Legal Codes as Semantic Space: Toward a Value-Ontological Infrastructure for Preventive Justice

J. Konstapel Leiden, 1-3-2026

With the help of Claude.

Abstract

Contemporary legal AI systems treat statutory codes as rule databases optimized for retrieval efficiency. This paper proposes a fundamentally different paradigm: legal codes as layered semantic spaces encoding decades of collective value negotiation. Drawing on legal ontology, Homotopy Type Theory (HTT), deliberative democracy theory, and collective intelligence frameworks, we present an architecture in which AI mediates value discovery rather than rule application. A Monitor subsystem aggregates individual legal reflections into collective pattern maps — rendering visible the pre-escalation layer of social conflict that current legal and political systems cannot perceive. We argue this constitutes a new category of instrument: **preventive justice infrastructure**, distinct from both legal tech and policy analytics. A pilot design for municipal deployment is outlined.

Keywords: legal ontology, Homotopy Type Theory, deliberative democracy, collective intelligence, preventive justice, semantic law, value discovery

1. Introduction

The dominant trajectory of legal technology optimizes for speed and scale: faster contract analysis, more efficient document retrieval, lower-cost jurisprudence search (Surden, 2014; Katz, 2013). This trajectory treats law as a formal system whose primary challenge is computational throughput.

We contest this framing at the foundational level. Legal codes did not originate as rule databases. They emerged as instruments for encoding collective judgments about values — what constitutes fair exchange, adequate protection, legitimate authority (Fuller, 1969; Dworkin, 1986). The formalization of law into retrievable rules is a historically recent and functionally partial representation of what legal codes actually contain.

This paper develops an alternative architecture grounded in three convergent theoretical traditions:

1. **Legal ontology** — the study of how juridical concepts are semantically structured and relationally embedded (Valente, 1995; Bench-Capon & Sartor, 2003)

2. **Homotopy Type Theory** — a mathematical framework that treats equality as structural equivalence rather than identity, enabling richer modeling of normative relationships (Univalent Foundations Program, 2013)

3. **Deliberative democracy and collective intelligence** — frameworks for understanding how distributed individual judgment aggregates into legitimate collective meaning (Habermas, 1996; Dryzek, 2000; Surowiecki, 2004)

The synthesis of these three traditions yields a practical architecture: a conversational AI system that surfaces the value layer beneath legal rules, combined with a Monitor subsystem that aggregates individual value-reflections into collective pattern maps. The result is an instrument that makes visible a previously imperceptible social layer — the zone of felt injustice that has not yet escalated into formal legal conflict.

2. Theoretical Foundations

2.1 Legal Ontology and Semantic Structure

Legal ontology examines the deep structure of juridical concepts — not what rules say, but how legal concepts relate to each other and to underlying principles (Valente, 1995). The LKIF (Legal Knowledge Interchange Format) Core ontology, developed within the ESTRELLA project, demonstrated that legal concepts function as nodes in semantic networks rather than as isolated propositional units (Hoekstra et al., 2007).

Bench-Capon and Sartor (2003) showed that legal reasoning is fundamentally argument-theoretic: legal conclusions emerge from structured dialogues between competing value frameworks, not from mechanical rule application. This has significant implications for AI design. A system that retrieves the applicable rule without surfacing the underlying value framework provides formally correct but contextually impoverished guidance.

McCarty's (1977) foundational work on TAXMAN demonstrated that even in the apparently rule-bound domain of tax law, edge cases require reasoning about the *purposes* rules serve — what values they protect — rather than their literal formulation. This observation, made nearly fifty years ago, has been insufficiently integrated into contemporary legal AI design.

2.2 Homotopy Type Theory as a Framework for Legal Equivalence

Standard first-order logic, which underlies most formal legal reasoning systems, treats equality as identity: two legal norms are either identical or they are not. This is inadequate for legal interpretation, where different statutory provisions may express the same underlying principle through different formulations, in different contexts, across different historical periods.

Homotopy Type Theory (HoTT), developed at the Institute for Advanced Study (Univalent Foundations Program, 2013), offers a richer framework. In HoTT, equality is treated as a *path* — a structural relationship between terms that preserves properties while allowing formal differences. The Univalence Axiom states that equivalent structures are interchangeable.

Applied to legal semantics, this means we can formally represent the relationship between, for example, Article 7:247 BW (Dutch Civil Code protection against unreasonable rent increases) and Article 6:248 BW (good faith in contractual relations) as a *homotopy path* through the underlying principle of *equality of bargaining position*. Both articles are formal expressions of the same deeper value, arrived at through different legislative routes.

This has concrete computational implications. A HTT-structured legal knowledge base can:

- Identify when apparently unrelated statutory provisions share a common value basis
- Model how legislative evolution represents path-deformation rather than value replacement
- Support analogical reasoning across legal domains through shared value-type membership

2.3 Deliberative Democracy and the Pre-Escalation Layer

Habermas's (1996) theory of deliberative democracy holds that legitimate legal norms emerge from rational communicative processes in which affected parties can voice their perspectives. A critical weakness of existing legal systems is that this communicative process effectively terminates once norms are codified. Citizens interact with law as a finished product rather than as an ongoing conversation.

