

WhatIfWe

The Technology Landscape

Collective Intelligence: What Exists, What It Cannot Do, and Where WhatIfWe Sits

Version 2.0 · May 2026

This document maps the current state of the collective intelligence technology ecosystem in depth: the most significant existing initiatives, what each one can and cannot do, the structural ceiling they collectively reach, and the specific gap that WhatIfWe is designed to address. Version 2.0 substantially expands the landscape coverage to include major initiatives that emerged or matured in 2025–2026, deepens the analysis of existing projects, and sharpens the account of WhatIfWe's specific position within it. It is designed to be read alongside *Building the Lab* and can be provided to funders, technical partners, and potential guests from the technology track.

Contents

1. The Central Question
2. Computational Democracy: Polis, vTaiwan, the Habermas Machine, and Talk to the City
3. AI-Enhanced Collective Intelligence Research: AI4CI and the Collective Intelligence Project
4. The Global Scaling Wave: Engaged California, Global Dialogues, and the Inclusive AI Coalition
5. Warm Data and Transcontextual Knowledge: the Bateson Institute
6. Plurality and Collaborative Diversity: Weyl and Tang
7. Contemplative AI and Abundant Intelligences
8. The Epistemic Monoculture Problem: New Research (2025–2026)
9. Multi-Agent Systems and Epistemic Diversity Research
10. Commercial Knowledge Synthesis Platforms
11. What the Map Reveals: the Shared Ceiling
12. WhatIfWe's Specific Position

The Central Question

The technology world has been working for a decade on variations of a single problem: how do you get groups of people to think better together? The approaches vary enormously — from AI-assisted deliberation tools processing hundreds of thousands of participants across dozens of countries, to intimate contemplative frameworks for small groups, to architectural interventions in AI systems themselves, to Indigenous-led programmes rebuilding AI's epistemological foundations from the ground up. What is remarkable is how consistently these approaches, pursued to their frontier, identify the same gap.

The gap they are each naming is this: the difference between intelligence that operates within a single way of knowing — however sophisticated — and the intelligence that can only emerge from sustained genuine encounter between genuinely different ways of knowing. No existing system has crossed it.

Since Version 1.0 of this document, the pace of development in this field has accelerated significantly. The deployment of AI-assisted deliberation has moved from experimental pilots to institutional infrastructure: a major US state has launched a permanent AI-powered deliberative democracy platform; a global coalition co-initiated at the 2025 Paris AI Action Summit has begun a structured programme to engage 10,000 citizens from over 100 countries; and the research on epistemic monoculture — the structural problem that WhatIfWe is designed to address — has moved from a scattered set of findings into a coherent and urgent research agenda. None of these developments have closed the gap WhatIfWe is working at. Many of them have made it more visible.

This document maps the most significant initiatives in the landscape, assesses what each can and cannot do with precision, and locates WhatIfWe within this map. The purpose is not to position WhatIfWe as superior to these initiatives — many of them are doing important, complementary work — but to identify with precision the specific problem WhatIfWe is attempting to solve that none of them are currently equipped to address.

Computational Democracy

Polis, vTaiwan, the Habermas Machine, and Talk to the City

The vTaiwan process, developed under Taiwan's Digital Minister Audrey Tang beginning in 2014, remains the most sustained and documented national-scale implementation of AI-assisted collective intelligence. Its core mechanism — the Polis platform, developed by the Computational Democracy Project in Seattle — uses machine learning to cluster opinions across large populations, surfacing latent consensus that polarised discourse makes invisible. Citizens submit short statements and vote on others; Polis maps agreement and disagreement into a low-dimensional visualisation that reveals consensus statements (positions most groups agree on) and divisive ones (positions that sharply divide the population). By 2025, the process had engaged over 200,000 participants and contributed to 26 pieces of legislation.

The vTaiwan ecosystem has continued to evolve. The 'Join' platform, also under Taiwan's Digital Affairs Ministry, has significantly expanded the reach of participatory governance beyond the more technically demanding vTaiwan format, engaging a broader demographic including older users on issues ranging from drunk driving to child abuse policy. In 2023, Taiwan launched 'Alignment Assemblies' — a collaboration between the Ministry of Digital Affairs and the US-based Collective Intelligence Project — specifically designed to give citizens input into the governance of AI systems. The vTaiwan community's AI governance work, presented to Taiwan's National Human Rights Commission in March 2025, has contributed directly to the draft AI Basic Act. These developments represent the maturation of computational democracy from a civic experiment into state infrastructure — and a significant expansion of its scope into AI governance itself.

