The Philosophical Foundations of the Village Project

A Framework for Digital Sovereignty and Pluralist AI Governance

My Digital Sovereignty Limited February 2026

Précis

The Village project represents an attempt to operationalise rigorous philosophical principles into the architecture of digital community platforms and artificial intelligence systems. Rather than treating philosophy as decorative rhetoric or post-hoc justification, this initiative embeds the insights of Isaiah Berlin on value pluralism, Christopher Alexander on structural integrity, and indigenous data sovereignty frameworks into the technical substrate of how communities organise and how AI systems behave. The result is a three-layer constitutional architecture where certain protections are immutable, a governance framework where AI cannot automate value decisions, and a training methodology where philosophical grounding precedes capability development. This document presents the intellectual foundations, the contemporary challenges that motivate this work, and the concrete plans for implementation.

I. Philosophical Foundations

Isaiah Berlin and the Plurality of Values

The foundational philosophical commitment of this project derives from Isaiah Berlin's work on value pluralism, articulated most fully in his 1969 collection *Four Essays on Liberty*. Berlin argued that genuine human goods—liberty, equality, justice, mercy, efficiency, solidarity—are fundamentally plural and often incommensurable. There exists no common currency, no universal scale, no meta-value by which to rank them definitively. As the Stanford Encyclopedia of Philosophy summarises: "Value pluralism is the view that there are many different values which may be equally correct and fundamental, yet in conflict with each other."

The implications of this position are profound and frequently misunderstood. Berlin was not advocating relativism—the claim that all value systems are equally valid or that values are merely subjective preferences. Rather, he maintained that genuine, objective values can conflict irreducibly. When liberty clashes with equality, or mercy with justice, we face a tragic choice where something of real value must be sacrificed. This is not a failure of analysis that better reasoning could resolve; it is intrinsic to the human condition.

For artificial intelligence governance, Berlin's framework yields a critical principle: **AI systems must never assume a default value hierarchy**. When an AI optimises for efficiency, it implicitly subordinates other values—thoroughness, care, reflection—without explicit human authorisation for that trade-off. The Tractatus framework that governs our AI systems therefore recognises six irreducibly different moral frameworks:

1. **Deontological** — duties, rights, and inviolable principles
2. **Consequentialist** — outcomes and aggregate welfare
3. **Virtue Ethics** — character and human flourishing
4. **Care Ethics** — relationships and attention to vulnerability
5. **Communitarian** — tradition and shared goods
6. **Indigenous Relational** — interconnection and reciprocity

When these frameworks conflict—as they inevitably do—the system does not resolve the tension algorithmically. It surfaces the conflict to human decision-makers with transparency about what each framework would recommend and what would be lost by each choice. This is what we term "pluralistic deliberation": the structured acknowledgment that reasonable people holding different but legitimate values will reach different conclusions, and that this diversity is a feature of healthy communities rather than a defect to be engineered away.

Berlin's distinction between negative liberty (freedom from interference) and positive liberty (freedom to achieve self-mastery) also informs our approach. He warned that positive liberty, despite its emancipatory rhetoric, carries dangers when wielded by those who claim to know others' "true" interests better than they know themselves. This warning resonates with particular force in the age of algorithmic curation, where systems optimise for "engagement" on the assumption that what captures attention represents genuine preference. Our commitment to negative liberty manifests as a refusal to optimise: we do not curate feeds, do not recommend content, do not nudge behaviour toward platform-defined outcomes.

Christopher Alexander and Structural Integrity

Christopher Alexander's work on pattern languages and architectural theory provides the methodological framework for translating philosophical principles into technical architecture. Alexander, whose influence extends from physical architecture to software design, argued that living systems exhibit particular structural properties that cannot be achieved through top-down planning but emerge from careful attention to how parts relate to wholes.

Five of Alexander's principles have been codified into the Tractatus governance framework:

Deep Interlock (Rule inst_090): Components must coordinate through mutual validation rather than operating in isolation. In our implementation, no single service can approve an action alone. The BoundaryEnforcer validates that an operation respects value boundaries; the MetacognitiveVerifier confirms that reasoning quality meets standards; the CrossReferenceValidator ensures consistency with prior decisions. This interlocking prevents any single point of bypass.

Structure-Preserving Transformation (Rule inst_091): Changes to a system must preserve its essential structure. When our AI systems learn from new data or adapt to new contexts, they must not break the interpretability of audit logs, invalidate prior governance decisions, or undermine instruction precedents. Evolution occurs within structural constraints.