Dryzek (2000) extends this by emphasizing *discursive representation* — the need for governance systems to represent not just organized interests but the diverse discourse-types through which citizens make sense of their social world. Current legal AI amplifies the efficiency of accessing finished legal products; it does not support discursive representation.

Our framework identifies a specific social layer that existing systems cannot perceive: the **pre-escalation zone** — the space of felt injustice, experienced value violation, and unresolved normative conflict that has not yet taken formal legal shape. This zone is sociologically significant. Research on legal consciousness (Ewick & Silbey, 1998) demonstrates that most experiences of injustice never reach formal legal institutions — not because they are trivial, but because the transaction costs (financial, cognitive, relational) of formal legal engagement are prohibitive.

The pre-escalation zone is where social cohesion is actually negotiated, damaged, and potentially repaired. Making this zone visible and actionable is the core contribution of the architecture we propose.

2.4 Collective Intelligence and Pattern Aggregation

Surowiecki (2004) established that aggregated independent judgments frequently outperform expert opinion, provided judgments are sufficiently independent and diverse. Lévy (1997) developed the concept of *collective intelligence* as a distributed cognitive capacity irreducible to individual contributors.

More recently, Woolley et al. (2010) demonstrated that groups exhibit a measurable *collective intelligence factor* — a general capacity for coordinated problem-solving that is distinct from the average or maximum intelligence of group members. This factor correlates strongly with social sensitivity and turn-taking in conversation.

These findings support the design of a Monitor subsystem that aggregates individual value-reflections. When multiple individuals, independently engaging with a legal domain, arrive at similar value framings, this convergence constitutes evidence about the normative landscape — evidence that is invisible to both individual legal practitioners and to aggregate statistical measures like opinion polls.

3. Architecture

3.1 Overview

The proposed system consists of three integrated components:

Component 1: The Semantic Legal Model Statutory codes are annotated as type-hierarchies with dependent relations:

```
Legal Domain
└── Code (e.g., Dutch Civil Code, Criminal Code)
    └── Article
        └── Legal Norm
            ├── Underlying Value (protection, equality, freedom, dignity)
            ├── Social Context (historical, demographic, relational)
            └── Application Domain (contract, property, tort, family)
```

HTT path-relationships are mapped between articles sharing value-type membership. This yields not a flat rule database but a *value topology* — a structured landscape in which legal provisions are positioned relative to their normative foundations.

Component 2: The Conversational Value Discovery Interface A conversational AI engages users through a structured three-stage dialogue:

Stage 1 — Perspective Elicitation: The system establishes the user's relational position (tenant, landlord, policymaker, researcher), experiential basis (personal experience vs. general inquiry), and the affective quality of their situation.

Stage 2 — Value Surfacing: Through reflective questioning, the system guides the user toward articulating the underlying value at stake, distinct from the specific rule they seek.

Stage 3 — Normative Mapping: The system maps the user's value-framing onto the semantic legal model, identifying relevant provisions as expressions of the underlying principle rather than as isolated rules.

The output is a structured reflection:

```
PERSPECTIVE: [relational position]

SOURCE: [experiential basis]

CORE FEELING: [affective characterization]

UNDERLYING PRINCIPLE: [value articulation]
```

```
RELEVANT PROVISIONS: [mapped statutory references]
```

Component 3: The Monitor Individual structured reflections are anonymized and aggregated by the Monitor subsystem. Pattern recognition operates across three dimensions:

- *Thematic clustering*: Which underlying values are most frequently activated?
- *Geographic distribution*: Where do specific value-tensions concentrate?
- *Temporal dynamics*: Are specific value-tensions intensifying or resolving over time?

Monitor outputs are designed for three distinct audiences: citizens (accessible summaries of shared experiences), legal professionals (pattern maps for anticipatory legal planning), and policymakers (early-warning signals for pre-escalation intervention).

3.2 Privacy Architecture

Given the sensitivity of pre-legal conflict data, privacy-by-design is architecturally foundational. Individual conversations are processed through a differential privacy pipeline before entering the Monitor. No individual reflection is stored in identifiable form. Geographic aggregation operates at minimum resolution of postal code district. Temporal data is retained for rolling 24-month windows only.

This design follows the requirements of GDPR Article 25 (data protection by design) and aligns with the Dutch Autoriteit Persoonsgegevens guidelines for AI systems processing sensitive personal data.

4. Pilot Design

4.1 Domain Selection

Housing conflict is selected as the pilot domain for three reasons: (1) it affects a demographically broad population including vulnerable groups; (2) existing legal protection (BW Articles 7:247-252) is technically adequate but practically inaccessible to many tenants; (3) the pre-escalation zone in housing is demonstrably large — research by the Dutch Huurcommissie indicates that a significant majority of potentially actionable rent disputes are never formally raised.