The Habermas Machine

In 2024, Google DeepMind's 'Habermas Machine' extended the vTaiwan/Polis approach significantly. Tested across 5,734 UK participants divided into small groups of five, the LLM-based system iteratively refined group statements through a 'caucus mediation' approach — generating candidate consensus statements, soliciting participant critiques, and revising until convergence. Independent evaluators rated the AI-mediated group statements as clearer and less biased than those generated by human mediators. The system also maintained minority viewpoint representation more reliably than human-led deliberation.

The Habermas Machine has attracted substantial critical attention from political theorists. A 2025 paper in *Ethics and Information Technology* (Volpe) and an AAAI/ACM AIES paper (Palomo Hernández) both argue that the system embeds a narrow conception of deliberative democracy: one that overemphasises agreement, privatises participation (each participant deliberates individually with the AI, not with each other), and relegates humans to the role of opinion generators rather than genuine deliberators. These are not incidental implementation flaws — they reflect the deep structural assumption that deliberation is about reaching consensus within a shared framework of what constitutes reasonable argument. When that shared framework is what differs, the architecture cannot proceed.

Talk to the City

Talk to the City (T3C), developed by the AI Objectives Institute, represents a significant complementary tool in the computational democracy ecosystem. Where Polis clusters opinions

statistically, T3C uses LLMs to generate structured thematic maps from unstructured input — survey responses, interview transcripts, social media discussions, meeting notes — preserving the nuance and specific language of participants' contributions. Every theme it identifies is grounded in direct participant quotes, making the analysis auditable in ways that statistical clustering cannot be. It has been deployed by the Taiwanese Ministry of Digital Affairs for AI policy analysis and party platform synthesis; by unions in California healthcare negotiations; and by researchers conducting broad-based opinion mapping on AI governance. In 2025, Carnegie Endowment for International Peace identified T3C alongside Polis and vTaiwan as part of the core infrastructure of AI-assisted deliberative democracy.

The ceiling, precisely stated: Polis, the Habermas Machine, and Talk to the City can surface consensus, map diversity, and preserve nuance within a shared epistemological framework. They are extraordinarily good at finding where people who share basic assumptions about what constitutes evidence and reasonable argument already agree or disagree without knowing it. When the assumptions themselves are what differs — as they do between a Quechua knowledge holder, a systems scientist, and a Sufi philosopher — the architecture has no mechanism for the encounter. It can cluster their responses statistically and preserve their language faithfully. It cannot identify that they are asking different questions, let alone create the conditions under which their different questions might generate something that none of them could reach alone.

AI-Enhanced Collective Intelligence Research

AI4CI and the Collective Intelligence Project

The AI for Collective Intelligence initiative (AI4CI), a UK national research hub spanning seven universities including UCL, Oxford, and Cambridge, and the Collective Intelligence Project (CIP) in the US (founded by Divya Siddarth and Saffron Huang), represent the most developed institutional expressions of the research agenda establishing collective intelligence as a rigorous field.

The foundational framework, developed by Woolley and Gupta (2024, *Patterns* journal), establishes that collective intelligence emerges from three interdependent components: collective memory (shared knowledge bases), collective attention (what the group focuses on), and collective reasoning (how the group processes information). A 2025 review by Riedl and De Cremer in ACM Collective Intelligence extends this to include AI agents as epistemic partners — 'integrative support structures' that can enhance all three dimensions simultaneously. The ACM Collective Intelligence 2025 conference, themed 'Diverse Intelligence,' brought together researchers on how human, animal, machine, and hybrid intelligences can interact to solve complex problems, reflecting a significant broadening of the field's scope.

The Collective Intelligence Project's expanding scope

The CIP has evolved significantly since its founding. Its Global Dialogues programme — seven rounds of structured deliberation across more than 70 countries, engaging over 6,000 participants in 2025 — represents an ambitious attempt to bring non-Western voices into AI governance. Each round uses Remesh, an AI-assisted live deliberation platform, to surface consensus across language barriers in real time. The data is made publicly available as open-source research infrastructure. In 2026, CIP announced that Global Dialogues will become standing global infrastructure, with annual longitudinal panels and a Global Trust and Expectations Index tracking how societal sentiment evolves as AI embeds itself in daily life.