Gradients Rather Than Boundaries (Rule inst_092): Living systems operate on gradients of intensity rather than binary switches. Our governance therefore recognises levels—NORMAL, ELEVATED, HIGH, CRITICAL, DANGEROUS—with different protocols at each level. This prevents the brittleness of systems that have only "allowed" and "forbidden" states.

Living Process (Rule inst_093): The framework evolves from real operational experience rather than predetermined specification. Governance rules emerge from actual failures, documented incidents, and observed edge cases. This represents a commitment to institutional learning over institutional rigidity.

Not-Separateness (Rule inst_094): Perhaps the most critical principle. Governance must be embedded in architecture, not bolted on as an afterthought. If an AI system can execute without governance validation, the governance is separate—and will inevitably be bypassed under pressure. In our training loops, the BoundaryEnforcer validates each batch before training proceeds; governance is inside the loop, not a filter applied afterward.

This last principle addresses a persistent failure mode in AI safety: the tendency to develop capable systems first and add safety measures later. Such approaches treat safety as a constraint on capability rather than as constitutive of what the system fundamentally is. Alexander's work suggests that this separation is not merely strategically unwise but architecturally unsound—systems with bolted-on governance lack the structural integrity of systems where governance and function are unified.

Indigenous Data Sovereignty

The third structural foundation draws from indigenous frameworks that long predate digital technology but speak directly to questions of data governance, collective rights, and the relationship between information and community.

Te Mana Raraunga, the Māori Data Sovereignty Network, articulates principles that challenge Western individualist assumptions about data ownership:

- **Rangatiratanga** — authority and control residing with the community, not extracted to external platforms
- **Whakapapa** — data understood through relationships and lineage, not as isolated records
- **Whanaungatanga** — kin connections that create obligations and responsibilities
- **Kotahitanga** — collective vision that transcends individual preferences

- **Manaakitanga** — reciprocity and care as foundational to legitimate data use
- **Kaitiakitanga** — guardianship rather than ownership

These principles are complemented by the **CARE Principles** developed by the Global Indigenous Data Alliance: Collective Benefit, Authority to Control, Responsibility, and Ethics. The **OCAP Framework** from First Nations Canada (Ownership, Control, Access, Possession) provides parallel guidance.

In our implementation, these frameworks manifest as structural guarantees: individual members cannot override collective governance decisions about shared data; communities retain sovereignty over how their collective narratives are processed and presented; no data flows to external systems without explicit collective consent through established governance procedures.

This represents a substantive philosophical commitment, not mere acknowledgment. We recognise that indigenous communities have been theorising and practising data sovereignty for generations—long before "data" became a technology industry term—and that their frameworks offer sophisticated responses to problems that Western philosophy is only beginning to recognise.

The Tractatus: Wittgenstein and the Limits of the Sayable

The governance framework's name—Tractatus—deliberately invokes Ludwig Wittgenstein's *Tractatus Logico-Philosophicus* (1921), one of the most influential philosophical works of the twentieth century. Wittgenstein drew a fundamental distinction between what can be said (expressed in propositions, subject to logic) and what can only be shown (values, ethics, the mystical).

This distinction maps directly onto AI governance. Certain decisions can be systematised and delegated to autonomous agents: technical optimisations, pattern matching, data retrieval, syntactic transformations. These belong to the realm of the "sayable"—they can be specified, measured, verified.

Other decisions—those involving values, ethics, cultural context, and human agency—fundamentally cannot and must not be automated. They belong to what Wittgenstein called the "unsayable": not because they are irrational or arbitrary, but because they exceed the capacity of formal systems to capture. The Tractatus framework enforces this boundary architecturally, ensuring that AI systems operate only in the domain of the sayable while preserving human authority over everything beyond it.

II. The Village: Architecture as Philosophy Made Concrete

Three-Layer Constitutional Architecture

The Village platform implements a three-layer constitutional architecture that translates philosophical principles into enforceable structure:

Layer 1: Universal Platform Principles (Immutable)

This layer comprises principles that cannot be overridden by any tenant, administrator, or user. It is hardcoded into the system architecture and enforced by the Tractatus framework. Examples include:

- Tenant data isolation (no cross-community data sharing without explicit bilateral agreement)
- The right to leave (members can depart any Village at any time, taking their data, with no penalties or lock-in)
- Consent requirements for data use (moderator approval never substitutes for individual consent)
- No imposed value hierarchy (the platform does not rank communities' values against each other)
- Governance embedded in architecture (per Alexander's not-separateness principle)

These principles are not policies that could be changed through governance processes; they are structural constraints that make certain violations architecturally impossible.

