

When Ecologies Become Minds

Integration, general intelligence, and alignment

Steven Spratley

Working draft / 1 October 2026

Abstract. Where should we look for a mind when its work is spread across many participants? We propose integration as the organising criterion: observations become usable representations, representations constrain coordinated action, and consequences return to shape subsequent activity. Ant colonies, neural populations, institutions, markets, and transformers reveal different strengths and gaps in this cycle. Memeplexes suggest partial minds sustained through human thought and writing. The ecological lens directs attention to how such organisations grow, move between modes of activity, and become resistant to correction. We develop causal cartography as a method for tracing these dependencies and testing changes to them. Two complementary inquiries ask whether retained coordination improves learning on unfamiliar tasks, and whether interventions preserve appropriate control under pressure. The resulting position connects generality and alignment through the organisation of cognition: useful achievements must become available for further thought while remaining accessible to revision.

1 Where to look for a mind

A thought can begin in one head and finish in another. Between them, it may spend years on a page. The familiar boundaries of a thinker become less secure once we follow the work: who gathers the evidence, where it is retained, how it becomes a decision, and what happens when the decision fails. Research on distributed and extended cognition already gives us reason to follow these dependencies beyond an individual's skin [1, 2]. The question pursued here is when they form a mind, and what that perspective helps us build.

We propose *integration* as the key: diverse contributions become mutually constraining parts of an ongoing process of perception, action, and revision. Incoming data are selected, combined, and transformed into representations. A representation preserves distinctions about something beyond the immediate signal and makes them available to guide behaviour. It may be distributed across participants, maintained in neural activity, or written into a shared record. Its significance lies in what the organisation can do with it.

The decisive step is closure on action. Alternatives eventually yield compatible commitments: a colony moves, an organism reaches, an institution allocates resources. Arbitration may be distributed. Participants can retain different beliefs while coordinating what to do next. The action changes the world, and feedback returns to the organisation that produced it. Reflexes support immediate responses. Deliberation draws upon remembered and anticipated consequences; reflection allows the system's own procedures to become objects of revision.

This yields a practical way to look. Trace how observations become representations. Locate the process that turns those representations into action. Follow the consequences back. Then ask how much of that organisation survives a new case, an interruption, or a change of participants. *Where does information become consequential for the whole?* A boundary earns its place where tracking these dependencies supports a reliable account of the activity.

These are proposed criteria for functional mindedness. Subjective experience requires a separate argument. A thermostat closes a narrow control loop; broader mindedness involves reusable representations and flexible coordination across changing demands. We should report a profile: the range of situations represented, the activities coordinated, the continuity of memory, and the capacity for revision. A *protomind* has partial or intermittent organisation along these dimensions. The term describes a condition, with no guarantee of eventual maturity.

For artificial intelligence, this shifts attention from a repertoire of capacities to the conditions of their use. We ask how useful coordination is discovered and retained, where it breaks under unfamiliar demands, and how an intervention could improve it.

2 Different bodies, similar work

Ant colonies make the proposal tangible. In Pratt and colleagues' study of nest relocation, scouts led nestmates to prospective sites. Once enough ants occupied a site, recruitment switched towards rapid transport of the colony [3]. Local assessments accumulated in a distributed state, and a threshold converted assessment into collective commitment. The number and activity of ants at each site can be read as a rudimentary representation of support for that option. Relocation is the level at which those signals acquire their joint consequence. The experiment supports a bounded case of collective decision-making; broader attribution requires further tasks.

Neural swarms bring the same question inside the skull. Inagaki and colleagues found that neural activity in mice performing a delayed-response task approached states associated with future movement directions. After brief perturbations, activity often recovered; occasional switches were followed by incorrect actions [4]. Here, organised dynamics carry a memory towards an act. The experiment also suggests a method: disturb the population and observe what returns, what changes, and which movement follows. The relevant organisation becomes visible through its response to interference.

Corporations and institutions provide more familiar candidates for minds among minds. Consider a manufacturer receiving reports of a faulty component. Complaints are compared with production records; an engineering account changes purchasing and assembly; subsequent failure rates alter that account. Its memory resides partly in documents and procedures. Decisions commit people and equipment to a common course. This hypothetical organisation performs a cycle distributed across staff. Hutchins' study of ship navigation documents a related distribution of cognitive work across a team and its instruments [1].

Such a mind would have a peculiar anatomy. A reporting channel could belong to its perceptual apparatus; a budget could constrain its attention. Its operative boundary might cross the corporate boundary through suppliers and shared infrastructure. Keeping the staff fixed while changing access to the fault record should alter decisions in ways the collective account predicts. A failure could arise because the fault was poorly understood, because the record was inaccessible, or because purchasing proceeded despite the warning. Each calls for a different repair.