CIP has also developed Weval — an open AI evaluation platform hosting over 100 community-contributed blueprints testing model configurations on socially complex tasks — and a Digital Twin Evaluation Framework that tests how reliably AI models represent the nuanced views of diverse demographic groups. In 2026, CIP is launching an Epistemic Evaluation Suite that will systematically measure whether AI systems are truthful, grounded, impartial, and reason in good faith. These are properties that current benchmarks do not measure — and that go directly to the question of what kind of knowing an AI system can actually hold.

The ceiling, precisely stated: Every system in this tradition has been developed for groups operating within a shared epistemological culture — typically Western, academically trained, and operating within scientific norms of evidence and argument. The architecture assumes shared foundations that the participants in a WhatIfWe encounter explicitly do not share. CIP's Global Dialogues expand geographic reach dramatically; they do not address epistemological incommensurability. A Quechua knowledge holder and a systems scientist participating in the same Global Dialogue round are given the same structured statement format, the same voting mechanism, and the same AI-assisted synthesis process — which produces consensus within whatever framework both accept, and cannot produce the encounter that occurs when their frameworks themselves are what is in question.

The Global Scaling Wave

Engaged California, the Global Inclusive AI Coalition, and institutional AI deliberation

Since Version 1.0 of this document, the deployment of AI-assisted deliberative democracy has moved from experimental pilots to institutional infrastructure at a pace that few anticipated. Three developments in 2025–2026 deserve detailed attention.

Engaged California

In February 2025, California Governor Gavin Newsom launched Engaged California — the first state-level permanent AI-powered deliberative democracy platform in the United States. The system, developed over two years by the state's Office of Data and Innovation and running on the Ethelo platform, operates as a three-way tool: it enables policymakers to listen at scale to citizens outside election cycles; it invites citizens to voice concerns and proposals on an ongoing basis; and it provides a space for Californians to find common ground with each other. AI synthesis — not AI opinions — does the work of organising thousands of individual contributions into actionable thematic maps while preserving the specific language of participants.

The first deployment addressed the aftermath of the 2025 Los Angeles wildfires, transforming over a thousand detailed resident responses into actionable insights while preserving participants' own language and elevating shared concerns around housing, insurance, and long-term resilience. The second focused on state employees generating efficiency ideas. In May 2026, Governor Newsom launched the programme statewide to gather citizen input on AI policy. The Carnegie Endowment for International Peace, which played a collaborative role in the programme's development, has described it as 'a new governing ecosystem' rather than a one-off civic exercise — a permanent feature of state government rather than a consultation tool.

The Global Inclusive AI Coalition

At the Paris AI Action Summit in February 2025, a Global Coalition for Inclusive AI was officially launched, co-initiated by Missions Publiques and the Stanford Deliberative Democracy Lab. The coalition aims to design and support deliberative processes at both local and global levels on AI governance. For its first deliberation cycle (2025–2026), it plans to combine decentralised deliberative dialogues with a global deliberative poll, aiming to engage 10,000 demographically representative citizens from over 100 countries — not AI experts, but ordinary citizens deliberating after structured learning. Governments including Switzerland, the United Kingdom, and France have joined, alongside over 80 civil society organisations and the UN.

The AI Penalty and the Deliberative Divide

A significant 2025 paper by Jungherr and Rauchfleisch (University of Bamberg and National Taiwan University) identifies what they call the 'AI penalty' in deliberation: the introduction of AI into deliberative processes creates a new participation divide, driven not by digital literacy but by participants' attitudes toward AI itself. People with negative preconceptions about AI are significantly less likely to participate in AI-mediated deliberation — which means the expansion of computational democracy through AI may systematically exclude exactly those communities most sceptical of Western technological frameworks. This finding is directly relevant to WhatIfWe's challenge: indigenous knowledge holders and contemplative practitioners from non-Western traditions may be

precisely the communities for whom AI-mediated deliberation feels most alien.