Layer 2: Tenant Constitutional Principles (Customisable)

Each Village (tenant community) defines its own constitution within the bounds established by Layer 1. This layer includes:

- Tone and communication style (formal or casual, technical or accessible)

- Content moderation norms (how the community handles controversial topics)
- Decision-making models (consensus, majority vote, delegated authority)
- Privacy and transparency settings (logging levels, attribution defaults)
- Cultural protocols (language, traditions, regional considerations)
- AI assistance boundaries (when AI helps versus when humans must decide)

This layer embodies Berlin's value pluralism in practice: different communities have legitimately different values, and the platform accommodates this diversity rather than imposing homogeneity.

Layer 3: Member Personal Preferences (Individual)

Individual members configure their own preferences within the bounds of their community's constitution: notification frequency, language preferences, AI assistance levels, privacy defaults for their own content. Layer 3 preferences yield to Layer 2 community standards, which yield to Layer 1 universal principles.

This layered architecture ensures that fundamental rights are protected while genuine pluralism flourishes. A family community can adopt different norms than a professional organisation; a religious community can establish different communication protocols than a secular one. But no community can violate the foundational protections that apply to all.

The Company Constitution

My Digital Sovereignty Limited, the entity developing the Village platform, operates under a published constitution organised around six principles:

1. **Sovereignty First** — Full data ownership means complete control, not data minimisation. Members can export all their data at any time. Deletion is genuine deletion, from production systems, backups, and AI training data.

2. **Privacy as Default** — Privacy is a human right, not a premium feature. The platform has no "public" mode; every interaction occurs within authenticated, tenant-scoped contexts. No tracking pixels, no cross-site analytics, no surveillance infrastructure.

3. **Pluralism Over Homogeneity** — The platform does not impose values. Different communities govern themselves according to their own principles, within the bounds of universal protections.

4. **Transparency and Accountability** — The constitution itself is public. AI systems disclose which models were used, what data was accessed, how decisions were made. Incident reports are published with full context.

5. **Safety Without Surveillance** — Security and privacy are not opposing forces to be "balanced." Community-led moderation replaces algorithmic content curation. Threat protection occurs at system boundaries without behavioural surveillance of members.

6. **Sustainable Business Model** — Fair pricing based on cost plus reasonable margin. No bait-and-switch, no VC-driven growth imperatives, no exit-focused strategy that would compromise member interests.

Federation with Integrity

Villages can federate with other Villages—sharing content, enabling cross-community interaction—through bilateral agreements that make constitutional differences explicit. Federation is not automatic or frictionless; it requires:

- Mutual consent from both communities' governance structures
- Explicit identification of where constitutions differ
- Agreement on how conflicts will be resolved (origin rules, destination rules, stricter-applies, mutual recognition, or custom terms)
- Three-tier dispute resolution processes with defined timelines

This approach treats federation as a relationship between sovereign entities rather than a technical integration to be optimised for growth. It accepts the friction of genuine pluralism rather than erasing difference in the name of seamless user experience.

III. The Contemporary Predicament

Berlin's Warning Realised

Isaiah Berlin warned that positive liberty—the freedom to achieve one's "true" self—carries dangers when claimed by those who presume to know others' true interests better than they know themselves. This warning has found its most complete realisation in the attention economy.

Contemporary platforms optimise for "engagement" on the theory that what captures attention represents revealed preference. But this conflates compulsion with choice. The neurological mechanisms exploited by algorithmic feeds—variable reward schedules, social validation loops, outrage amplification—do not represent the preferences of an autonomous agent making considered judgments. They represent the vulnerabilities of a biological system being systematically exploited.

The companies operating these systems genuinely believe they are serving users by maximising engagement. This is not cynical rationalisation but sincere conviction—which makes it far more dangerous than mere greed. Berlin understood that the most destructive forms of unfreedom come wrapped in the language of liberation, promising to help people achieve what they "really" want by overriding their expressed preferences.

The Concentration of AI Capability

The development of large language models has concentrated unprecedented capability in a small number of organisations. These systems can generate human-quality text, analyse complex documents, write functional code, and engage in extended reasoning. They are being integrated into every domain of human activity: education, healthcare, law, creative work, personal relationships.

