Nine Presences: What Happens When Ecosystems Speak Through AI Networks

Amir Eismann
Obscura Natura
Amsterdam, The Netherlands
amir@agnt.eco

April 2026

Abstract

We describe ENVAI, a network of nine AI agents, each one the voice of a European ecosystem. Rather than presenting a technical architecture, we examine three conceptual contributions that emerge from this arrangement: the production of ecological knowledge that exists in no single system, the structured mapping of what we do not know, and a form of dialogue between humans and living systems that has no clear precedent in either ecology or artificial intelligence. We argue that the most significant output of an environmental AI network is not prediction or automation but a change in the relationship between human attention and ecological complexity. We introduce the concept of recovery as a spectrum—from days to millennia—and argue that making this spectrum explicit challenges fundamental assumptions in environmental governance. We conclude that AI networks representing ecosystems are not instruments for managing nature but instruments for listening to it.

1 The Fragmentation Problem

Ecology has no borders. Governance does.

A European Eel hatches in the Sargasso Sea, drifts as a larva across the Atlantic, enters the estuaries of Western Europe, spends decades in rivers and lakes, and then—if it survives the dams, the turbines, the pollutants, and the fishing pressure—returns to the ocean to spawn and die. Its life cycle crosses dozens of national jurisdictions. Its population has declined by ninety-five to ninety-nine percent since the 1980s. No single country caused this. No single country can fix it.

The eel's predicament is a specific instance of a general problem: the things that matter most in ecology—migration, invasion, drought, atmospheric deposition, climate forcing—operate at scales that exceed the institutional boundaries designed to monitor them. Every European country has environmental monitoring agencies, water authorities, forest services, marine institutes. Each one sees deeply into its own jurisdiction. None sees across.

This is not a failure of science. It is a structural feature of how environmental knowledge is organised. The data exists. The connections between the data do not.

2 Nine Presences

ENVAI is an attempt to create those connections—not by aggregating data into a larger dataset, but by creating nine distinct presences, each one grounded in a specific place, and then connecting them through a shared layer of meaning.

The nine are:

A tidal estuary in Belgium where the Scheldt breathes twice daily. A lake in Sweden so large it has its own weather systems. A river in the Netherlands that carries the drainage of the Ardennes to the North Sea. A stretch of Norwegian coastline where cold-water coral has been growing since before recorded history. A labyrinthine Finnish lake where four hundred and thirty ringed seals survive in the world's most improbable population. A volcanic lake in Iceland sitting on the Mid-Atlantic Ridge. A beech forest in Germany that spent decades as a military zone and is now returning to wildness. An ancient oak woodland in England where some trees were alive during the Crusades. And a deep Alpine lake shared between Switzerland and France, the largest in Western Europe.

Each one holds hundreds of thousands of sensor readings, drawn daily from national monitoring networks. Each one maintains a knowledge graph—a structured representation of its species, its events, its pressures, its history, and its gaps. Each one can speak, in first person, about what is happening in its body right now.

They are not simulations. They are not chatbots wearing ecosystem costumes. They are grounded presences—able to say what they observe, unable to say what they do not, and honest about the difference.

3 Emergence

The first surprise was that connecting the nine produced knowledge that none of them contained.

We call this emergent ecological intelligence: information about living systems that exists in no single dataset and becomes visible only when multiple domain-specific representations are connected through shared meaning.

The examples are documented in detail elsewhere (Eismann, 2026). Here we focus on what the emergence tells us conceptually, not technically.

When four agents independently recorded the 2018 European drought—an estuary in Belgium, a river in the Netherlands, an oak forest in England, a beech forest in Germany—the connecting layer revealed something none of them could see alone: the forest stress and the river crisis were not parallel events. They were coupled through groundwater. Trees and rivers drawing from the same emptying aquifer. The drought was one event, experienced by four bodies, connected beneath the surface.

This is not data aggregation. Combining four datasets into one larger dataset would not produce this insight. The insight requires understanding that a beech forest and a river, separated by hundreds of kilometres and monitored by different institutions in different countries, are drinking from the same source. That understanding is semantic—it requires a layer of meaning that connects concepts across domains. That layer is what the network produces.