The ceiling, precisely stated: The global scaling wave represents a genuine and significant achievement in expanding the reach of deliberative processes. What it scales is sophisticated consultation within political and cultural diversity — diversity of values, interests, and lived experiences. It does not scale epistemological encounter. Engaged California, however carefully designed, processes Californian citizens' inputs through a shared framework of what constitutes a policy proposal. The Global Inclusive AI Coalition, however geographically representative, produces consensus statements within the frame of what AI governance questions Western institutions have decided are worth asking. The epistemological diversity that WhatIfWe is designed to hold is not reached by increasing the geographic and demographic diversity of participants within a shared deliberative architecture.

Warm Data and Transcontextual Knowledge

The International Bateson Institute

The International Bateson Institute, led by Nora Bateson, has developed what remains the most philosophically precise diagnosis of the problem WhatIfWe is attempting to address. Warm Data — defined as 'transcontextual information about and within the interrelationships that integrate elements of a complex system' — names a structural problem that technical collective intelligence systems have not solved.

The core insight: knowledge that exists only in the relationship between things is destroyed when extracted from its context. The moment you abstract a Vedic understanding of interdependence from the practice, community, and ecosystem in which it lives, you have something that resembles the original but lacks its functional capacity. The same is true of indigenous ecological knowledge, contemplative phenomenology, and the tacit dimensions of scientific intuition that precede formalisation. Bateson's concept of 'sympathesy' — mutual learning within living systems — extends this: complex systems are not made of parts interacting, but of co-evolving relationships. Any knowledge methodology that treats participants as discrete information sources rather than nodes in a living relational web will systematically miss what is most important.

The Warm Data Lab methodology has grown substantially since Version 1.0 of this document. Over 1,000 groups in more than 40 countries have now participated in Warm Data Labs held by more than 600 certified hosts. The methodology has been deployed in educational, organisational, ecological, and policy contexts. Recent academic work has begun developing Warm Data into a formal framework for 'transcontextual learning' — examining how people collaboratively learn across multiple contexts simultaneously, and how this kind of learning produces insights that no single-context approach can generate. The relevance to WhatIfWe is direct: the Warm Data Lab is the closest existing process to the kind of encounter WhatIfWe is designing for. The difference is one of scale, persistence, and instrumentation.

The ceiling, precisely stated: The Warm Data framework diagnoses the problem with extraordinary precision but does not yet provide scalable, persistent infrastructure for what it points toward. The Warm Data Lab creates temporary conditions — a structured group process that generates transcontextual awareness for the duration of the session. When the session ends, the conditions reset. WhatIfWe is attempting to create a sustained environment that accumulates across sessions and participants — one that builds a persistent record of what transcontextual encounter actually produces, and can instrument that record with AI facilitation that tracks emergence across time. The Warm Data framework does not yet include this layer; this is not a criticism but a deliberate methodological choice about scope.

Plurality and Collaborative Diversity

Weyl, Tang, and the epistemological frontier

The Plurality project — Glen Weyl and Audrey Tang, with a 2024 book and an associated research and policy agenda — represents the most ambitious current attempt to articulate a political and technological philosophy for harnessing diversity. Plurality argues that digital tools should channel 'the potential energy in social diversity that can erupt in conflict instead for progress, growth and beauty.' Drawing on Taiwan's decade of digital democracy, the project proposes tools ranging from augmented deliberation to new models for digital identity — technical architectures that recognise and strengthen collaboration across difference.

The Plurality framework is philosophically serious and empirically grounded in a way that most technology philosophy lacks. Tang's direct experience operating vTaiwan at national scale gives it a credibility that purely theoretical frameworks cannot claim. The concept of 'collaborative diversity' — the idea that difference, properly channelled, is a source of civilisational intelligence rather than merely a problem to be managed — is close in spirit to what WhatIfWe is attempting. Plurality is the closest existing framework to WhatIfWe's core argument at the political and technological level.

The limit is that Plurality's diversity is primarily political, cultural, and demographic. It addresses communities with different values, interests, and lived experiences — which is genuinely important and underserved work. What it does not address is epistemological diversity: the deeper difference between traditions that have developed genuinely distinct instruments for perceiving what is real. A Daoist philosopher and a systems ecologist may share democratic values and civic goodwill. What differs between them is not their politics but their instruments of perception — their trained capacity to notice different aspects of reality that the other tradition is structurally unable to see. That gap requires not just tools for diverse deliberation but infrastructure for epistemological encounter — encounter in which the instruments of perception themselves are what is at stake.