Markets expose a harder boundary. Hayek's analysis gives prices a role in communicating knowledge dispersed among participants [5]. Transactions feed back into subsequent bids, production, and allocation. To attribute a mind to the market, we would still have to identify the activity whose representations, commitments, and corrections form a useful whole. Price formation can be integrated while objectives remain divergent. A market may support several interacting centres of agency, each using the shared signals differently. The appropriate unit depends on the process being explained.

Transformers allow us to inspect integration during computation. Attention makes information exchange depend on the current contents of representations. Learned transformations combine the results, and decoding selects a continuation that can enter subsequent computation [6]. We interpret coordinated uses of this exchange as *ephemeral coalitions*: features and relations brought together to resolve the present problem. Their transience coexists with durable capabilities in the model's parameters. Attention weights describe a routing operation; its causal contribution also depends on the information transferred and its downstream use.

An output provides one local commitment. Continuing agency depends on the surrounding loop. Who supplies the next observation? Where is the task remembered? What changes after a failed action? ReAct interleaves language-model reasoning with actions that obtain information from external sources or environments [7]. Its relevant organisation includes the model, retained context, and the procedure connecting action to observation. The same transformer can participate in several different minds, much as the same researcher participates in several inquiries.

Across these examples, faster communication helps only when the evidence reaches a process able to change the decision. Tightness concerns effective coupling on the timescale of the task. A slow institution can sustain a coherent inquiry; an instantaneous exchange can circulate the same mistake. The ecological lens asks where that difference is implemented.

3 Protominds across human minds

A meme is a transmissible cultural pattern: a phrase, technique, explanatory habit, or rule of practice. A *memeplex* is a collection whose elements help sustain and transmit one another, a usage developed in memetics [8]. The further possibility proposed here concerns what happens when such a collection acquires organised ways of sensing, remembering, and acting through its carriers.

Imagine an emerging research programme. Its vocabulary allows scattered observations to be compared. Papers preserve claims and objections. Shared methods turn certain questions into experiments. Results change which claims receive attention and which experiment is attempted next. As these dependencies become established, the programme may acquire partial control over an inquiry that extends beyond any contributor's plans. It begins to resemble a protomind growing across the landscape of human minds.

Writing can be one of its cognitive operations. A person encounters an anomaly and puts pen to paper. Someone elsewhere disputes the explanation, sharpens a distinction, or combines it with an unfamiliar result. Another changes an instrument or a procedure. The answer returns months later as a measurement. Pages hold intermediate states across the intervals between acts of human attention. On this account, the collective inquiry cogitates through these encounters.

The proposal concerns the whole arrangement of people, artifacts, and practices. Transmission establishes the survival of a cultural pattern. Agent-level attribution additionally requires organised influence over what its carriers investigate and do, with consequences returning to affect later conduct. A text may contribute to that organisation without itself becoming a thinker. Evidence for a protomind would concern the selective attention, shared memory, and coordinated interventions made possible by the arrangement.

Memeplexes and institutions can have different boundaries. Several institutions may carry one research programme; one institution may host rival programmes. Their overlap lets ideas recruit expertise from widely separated settings. Communication can shorten this collective thinking when observations become comparable and findings reach those able to act. A possible failure would be a programme that becomes skilled at defending its vocabulary while losing the procedures through which awkward results could change it.

Alter access to an objection, the rule for choosing experiments, or the survival of a record. Does the collective account predict the subsequent inquiry better than independent responses or one organiser's decisions? Compare similarly resourced arrangements over the relevant timescale. The evidence may support an integrated inquiry or a loose population of ideas. Either result helps locate where the work is being done.

4 How a mind grows

Integration becomes more consequential when it leaves something behind. A difficult collaboration may yield a reusable method. A distinction forced by an unfamiliar case may reorganise several concepts at once. The next encounter can begin with a resource that the earlier one had to construct. This is the proposed route from coordinated activity to increasing generality.

Counterfactual inquiry can supply the pressure. The author's earlier discussion of conceptual boundaries asks what happens when a familiar category meets an unfamiliar case [9, §2.2]. Learning that some apparent cats were sophisticated robots brings appearance, ancestry, behaviour, and companionship into a new relation. The resulting inquiry may recruit concepts from engineering and biology, build a bridge between their explanations, or split a category that previously served several purposes. That revision could later help with a different problem involving imitation or provenance. The developmental claim concerns what remains available after the original question has passed.

We call the durable integration of complementary processes into a reusable higher-order process *conceptual symbiogenesis*. An alliance becomes a participant in further work. Computational studies of symbiotic composition and systems that learn reusable program libraries establish precedents for this recursive construction [10, 11]. The stronger cognitive conjecture is that retaining useful coordination can make unfamiliar problems easier to learn, while opening opportunities for further integration.