4.2 Municipal Partnership

A single municipal government serves as pilot partner. The municipality receives Monitor output as a policy intelligence service: quarterly pattern reports identifying emerging value-tensions in the local housing domain, mapped to geographic distribution and temporal trend. Individual conversations remain fully confidential.

4.3 Evaluation Criteria

Success is evaluated on three dimensions:

- **Accessibility**: Do users report improved understanding of the value basis of relevant legal provisions? (pre/post conversation survey)
- **Resolution rate**: What proportion of conversations result in user-reported clarity or resolution without formal legal escalation?
- **Policy utility**: Do municipal policy officers rate Monitor reports as containing actionable information not available through existing channels?

5. Discussion

5.1 Relationship to Existing Legal AI

This architecture is not in competition with efficiency-oriented legal AI. Document review, contract analysis, and jurisprudence retrieval serve legitimate functions that this system does not address. The claim is narrower: that a category of legal interaction — the pre-escalation zone of value conflict — is currently unserved by existing tools, and that this gap has significant social consequences.

5.2 The Contestability of Values

A legitimate challenge concerns the apparent neutrality of "value discovery." Legal values are not pre-political givens — they are contested, historically situated, and distributed unequally across social positions. The architecture must not present value-framing as neutral when it inevitably reflects choices about which values are represented in the semantic model.

We address this through two design commitments: (1) the semantic legal model is built through participatory annotation involving legal scholars, community organizations, and citizen panels; (2) the conversational interface explicitly presents multiple competing value framings for any given conflict domain, rather than converging on a single "correct" underlying principle.

5.3 The Governance of the Monitor

The Monitor itself requires a governance structure that is itself value-consistent. We propose a multi-stakeholder oversight body including citizen representatives, legal professionals, and municipal partners, with quarterly public reporting on Monitor operation and output. This is consistent with the participatory governance models advocated by Ostrom (1990) for commons management — the Monitor output is, in effect, a commons: collectively produced knowledge about shared normative experience.

6. Conclusion

Legal codes are not rule databases. They are encoded records of collective value negotiation — living semantic structures that contain more than formal legal systems have learned to retrieve. By applying legal ontology, Homotopy Type Theory, and collective intelligence frameworks to statutory interpretation, we can build AI systems that help citizens discover the value landscape

beneath legal rules, and that aggregate individual value-reflections into actionable collective intelligence.

The practical result is a new category of instrument: preventive justice infrastructure. It does not replace formal legal institutions. It makes visible and actionable the pre-escalation zone where social cohesion is actually formed and damaged — the layer that current legal and political systems are structurally unable to perceive.

This is law returned to its relational origin: not as command, but as instrument for building shared understanding of what we owe each other.

Based on https://constable.blog/2025/06/01/wetboeken-als-betekenisruimte-een-nieuwe-juridische-infrastructuur/

References

Bench-Capon, T. J. M., & Sartor, G. (2003). A model of legal reasoning with cases incorporating theories and values. *Artificial Intelligence*, 150(1–2), 97–143.

Dryzek, J. S. (2000). *Deliberative Democracy and Beyond: Liberals, Critics, Contestations*. Oxford University Press.

Dworkin, R. (1986). *Law's Empire*. Harvard University Press.

Ewick, P., & Silbey, S. S. (1998). *The Common Place of Law: Stories from Everyday Life*. University of Chicago Press.

Fuller, L. L. (1969). *The Morality of Law* (rev. ed.). Yale University Press.

Habermas, J. (1996). *Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy*. MIT Press.

Hoekstra, R., Breuker, J., Di Bello, M., & Boer, A. (2007). The LKIF Core ontology of basic legal concepts. *Proceedings of the Workshop on Legal Ontologies and Artificial Intelligence Techniques (LOAIT)*, 43–63.

Katz, D. M. (2013). Quantitative legal prediction — or — how I learned to stop worrying and start preparing for the data-driven future of the legal services industry. *Emory Law Journal*, 62(4), 909–966.

Lévy, P. (1997). *Collective Intelligence: Mankind's Emerging World in Cyberspace*. Plenum Trade.

McCarty, L. T. (1977). Reflections on TAXMAN: An experiment in artificial intelligence and legal reasoning. *Harvard Law Review*, 90(5), 837–893.

Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press.

Surden, H. (2014). Machine learning and law. *Washington Law Review*, 89(1), 87–115.

Surowiecki, J. (2004). *The Wisdom of Crowds*. Doubleday.

Univalent Foundations Program. (2013). *Homotopy Type Theory: Univalent Foundations of Mathematics*. Institute for Advanced Study. https://homotopytypetheory.org/book/

Valente, A. (1995). *Legal Knowledge Engineering: A Modelling Approach*. IOS Press.

Woolley, A. W., Chabris, C. F., Pentland, A., Hashmi, N., & Malone, T. W. (2010). Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004), 686–688.