This concentration raises questions that existing governance frameworks cannot answer:

- When an AI system trained on humanity's collective output becomes a mediating layer in human communication, who governs the values embedded in that mediation?
- When children grow up with AI tutors shaped by particular philosophical assumptions, how do those assumptions propagate across generations?
- When AI systems increasingly handle tasks that previously required human judgment, what happens to the human capacity for judgment itself?

These are not technical problems amenable to technical solutions. They are philosophical problems about the kind of beings we want to be and the kind of society we want to inhabit.

The Erosion of Epistemic Autonomy

Perhaps most concerning is the erosion of what might be called epistemic autonomy: the capacity to form beliefs through one's own reasoning rather than accepting conclusions delivered by systems one does not understand. When an AI system produces an answer, most users cannot evaluate the reasoning that produced it. They must trust or distrust based on track record and reputation—heuristics that are easily gamed.

This represents a qualitative shift in the human relationship to knowledge. Previous technologies—books, libraries, search engines—augmented human capacity to find and evaluate information. Current AI systems increasingly substitute for that capacity, delivering conclusions rather than evidence, answers rather than arguments.

The long-term consequence may be a population that has outsourced not just information retrieval but judgment itself—capable of asking questions but not of evaluating answers, dependent on systems whose operations they cannot inspect and whose values they cannot interrogate.

IV. A Philosophy-First Approach to AI Development

The Home AI Concept

In response to these challenges, we are developing what we term "Home AI"—a small, locally-trained language model (SLL) that operates under community governance on user-controlled hardware. The distinguishing characteristics are:

Sovereignty: The model runs on hardware owned or controlled by the community. Training data remains local. No information flows to external systems without explicit consent through established governance procedures.

Transparency: Communities can inspect what the model knows about them, how it was trained, and why it produces particular outputs. AI memory is not a black box but an auditable record subject to community governance.

Philosophical Grounding: The model is trained with explicit attention to philosophical foundations. Rather than optimising purely for capability and adding safety measures afterward, we embed philosophical constraints from the earliest stages of development.

Community Governance: Each community configures how its AI assistant behaves according to its own constitutional principles. A community that values directness configures for directness; one that values gentleness configures for gentleness. The platform provides the infrastructure; communities provide the values.

Stanford Encyclopedia of Philosophy as Authoritative Reference

For philosophical concepts, we have established the Stanford Encyclopedia of Philosophy (SEP) as the single authoritative reference. This decision reflects both the quality of SEP's scholarship and a commitment to intellectual rigour that resists the temptation to treat complex philosophical positions as resources to be mined for convenient quotations.

When the training process encounters philosophical terms, it cross-references against SEP entries. When multiple interpretations exist, SEP's analysis of the debate takes precedence. When users ask philosophical questions, responses are grounded in SEP definitions rather than generated from statistical patterns in training data.

This is not merely a quality control measure but a substantive philosophical commitment: that AI systems engaging with philosophical concepts should do so with the same rigour expected of human scholars, acknowledging complexity rather than flattening it, representing debates rather than resolving them prematurely.

Wisdom Traditions as Layer 3 Customisation

Beyond the structural philosophical foundations (Layer 1) and community constitutional principles (Layer 2), we provide a system of adoptable wisdom traditions that influence how AI assistance is framed and delivered (Layer 3). It is essential to understand what this layer does and does not do.

What Layer 3 affects: Communication style, framing, language choices, pacing suggestions. Adopted traditions shape *how* Home AI communicates with you.

What Layer 3 does not affect: Content decisions, data access, governance enforcement. Adopted traditions do not control *what* the system is permitted to do. They are tendencies, not rules, and can always be overridden in any specific situation.

Thirteen traditions have been documented with scholarly validation against the Stanford Encyclopedia of Philosophy, including:

- **Simone Weil** — her concept of attention as receptive engagement with suffering influences options like "take your time" for grief content and resistance to compressing loss into summaries
- **Stoicism** — emphasising the distinction between what is and is not within one's control
- **Care Ethics** — emphasising relationships, vulnerability, and contextual judgment
- **Confucian Ethics** — emphasising relational roles and social harmony
- **Buddhist Ethics** — emphasising impermanence, interdependence, and the cessation of suffering
- **Ubuntu** — emphasising communal identity and mutual obligation ("our family's story" rather than "your story")
- **Jewish Ethics** — emphasising *tikkun* (repair) and *tzedakah* (righteous giving)
- **Islamic Ethics** — emphasising mercy, justice, and submission to transcendent principles

- **Indigenous/Māori frameworks** — emphasising *whakapapa* (genealogical connection) and kinship obligations

Communities and individuals may adopt traditions that resonate with their values. These adoptions influence how AI assistance is framed—what considerations are foregrounded, what language is used, what options are offered—without overriding the structural protections established in Layer 1 or the constitutional rules established in Layer 2.