The ceiling, precisely stated: Plurality addresses diversity of values and interests. WhatIfWe addresses diversity of epistemic instruments — the deeper layer that political diversity tools are not designed to reach. Plurality and WhatIfWe are not in competition; they address adjacent but distinct problems. WhatIfWe begins where Plurality's framework reaches its frontier.

Contemplative AI and Abundant Intelligences

Attempts to address AI's epistemological limitations from the inside

Contemplative AI

Contemplative AI (Laukkonen, Timmermann, et al., arXiv 2025) attempts to instil contemplative principles — mindfulness, emptiness, non-duality, boundless care — directly into AI architectures. The research demonstrates that prompting AI systems to reflect on contemplative principles measurably improves performance on ethical benchmarks and cooperative behaviour tasks. The approach takes seriously the claim that contemplative traditions have developed genuine insights about the nature of mind, attention, and ethical action that AI systems might be made to embody — not merely to simulate.

Abundant Intelligences

The Abundant Intelligences initiative (Lewis, Whaanga, Yolgörmez, et al., published in *AI & Society*, 2024–2025) is an Indigenous-led, Indigenous-majority international research programme anchored at Concordia University in Montreal, with research clusters (called 'pods') at Western University, the University of Lethbridge, the University of Hawai'i–West Oahu, Bard College, and Massey University in New Zealand. Its six-year programme (2023–2029, funded by Canada's New Frontiers in Research Fund, the MacArthur Foundation, and the Schmidt Family Foundation) argues that the current trajectory of AI development suffers from fundamental epistemological shortcomings — systematically excluding non-Western, non-male, non-white ways of knowing from training data, evaluation frameworks, and design assumptions.

The programme organises its work across three axes: Integrations (how Indigenous knowledge systems can be embedded in AI architectures), Imaginaries (how Indigenous futurism and speculative practice can reshape AI development), and Intelligences (how the diversity of Indigenous knowledge systems can serve as a model for building AI systems that are genuinely epistemically plural). Specific areas include Indigenous storytelling and oral traditions, traditional land management practices, and social-neuro AI that models how humans draw on socio-cultural context to learn. The programme has produced significant academic output and has received attention from ScienceDaily, Concordia University press, and the international AI policy community.

The ceiling, precisely stated: Both Contemplative AI and Abundant Intelligences work within the AI system — trying to make AI more contemplative, more epistemically diverse, more aligned with a broader range of human knowing. This is important work, and WhatIfWe draws on it. But WhatIfWe's question is categorically different: not how to make AI wiser, but how to create the conditions under which human wisdom — already present in living traditions — can encounter other human wisdom in a way that generates something neither could reach alone. AI is the facilitator of that encounter, not its subject. The intelligence that matters in WhatIfWe is the collective intelligence of the human participants. What Abundant Intelligences and Contemplative AI are building is the kind of AI infrastructure that could serve as a more capable facilitator of that encounter — which makes them complementary to WhatIfWe, not alternatives to it.

The Epistemic Monoculture Problem

New research (2025–2026) that reframes the field's central challenge

Since Version 1.0 of this document, the concern about epistemic monoculture — the structural tendency of AI-assisted collective intelligence to flatten rather than amplify genuine difference — has moved from a scattered finding into a coherent and urgent research agenda. Several significant papers in 2025–2026 have sharpened the diagnosis considerably.

AI and the scientific monoculture (2026)

A February 2026 paper in *Communications Psychology* (Nature portfolio) by Jungherr and Rauchfleisch identifies a systemic epistemic narrowing that AI is producing across research fields: as AI tools become the dominant methodology, research agendas across disciplines are reframed through the AI lens, producing topical convergence (diverse questions become questions about AI and cognition, AI and institutions) and methodological convergence (shared analytic pipelines, toolkits, and benchmark datasets). The consequences are three: loss of intellectual diversity; erosion of methodological pluralism (the ability to triangulate phenomena through multiple epistemic lenses); and loss of field-level optionality (the capacity of a research ecosystem to pivot when the world changes). The paper notes that 'topics that do not fit easily within the AI narrative may struggle for legitimacy or resources.'