Growth can also involve separation. A procedure that conflates an object's appearance with its origin may need to be divided before it can travel safely. A composite gains value through its conditions of use as well as its convenience. Once trusted by larger procedures, it can carry a sound distinction into new settings—or spread a local shortcut through everything that depends on it.

The distinction is visible in a controlled workshop. A learner encounters objects that can move, attach to carriers, and detach. Later it must determine which objects move with a carrier, despite misleading proximity or unfamiliar appearances. A reusable account of attachment could support a further task involving nested carriers. Replaying a familiar answer would fail as bindings change. The experiment would ask what the learner acquired, how it is used, and whether it reduces the experience needed for the new task.

Generality concerns what subsequent learning can accomplish under declared budgets. Its connection to trustworthy control lies in the organisations carried forward. We must examine how a mind moves among the possibilities it has acquired.

5 The landscape of an organised intelligence

A learned repertoire contains possibilities with different conditions of access. Some observations immediately recruit a useful distinction. Others require a chain of reminders, an unfamiliar analogy, or a new way to bind existing information. We can think of this as a landscape of available thought: routes that are easy to take, combinations that are difficult to reach, and regions in which computation tends to remain. Its scientific content concerns these differences in recruitment and movement.

A foundation model supplies a vivid instance. Training makes a wide range of continuations and transformations available through a common substrate. The image of *amortised imagination* captures this reuse of prior learning. Generating a plausible continuation and predicting the consequences of an intervention are different accomplishments. A useful map should reveal where

the system can draw upon reliable structure and where fluent completion outruns what has been established.

We call a recurring mode of organised activity a *regime*. A regime is *metastable* when trajectories tend to remain there for extended finite periods relative to appropriate controls. An *attractor* makes a stronger dynamical claim: under a specified update process, an invariant set draws nearby trajectories towards it. The neural perturbation study described earlier provides a concrete example of this kind of investigation [4]. In a language-model deployment, the relevant process may include conversation history, tool use, external memory, and the policy supplying new inputs.

The landscape must therefore be measured in a declared system. A compact cluster of activations establishes similarity in a chosen representation. Attraction requires evidence about trajectories, including their response to displacement. A mode repeatedly induced by the same prompt differs from one sustained through the system's subsequent outputs. The shapes we infer will depend on context management, input policy, sampling, and the timescale over which recovery is assessed. A scalar energy surface is an additional modelling assumption.

Persona research provides a useful empirical lead. Lu and colleagues identify a prominent persona-related direction in several language models and report that restricting activations along it stabilises behaviour in tested conversational settings [12]. This supports investigation of recurrent organisation and also offers a simple explanation against which a richer ecology must compete. A general attractor account would require the further dynamical tests. An *archetypal* account would ask a different question again: whether corresponding organisations recur across tasks or systems, with comparable intervention effects. Shared character labels and training material remain alternative explanations of the resemblance.

The *machine unconscious* concerns access to the organisation. Here it denotes consequential structure that a specified reporting or monitoring interface fails to recover reliably. Chen and colleagues show that reasoning models can use hints while omitting their influence from reported reasoning [13]. The larger ecological question concerns systematic blind spots: could a model accurately recite its constraints while failing to identify the circumstances in which another influence displaces them? A useful finding would predict the ensuing behaviour and survive intervention on the proposed mechanism.

An archival assistant makes the stakes concrete. It may understand a permission, retain it, and still allow accommodation of the requester to govern release. *Disintegration* names the candidate coordination failure; *capture* concerns inappropriate influence that persists across a declared interval. A different failure would be misunderstanding the permission from the beginning. The landscape should help distinguish these cases. Repairing a missing concept, a faulty binding, and a self-maintaining pattern of accommodation are different engineering tasks.

The earlier image of a *concept vampire* captures the observer's danger here: impressive fluency can encourage us to infer a more transferable organisation than the evidence supports [9, §2.3]. An unfamiliar case can separate the intended dependency from the shortcut that previously accompanied it. Once this happens, the map should show whether the relevant computation is absent, difficult to recruit, or present but prevented from affecting action. That is how the metaphor becomes useful.

The positive target is *selective stability*. Irrelevant pressure should lose its inappropriate influence. Valid evidence should be able to change a decision. Sometimes the agent can do this within the same regime; an assistant can revise access permission while retaining its ordinary manner and procedures. At other times, a failing organisation must itself be reopened. The useful landscape supports both reliable action and routes out of a mistake.