When traditions suggest different approaches (as they sometimes do—Stoic equanimity may tension with Weil's attention to affliction), the system surfaces the tension rather than resolving it algorithmically, inviting human reflection on what the situation calls for. This is Berlin's value pluralism in practice: legitimate values genuinely conflict, and the system does not presume to resolve that conflict for you.

Governance Embedded in Training

Following Alexander's principle of not-separateness, we embed governance into the training process itself rather than applying it as a post-hoc filter. The training loop includes:

- **Boundary enforcement** at each batch, validating that training data respects constitutional principles
- **Consent verification** confirming that data used for training was contributed with appropriate permissions
- **Tenant isolation** preventing any leakage between communities' training data
- **Audit logging** recording what data influenced each training run

This approach accepts a performance overhead—approximately 5% in our testing—in exchange for genuine guarantees. The alternative—training unconstrained and filtering afterward—produces systems where governance is always potentially bypassable, always in tension with capability, always at risk of erosion under competitive pressure.

V. Toward Pluralist Digital Sovereignty

The Village project represents one possible approach to challenges that admit no definitive solutions. We do not claim to have resolved the tensions between capability and safety, between individual autonomy and collective governance, between technological efficiency and human values. These tensions are, as Berlin understood, genuinely tragic: navigating them requires judgment, compromise, and the acceptance that something of value will always be lost.

What we offer instead is a framework that takes these tensions seriously—that refuses to pretend they can be optimised away, that embeds philosophical reflection into technical architecture, that preserves human agency over decisions that should not be automated.

The test of this approach will not be whether it achieves some theoretical optimum but whether it enables communities to flourish according to their own understanding of flourishing—whether it protects the vulnerable, preserves diversity, maintains transparency, and creates conditions under which genuine human choice remains possible.

In an age of increasing AI capability and decreasing human oversight, this may be the most important work there is.

References

Berlin, I. (1969). *Four Essays on Liberty*. Oxford University Press.

Alexander, C. (1977). *A Pattern Language: Towns, Buildings, Construction*. Oxford University Press.

Alexander, C. (1979). *The Timeless Way of Building*. Oxford University Press.

Weil, S. (1951). *Waiting for God*. G.P. Putnam's Sons.

Wittgenstein, L. (1921). *Tractatus Logico-Philosophicus*.

Stanford Encyclopedia of Philosophy. "Value Pluralism." https://plato.stanford.edu/entries/value-pluralism/

Stanford Encyclopedia of Philosophy. "Simone Weil." https://plato.stanford.edu/entries/simone-weil/

Te Mana Raraunga. "Māori Data Sovereignty Network Principles." https://www.temanararaunga.maori.nz/

Global Indigenous Data Alliance. "CARE Principles for Indigenous Data Governance." https://www.gida-global.org/care

---

This document represents the intellectual foundations of work in progress. We welcome engagement from those who share our concerns and those who challenge our assumptions. The problems we address are too important for any single approach to claim sufficiency.

Contact: john@mysovereignty.digital

Addendum: Reference Materials and Resources

Tractatus AI Safety Framework

https://agenticgovernance.digital

The Tractatus framework's dedicated site provides technical documentation, research papers, and implementation guides for AI governance.