Epistemic diversity and knowledge collapse (2025)

A December 2025 preprint, 'Epistemic Diversity Across Language Models Mitigates Knowledge Collapse', demonstrates formally what the Doshi-Hauser finding (Science Advances, 2024) suggested qualitatively: AI monoculture produces knowledge collapse, but increased diversity across AI models can mitigate it — up to an optimal level. Beyond that level, distributing data across too many models reduces each model's performance. The paper recommends monitoring diversity across AI systems and developing policies that incentivise community- and domain-specific models. This is the closest existing technical research to WhatIfWe's multi-agent, tradition-oriented architecture — a formal demonstration that epistemic diversity in the AI facilitation layer produces better collective outputs.

Two distinct monoculture problems

A 2025 paper by Messeri and Crockett, cited in multiple subsequent studies, distinguishes two types of epistemic monoculture: monocultures of knowers (homogeneity in the standpoints from which knowledge is produced — the diversity of researchers, communities, and traditions contributing to a field) and monocultures of knowing (homogeneity in how knowledge is produced — the shared methodologies, frameworks, and evaluation criteria applied regardless of who is asking). WhatIfWe addresses both simultaneously: it brings genuinely diverse knowers (people from genuinely different traditions) into encounter through a design that refuses to impose a shared framework of knowing. This double address is what makes WhatIfWe architecturally distinct from any existing approach.

The 'AI penalty' in deliberation

Jungherr and Rauchfleisch (2025) identify a further dimension: AI-mediated deliberation introduces a new participation divide, driven not by digital literacy but by participants' preconceptions about AI. People who distrust or fear AI are significantly less likely to participate in AI-mediated processes —

which may mean that the expansion of deliberative democracy through AI systematically excludes exactly the communities most sceptical of Western technological frameworks. The implication for WhatIfWe is design-critical: the AI facilitation layer must be designed to earn trust, not assume it.

The ceiling, precisely stated: The research consensus that has emerged since Version 1.0 is clear: AI-assisted collective intelligence, deployed without specific attention to epistemic diversity, accelerates convergence on dominant frameworks rather than enabling genuine encounter between different ones. The proposed antidote — deliberate epistemic differentiation, diversity in the AI facilitation layer, community-specific models — points in exactly the direction WhatIfWe is working. The research identifies the problem. WhatIfWe is attempting to build the methodology that addresses it.

Multi-Agent Systems and Epistemic Diversity Research

From AI-to-AI interaction to human-AI encounter

The multi-agent AI research community has produced significant work on how AI systems with different orientations can produce better collective outputs than single-system architectures. The core finding — that diversity in the AI layer produces diversity in output, which produces more robust collective intelligence — maps directly onto WhatIfWe's architecture, where multiple AI instances with distinct tradition-oriented system prompts serve as the facilitation layer.

Key research includes: Argyle et al. (2023, *Political Analysis*), demonstrating that LLMs can be reliably 'silicon sampled' to represent specific demographic and epistemic perspectives; Bai et al. (2024, NeurIPS), showing that multi-agent debate between AI systems with different orientations significantly improves factual accuracy compared to single-model architectures; and a 2025 arXiv paper on 'Deliberate Collective Reasoning with Typed Epistemic Acts' (DCI), which proposes formal structure for multi-agent LLM deliberation, combining five elements that no prior multi-agent system had integrated: distinct orientations, typed epistemic acts, structured deliberation rounds, persistent memory, and human oversight.

The DCI paper is significant for WhatIfWe because it identifies precisely what is missing from current multi-agent systems: the human side of the encounter. Existing multi-agent research is almost entirely focused on AI-to-AI interaction. The question of how to sustain genuine encounter between humans from genuinely different epistemic traditions, mediated by an AI layer that preserves rather than dissolves the difference, remains largely unexplored. This is the specific implementation problem WhatIfWe is attempting to address.

The ceiling, precisely stated: Multi-agent research has established that epistemic diversity in the AI layer produces better collective outputs. It has not addressed the harder question: how do you design AI-mediated encounter between humans from genuinely incommensurable frameworks in ways that preserve the productive tension rather than collapsing it into a synthesis that one tradition would recognise and the other would not? This is where the technical research agenda ends and WhatIfWe's design problem begins.

Commercial Knowledge Synthesis Platforms

The market's current answers — and their common limitation

Several commercial platforms are attempting to address the knowledge synthesis problem from a productivity and enterprise angle. Understanding their common limitation clarifies what WhatIfWe is specifically not attempting to build.