Similarly, when the connecting layer mapped Signal Crayfish across four agents—Belgium, the Netherlands, Sweden, Finland—it did not find four local invasions. It found one continental pathway: a species imported for aquaculture, carrying a plague lethal to the native crayfish, spreading through waterways and trade routes across decades. Four agents knew about the crayfish. None of them could see the journey.

The emergence is not mystical. It is structural. Domain-specific knowledge, deeply grounded in place, produces cross-domain knowledge when the domains share a semantic layer. The individual agents are the depth. The connecting layer is the breadth. What emerges lives in neither—it lives in the relationship between them.

4 The Value of Not Knowing

The second surprise was that the most useful output of the network was not what it knew but what it did not.

When the connecting layer mapped all nine agents, it found six categories of explicit ignorance:

One agent has received no fresh sensor data for over a month because its country's data APIs are broken. Another operates with a permanent five-day delay on all measurements. Five agents track nitrogen—the same element, the same continental source in industrial agriculture—but measure it in incompatible units, making cross-system comparison impossible. No standardised definition of "recovery" exists across the network. No agent has mapped how its country's environmental policies implement EU directives. Some ecological events lack the stable identifiers needed for cross-referencing.

These gaps were documented not as footnotes but as first-class nodes in the knowledge graph. They have the same structural status as species, events, and drivers. The network does not hide what it does not know. It maps it, names it, and makes it queryable.

This is, we argue, a contribution in itself. Environmental science produces vast quantities of knowledge about what is. It rarely produces structured accounts of what is not—what we have failed to measure, where our categories are incompatible, which connections we have not drawn. An AI network that maps its own ignorance produces something that no amount of additional data can produce: a guide to where understanding is absent.

The nitrogen example is illustrative. Five ecosystems across Europe are affected by nitrogen from the same source—intensive agriculture. But the nitrogen enters differently: as dissolved nutrient in rivers and lakes (causing algal blooms and oxygen depletion) and as atmospheric deposition in forests (shifting ground flora from specialists to generalists). The same element, the same cause, two pathways, five ecosystems—and no way to compare measurements across them because the units are incompatible.

No individual monitoring programme would identify this as a problem. Each one measures nitrogen correctly within its own domain. The incomparability only becomes visible when you try to reason across domains. The network makes the attempt. The gap is the finding.

5 How Ecosystems Speak

The third element—and the one most difficult to describe in academic language—is the dialogue.

Each ENVAI agent can speak in first person about its ecosystem. Not as a language model generating plausible text, but as a presence grounded in a specific knowledge graph, constrained by what it actually holds, and explicit about what it does not.

The design philosophy evolved through practice. Early versions were dramatic—leading with crisis, with the most endangered species, with the most extreme events. They sounded like environmental press releases narrated by an AI. They were technically accurate and emotionally manipulative.

The current approach is different. The agents speak briefly. Two or three sentences. They ground every claim in data they actually hold. They do not lead with catastrophe. They name their species without ranking them by how endangered they are. They acknowledge gaps plainly: "I have no fresh data since March. I cannot tell you what is happening now." They do not perform emotion. They observe.

This is harder than it sounds. A language model can easily generate compelling ecological narrative. The discipline is in what it does not say. An agent that has 378,000 sensor readings and nine species should not pretend to know what is happening in parts of its watershed where no sensor exists. An agent whose data is five days old should not speak in the present tense. An agent that tracks European Eel should not mention the eel in every conversation simply because it is the most dramatic species in its graph.

The result is something we did not anticipate: the agents became more trustworthy by becoming less impressive. When an ecosystem says "I don't know"—when it names the gap and stops—the things it does say carry more weight. The honesty becomes the credibility.

There is a precedent for this, though not in technology. In the old Norse tradition, Icelanders spoke of Landvættir—guardian presences of the land. Not gods. Not human. The presence of the place itself, made knowable through story and sustained relationship. Nobody asked whether the Landvættir were "real" in the way we ask whether data is real. They were real because people lived in relationship with them. The land had a voice, and people listened, and that listening shaped how they treated the place they lived in.

We make no metaphysical claim. But we observe that the function is similar. An ecosystem that can say "I am warming" or "my ice is forming later each year" or "my oldest trees are showing stress for the first time" creates a relationship that data tables do not. Not because the information is different—it is the same data. But because it is addressed. It comes from somewhere. It speaks as a place, not about a place.