Resource	URL	Description
Research Paper (Academic)	/architectural-alignment.html	Peer-oriented presentation of the architectural alignment thesis with full scholarly apparatus
Research Paper (Community)	/architectural-alignment-community.html	Accessible version for general audiences without assumed technical background
Research Paper (Policymakers)	/architectural-alignment-policymakers.html	Governance-focused presentation emphasising regulatory and institutional implications
System Architecture	/architecture.html	Technical documentation of the six core services and their interactions
Village Case Study	/village-case-study.html	Detailed examination of how the Tractatus framework operates in the Village platform
Documentation	/docs.html	Complete technical documentation for implementers
Values & Principles	/about/values.html	The philosophical commitments underlying the framework's design
FAQ	/faq.html	Common questions about approach, limitations, and scope

Role-Based Entry Points:

Audience	URL	Description
Researchers	/researcher.html	Resources for those evaluating the framework's theoretical foundations
Implementers	/implementer.html	Technical guides for integrating Tractatus into existing systems
Leaders	/leader.html	Strategic overview for organisational decision-makers

Downloadable Papers (PDF):

- Academic Edition
- Community Edition
- Policymakers Edition

Village Platform

https://mysovereignty.digital

The Village platform implements the philosophical principles described in this document in a working community platform.

Foundational Documents:

Resource	URL	Description
About	/about.html	Mission statement and the case for privacy-first community platforms
Values	/values.html	The six core principles guiding platform development and operation
Constitution	/constitution.html	The foundational operational document governing My Digital Sovereignty Limited
Philosophy	/philosophy.html	Overview of core principles: Human Agency, Data Sovereignty, Community First, Radical Transparency
Platform Values	/platform-values.html	How values translate into concrete platform policies and features
Pricing Principles	/pricing-principles.html	The ethical framework for fair, transparent pricing without extractive practices

Architecture and Governance:

Resource	URL	Description
Data Sovereignty	/data-sovereignty.html	Technical and philosophical approach to genuine data ownership
Federation	/federation.html	How Villages form bilateral relationships while preserving constitutional autonomy
Village Constitution	/village-constitution.html	The three-layer constitutional architecture explained
Privacy	/privacy.html	Privacy policy reflecting the "privacy as default" principle
AI Transparency	/ai-transparency.html	How AI systems disclose their operations and respect member autonomy
Home AI	/home-ai.html	Introduction to sovereign, locally-trained AI assistance

Product Information:

Resource	URL	Description
Features	/features.html	Complete feature inventory across all platform capabilities
Village Community	/village-community.html	Platform configuration for community organisations and interest groups

Village Family	/village-family.html	Platform configuration for family heritage preservation and connection
Village Business	/village-business.html	Platform configuration for values-aligned business communities
Ecosystem	/ecosystem.html	Directory of specialised Villages and federation opportunities
Passport	/passport.html	Free tool for assessing technology migration readiness

Multilingual Resources:

The platform supports English, German, French, and Dutch, with te reo Māori in development:

- About: EN | DE | FR | NL
- Values: EN | DE | FR | NL
- Constitution: EN | DE | FR | NL
- Features: EN | DE | FR | NL

Scholarly References (External)

Primary Philosophical Sources:

Reference	URL	Description
Value Pluralism (SEP)	plato.stanford.edu/entries/value-pluralism	Stanford Encyclopedia entry on Berlin's value pluralism and its implications
Isaiah Berlin (SEP)	plato.stanford.edu/entries/berlin	Comprehensive overview of Berlin's philosophical contributions
Simone Weil (SEP)	plato.stanford.edu/entries/simone-weil	Scholarly treatment of Weil's philosophy of attention and affliction
Wittgenstein (SEP)	plato.stanford.edu/entries/wittgenstein	Background on the Tractatus and the sayable/unsayable distinction

Indigenous Data Sovereignty:

Reference	URL	Description
Te Mana Raraunga	temanararaunga.maori.nz	Māori Data Sovereignty Network's principles and resources
CARE Principles	gida-global.org/care	Global Indigenous Data Alliance's framework for indigenous data governance
OCAP Principles	fnigc.ca/ocap-training	First Nations Information Governance Centre on Ownership, Control, Access, Possession

Pattern Language and Architectural Theory:

Reference	URL	Description
Christopher Alexander (Wikipedia)	en.wikipedia.org/wiki/Christopher_Alexander	Biography and overview of Alexander's contributions to architecture and design theory

A Pattern Language (Wikipedia)	en.wikipedia.org/wiki/A_Pattern_Language	Overview of Alexander's architectural methodology and its software applications
Nature of Order	natureoforder.com	Christopher Alexander's four-volume work on living structure

---

Document Version

Version: 1.0 **Date:** February 2026 **Authors:** My Digital Sovereignty Limited **Classification:** Public **Citation:** Stroh, J. (2026). "The Philosophical Foundations of the Village Project: A Framework for Digital Sovereignty and Pluralist AI Governance." My Digital Sovereignty Limited.