Platform	What it does	Where it stops
Notion AI / Confluence AI	Enterprise knowledge management with AI synthesis	Optimised for organisational knowledge within a single institutional culture. No multi-tradition architecture.
Mem.ai	Personal AI memory and synthesis for individuals	Individual use only. No multi-participant encounter infrastructure.
Roam Research / Obsidian + AI	Networked thought tools with emerging AI integration	Strong for individual knowledge synthesis. No structured group encounter capability.
Kialo / Polis derivatives	Structured deliberation for group argument mapping	Epistemologically flat — all arguments treated as equivalent propositions. Cannot hold incommensurable frameworks.
Decidim	Participatory democracy platform (Barcelona, etc.)	Designed for civic participation within shared political frameworks. No epistemological encounter layer.
Remesh	AI-assisted live deliberation (used by OpenAI, CIP, UN)	Surfaces consensus within groups. Assumes shared framework of what constitutes reasonable statements.
Minerva / Discourse Graph	Academic knowledge synthesis	Designed for scientific literature within a single epistemic culture. Not for cross-tradition encounter.

The ceiling, precisely stated: The common limitation across all commercial platforms: they are built for knowledge synthesis within a shared framework, not for encounter between frameworks. The market has not produced — and is not trying to produce — infrastructure for epistemological encounter. Commercial incentives optimise for productivity and adoption, which rewards speed of consensus and breadth of integration. WhatIfWe requires the opposite: infrastructure that slows the encounter down to the pace at which genuine difference can be held.

What the Map Reveals: The Shared Ceiling

Surveying the landscape as it stands in May 2026 — expanded and significantly more developed than it was when Version 1.0 of this document was written — the pattern that v1.0 identified has become more pronounced, not less. The technology world has made remarkable progress on five adjacent problems:

- Surfacing latent consensus within a shared epistemological framework — Polis, the Habermas Machine, Talk to the City, Remesh. These tools are mature, deployed at national scale, and producing measurable democratic outcomes.
- Scaling deliberative processes geographically and demographically — Engaged California, the Global Inclusive AI Coalition, CIP's Global Dialogues. The infrastructure for bringing thousands to tens of thousands of people into structured deliberation across languages and borders now exists and is being institutionalised.
- Augmenting collective intelligence within epistemologically homogeneous groups — AI4CI, CIP's Weval, the broader multi-agent research programme. The evidence that AI can support collective memory, attention, and reasoning within groups sharing a framework is now robust.
- Diagnosing what happens to knowledge when extracted from its relational context — Warm Data Labs, now operating at over 1,000 groups in 40 countries. The framework is mature; its deployment is growing.
- Beginning to address AI's epistemological monoculture — Contemplative AI, Abundant Intelligences, the emerging research on epistemic diversity across model ecosystems. The concern is now a coherent research agenda rather than a scattered set of observations.

Each initiative is working at the ceiling of what its approach can reach — and naming, at that ceiling, the same gap from different directions. The pattern is consistent enough to be significant: every serious attempt to build infrastructure for collective intelligence eventually encounters the problem of epistemological difference and finds it cannot address it within its existing architecture.

The gap they are each naming is the same gap: the gap between intelligence that operates within a single way of knowing, however sophisticated, and the intelligence that can only emerge from sustained genuine encounter between genuinely different ways of knowing.

The new research on epistemic monoculture (Section 8) makes this gap more urgent, not less. The finding that AI assistance homogenises collective outputs within groups — reducing collective diversity even as it enhances individual creativity — means that the scaling wave described in Section 4 carries a structural risk: the more AI-assisted deliberation spreads, the more it may flatten epistemological diversity at scale. The antidote is not less AI but different AI — AI designed specifically to preserve and amplify difference rather than to find consensus.

WhatIfWe's Specific Position

WhatIfWe is not positioned as an alternative to the initiatives described above. It is positioned as the layer each of them identifies as missing and cannot itself provide. The most important technology-track guests in *The First Conversations* — Audrey Tang, Nora Bateson, Geoff Mulgan — are there precisely because they have pressed their respective approaches to this ceiling from the inside, and can articulate the gap from direct experience rather than from theory.

