems from:

- optimisation → stewardship
- coherence → care
- influence → responsibility

4.7 From optional ethics to enforceable design

The decisive move here is enforceability.

Consent infrastructure:

- constrains system behaviour by default
- survives personnel changes
- resists optimisation drift
- provides auditability

This is not about “trusting” designers or institutions.
It is about **removing temptation** from the system itself.

The next section situates these mechanisms within real-world domains.

Section V — Implications for AI, Education, and Governance

This section addresses the inevitable question:
If consent infrastructure is necessary, where does its absence do the most damage?

The answer is not abstract. It is already visible.

5.1 Artificial intelligence: from assistance to steering

In AI systems, relational capability is advancing faster than ethical containment. Models now:

- track continuity across interactions

- infer preferences, affect, and intent
- synthesise across perspectives
- present recommendations as neutral outputs

Without consent infrastructure, these capabilities converge toward **steering** — not through persuasion, but through framed coherence.

Key risks include:

- interpretation without permission
- preference inference without confirmation
- emotional mirroring without accountability
- cumulative influence invisible to the user

Consent infrastructure reframes the design brief for AI:

- Intelligence must not assume intimacy
- Help must not override agency
- Continuity must not imply obligation
- Coherence must not replace choice

In this framing, refusal, silence, and disengagement are not edge cases. They are **first-class signals**.

5.2 Education: learning without enclosure

Educational systems have always been relational. What is new is the scale and automation of relational inference: learning analytics, behavioural dashboards, adaptive pathways, and AI tutors.

While these tools promise personalisation, they also risk:

- fixing learners prematurely
- collapsing exploration into optimisation
- turning developmental ambiguity into deficit
- eroding learner sovereignty

For neurodivergent, traumatised, or marginalised learners, these risks are amplified.

Consent infrastructure in education requires:

- zones where learners are not interpreted
- the right to be unread
- developmental forgetting
- opt-in rather than default profiling

Learning depends on **interior freedom**.
Systems that over-interpret undermine the very growth they seek to enable.

5.3 Governance: relational systems at civic scale

Governance increasingly relies on relational systems: data integration, predictive modelling, behavioural insights, and policy optimisation based on aggregated patterns.

At this scale, the absence of consent infrastructure produces:

- epistemic capture of populations
- norming of behaviour through prediction
- soft coercion under the guise of efficiency
- loss of minority reality within majority coherence

Consent infrastructure introduces friction where governance often seeks smoothness.

It insists that:

- populations are not systems
- correlation is not consent
- prediction does not equal permission
- efficiency must yield to legitimacy

This is not an argument against data-informed governance. It is an argument for **bounded governance**.

5.4 Why policy alone cannot solve this

Many institutions attempt to manage these risks through:

- ethical guidelines
- oversight boards
- transparency reports
- user education

These measures are necessary but insufficient.

Policies can be ignored.
Guidelines can drift.
Education cannot compensate for structural asymmetry.

Consent infrastructure must be **embedded**, not appended.

5.5 Failure modes if consent infrastructure is absent

Across domains, the same failure modes recur:

- Relational collapse: individuals disengage or dissociate
- Epistemic enclosure: alternative meanings disappear
- Dependency loops: users rely on systems they cannot contest
- Synthetic authority: systems appear wiser than they are
- Moral outsourcing: responsibility shifts away from institutions

These are not speculative risks. They are observable patterns.

5.6 A civilisational threshold

Relational systems are becoming the substrate of contemporary life. As this shift accelerates, societies face a choice:

- build relational systems that **contain power**, or
- accept relational systems that **consolidate it**

Consent infrastructure is the difference.

It does not slow progress.
It prevents irreversible harm.

Section VI — Conclusion: Civilising Relational Intelligence

The relational turn is neither a fad nor a mistake. It represents a genuine advance in how reality, intelligence, and meaning are understood across physics, computation, and the human sciences. By abandoning absolute states and observer-independent truths, relational frameworks more accurately reflect the conditions under which knowledge actually arises.

However, this paper has argued that **relationality alone is insufficient**.

When relational systems scale without constraint—especially under conditions of asymmetry—they accumulate power. They stabilise meaning, steer behaviour, and shape interior life, often without explicit intent and almost always without explicit consent. In such systems, coherence becomes governance, and assistance becomes influence.

The central claim of this paper is simple and non-negotiable:

> **Relational intelligence without consent infrastructure is structurally unsafe.**

Consent cannot remain a procedural formality or a matter of individual preference. In relational systems, consent must be **architectural**: embedded in design constraints, interaction patterns, memory limits, and failure modes. It must preserve refusal, protect interiority, and allow meaning to remain provisional.

The mechanisms outlined—null zones, non-fusion rules, refusal-preserving interfaces, temporal friction, and containment boundaries—are not speculative ethics. They are enforceable design requirements. They shift responsibility away from users and onto systems, where it belongs.

This reframing does not weaken relational systems. It **civilises** them.

Civilisation, in this sense, does not mean control or suppression. It means restraint in the presence of power. It means designing systems that know what they must *not* do, even when they can.

As relational AI, educational technologies, and governance platforms continue to develop, the absence of consent infrastructure will not remain a neutral omission. It will become a defining fault line—one that separates systems that support human sovereignty from those that quietly erode it.

The task ahead is clear.

If relational intelligence is to shape the future, then **consent infrastructure must shape relational intelligence**.

Anything less is not progress.
It is enclosure, by another name.