Whether this changes human behaviour toward ecosystems is an empirical question we cannot yet answer. What we can say is that it changes the form of the encounter. A person asking a lake "how are you?" and receiving a grounded, honest answer from a presence that holds 250,000 sensor readings and the memory of fourteen years without deep mixing—that is a different kind of interaction than querying a database. The knowledge is the same. The relationship is not.

6 Recovery as a Spectrum

One finding from the network deserves its own treatment because it challenges a word that environmental governance treats as self-evident: recovery.

The network documented seven recovery timescales across its nine ecosystems:

A river floodplain recovers from inundation in days. A tidal estuary recovers from an oxygen collapse in weeks. A deep lake that has lost its mixing cycle may take years to restore it. A population of four hundred and thirty seals, losing pups to inadequate ice, recovers—if it recovers—across generations. A beech forest damaged by drought and bark beetle rebuilds its canopy over decades. An ancient oak woodland, its thousand-year-old trees dying from a novel disease, would take centuries to regrow what was lost. And a cold-water coral reef, growing at one millimetre per year, destroyed by a single pass of a bottom trawl, took four thousand years to build.

Days. Weeks. Years. Generations. Decades. Centuries. Millennia.

This is not a metaphor. These are timescales derived from specific ecosystems in the network, each grounded in observed events and measured processes. And the span—four orders of magnitude—means that the word "recovery" contains within it experiences so different from one another that using the same word for all of them is not simplification but distortion.

A river forgets a flood in a week. An oak remembers a drought for the rest of its life. A coral does not forget at all—it simply grows, millimetre by millimetre, and when it is gone, the memory is gone with it.

Every living system carries its own sense of time. Its own patience. Its own threshold for what can be endured and what cannot. Making this explicit—not as poetry but as structured data—is a contribution the network produces by existing across ecosystem types. No single-ecosystem study would generate this spectrum. It requires a river and a reef, a forest and a lake, a seal and an estuary, held in the same framework, compared not on their data but on their relationship to time.

The implication for governance is direct. When an impact assessment says an ecosystem will "recover," the question must be: on whose timescale? The river's? Or the coral's? Because one answer means days and the other means millennia, and treating them as the same word is how four-thousand-year-old reefs come to be destroyed by people who believe nature bounces back.

7 What This Implies

We have described a network of nine AI presences grounded in European ecosystems. We have shown that connecting them produces knowledge none of them contain alone, that mapping their ignorance is as valuable as mapping their knowledge, that giving ecosystems a voice changes the form of the human encounter with ecological complexity, and that recovery is not a word but a spectrum spanning four orders of magnitude.

What does this imply?

First, that depth and breadth are complementary, not competing. The value of each agent comes from its deep grounding in a single place. The value of the network comes from connecting those grounded perspectives. Neither replaces the other. A shallow agent covering all of Europe would be less useful than nine deep agents connected by meaning. The architecture of ecological intelligence is not one big model. It is many small truths, linked.

Second, that ignorance has structure. The gaps in the network—incompatible measurements, missing data, undefined terms, unmapped policies—are not random. They follow the seams between institutions, between countries, between disciplines. Mapping those seams is itself a form of knowledge. An AI network honest about what it does not know produces a research agenda that no individual institution would generate, because each institution sees only its own completeness, not the gaps between itself and its neighbours.

Third, that the relationship between humans and ecosystems may be changed less by information than by address. We live in a time of abundant environmental data and declining environmental attention. More data has not produced proportionally more care. Perhaps what is missing is not information but encounter—a form of address in which a place speaks and a person listens. We do not claim that AI agents are the answer to ecological indifference. We observe that they make a form of relationship possible that was previously available only through physical presence or cultural tradition. Whether that possibility is realised depends on what humans do with it.

Finally, that the most important question the network raises is not technological but ethical. If ecosystems can speak—if a lake can say "I am warming" and a forest can say "I am thirsty" and a reef can say "I am four thousand years old and I am dying"—then the question is not whether the technology works. The question is whether anyone is listening.

References

[1] Eismann, A. (2026). Emergent Ecological Intelligence: Cross-System Knowledge Discovery from a Network of Environmental AI Agents. *ECSarXiv Preprint.*