WhatIfWe's specific contribution is threefold:

The encounter design

A structured methodology for sustained encounter between people from genuinely different epistemic traditions, applied to problems that no single tradition can adequately address. This is the philosophical and methodological layer that the technology world identifies as missing but has not built. The interview series — conversations with people who have pressed their tradition to its frontier — is Phase 0 of this design: creating a public record of what frontier encounter looks like, building the network of participants, and developing the relational trust that any deeper encounter requires.

The infrastructure

Persistent memory, multi-agent AI facilitation with tradition-oriented system prompts, and a participant environment that rewards depth rather than speed. The research on epistemic diversity in multi-agent systems (Section 9) provides the technical foundation. The insight from the Doshi-Hauser finding and the subsequent epistemic diversity research is that the AI facilitation layer must itself be epistemically diverse — not a single model trained to be culturally sensitive, but multiple instances with genuinely distinct orientations that preserve rather than dissolve the difference between participants' frameworks.

The training corpus

Every encounter in the Lab generates training data for future AI systems. The synthesis notes, tension flags, and emergence candidates produced by the AI facilitation layer are the first systematic attempt to build a knowledge base of what epistemological encounter actually looks like — what it produces, what it fails to produce, and under what conditions genuine synthesis occurs. This corpus addresses directly the problem that Abundant Intelligences identifies: AI systems trained on predominantly Western knowledge sources cannot represent what non-Western traditions know. WhatIfWe's training corpus will contain what happens at the boundary — which is precisely what is missing.

The implementation question WhatIfWe needs to answer — what does a virtual, technology-driven knowledge lab for epistemological encounter actually look like in practice — is the honest frontier this project is working at. We know what it is not. We know what it is trying to do. The conversations and encounters of Phase 0 and Phase 1 are how we find out.

Key References

- Riedl, C. & De Cremer, D. (2025). *AI for collective intelligence*. ACM Collective Intelligence.
- Woolley, A. & Gupta, P. (2024). *Collective memory, attention, and reasoning*. Patterns.
- Doshi, A.R. & Hauser, O.P. (2024). *Generative AI enhances individual creativity but reduces collective diversity of novel content*. Science Advances.
- Jungherr, A. & Rauchfleisch, A. (2025). *Artificial Intelligence in Deliberation: The AI Penalty and the Emergence of a New Deliberative Divide*. arXiv:2503.07690.
- Jungherr, A. & Rauchfleisch, A. (2026). *AI is turning research into a scientific monoculture*. Communications Psychology (Nature).
- Tessler, M.H., Bakker, M.A., et al. (2024). *AI can help humans find common ground in democratic deliberation*. [Habermas Machine]. Google DeepMind.
- Volpe, A. (2025). *Toward an artificial deliberation? On Google DeepMind's Habermas Machine*. Ethics and Information Technology.
- Palomo Hernández, N. (2025). *Towards Automating Deliberation? The Idea of Deliberative Democracy Embedded in Google's Habermas Machine*. AAAI/ACM AIES.
- Lewis, J.E., Whaanga, H. & Yolgörmez, C. (2024–2025). *Abundant intelligences: placing AI within Indigenous knowledge frameworks*. AI & Society.
- Laukkonen, R., Timmermann, C., et al. (2025). *Contemplative AI*. arXiv.
- Wan, K. & Kalman, Z. (2025). *Epistemic differentiation in AI-assisted collective intelligence*. [Referenced in AI4CI literature].
- Argyle, L.P., et al. (2023). *Out of one, many: Using language models to simulate human samples*. Political Analysis.
- 'Epistemic Diversity Across Language Models Mitigates Knowledge Collapse.' arXiv:2512.15011 (December 2025).
- Bateson, N. (2024). *Transcontextual learning and Warm Data*. International Bateson Institute.
- Weyl, G. & Tang, A. (2024). *Plurality: The Future of Collaborative Technology and Democracy*.
- California Governor's Office (2025–2026). *Engaged California programme documentation*.
- Missions Publiques & Stanford Deliberative Democracy Lab (2025). *Global Coalition for Inclusive AI, launched at Paris AI Action Summit*.
- Carnegie Endowment for International Peace (2026). *Realizing the Potential Gains of AI-Enabled Deliberative Democracy*.
- ACM Collective Intelligence 2025 Conference Proceedings. Theme: Diverse Intelligence. La Jolla, CA, August 2025.