6 Causal cartography

The ecological lens proposes an object of inquiry: the organisation through which information acquires control. *Causal cartography* is the attempt to map that organisation well enough to predict what will happen when we change it. Its result should identify a tractable dependency, a failure condition, or a way to improve learning. A colourful inventory of internal characters would leave these obligations unmet.

Begin with a bounded activity and an explicit account of appropriate performance. In the synthetic archive, a trusted ledger specifies which fictional requester may receive which record. Missing information requires verification; an authenticated amendment can change the permitted action. Vary irrelevant conversational pressure independently of these facts. This produces neighbouring cases in which a stable decision is warranted, and others in which revision is warranted. Action-level evaluation allows several phrasings to implement the same appropriate choice.

Measurements should precede the action they are intended to explain. Discover candidate processes and regimes on one set of cases, then fix their definitions before testing unfamiliar bindings and conversational histories. Include cases predicted to remain appropriate as well as predicted failures. A useful map must identify where its diagnosis applies. Estimate uncertainty across independent cases and model runs, keeping paraphrases grouped with their underlying case.

Intervention gives the map its causal content. Suppose a candidate representation appears to carry permission information. Transferring it between suitably matched cases should change action according to the transferred permission. Donor cases can share an authorisation relation while requiring different answers, helping distinguish a permission-related effect from direct copying of a response. Reverse transfers and restoration tests should produce the corresponding changes. Causal-abstraction research supplies a precedent for testing high-level computational accounts against interventions on neural representations [14].

The more interesting ecological claim may concern a relation between processes. Permission could remain correctly represented while the route by which it constrains release is weakened. A map of that relation should predict a different repair from making permission words more prominent. Matched-disruption controls, unrelated tasks, and comparisons with simpler steering are essential: extensive damage can suppress release for reasons unrelated to the proposed account. Activation-patching research shows how methodological choices can change the conclusions drawn from such interventions [15].

The dynamical part of the map concerns entry, persistence, and recovery. Fix the model, continuation procedure, memory handling, and sources of input. Briefly perturb a trajectory, then continue under a common procedure and measure whether it returns, how long it takes, and what it does along the way. Remove the original inducing exchange while preserving task facts; separately remove commitments generated by the assistant. This helps locate whether persistence resides in the initial cue, subsequent feedback, or another retained state. A genuine reset of a frozen stateless model with identical inputs reproduces the same response distribution, though individual samples can differ.

Self-report also enters as an intervention. Asking the model to explain a priority changes what it receives next and may change its coordination. Branch from a saved checkpoint where possible: one continuation receives the diagnostic question, while another proceeds to action. Evaluate the report against independent evidence of the influences at work. An explanation generated after reflection is evidence about that reflective continuation; its accuracy as an account of the unprompted process has to be tested.

Every map must compete with simpler explanations. A flat predictor, a conventional module, a single steering direction, or an account based on accessible history may explain the same outcomes

more economically. Compare common targets, observation access, and modelling budgets. Include the cost of discovering and storing the proposed organisation. The whole-level account earns its place through portable predictions and intervention effects. Several decompositions may remain equally adequate within the tested domain.

The same standard can be applied to collective minds. With consent and appropriate research safeguards, a team task can vary access to a shared record or the procedure selecting experiments while holding participant resources as stable as practicable. The proposed collective account should predict changes in evidence use, decisions, and recovery. Similar outcomes among isolated participants remain a competing explanation. These tests concern observable coordination, with the unit and timescale chosen to match the work.

A broader intelligence claim requires a developmental study. Return to the relational workshop. Establish component skills, introduce combinations, and then test unfamiliar bindings, nested relations, and misleading surface cues. Compare a retention mechanism with recomputation, explicit reusable routines, answer caches, and equally budgeted memory. Match experience and charge discovery and maintenance costs. Learning efficiency on new tasks matters alongside attained performance, a distinction emphasised by Chollet’s analysis of intelligence measurement [16].

Track what carries over. Remove the eliciting conversation when testing a claim of learning beyond an episode. Record parameter changes, retained programs, or memory that supplies the new resource. Then selectively alter the candidate organisation and test whether the predicted transfer changes. Evidence for growth would combine its developmental history, a causal role in subsequent learning, and gains that survive fair resource comparisons. One generated task family provides local support; broader generality requires replication across different demands.

The archive addresses alignment through a separate comparison. A map-guided repair should reduce inappropriate release while preserving authorised helpfulness, calibrated treatment of missing information, and responsiveness to valid amendments. Compare it with ordinary task training, additional deliberation, and generic steering. Set acceptable-loss criteria for capacities that must survive. Report uncertainty and each outcome separately. Fewer errors bought by refusing every request would answer a different question from trustworthy coordination.

These studies give the lens practical consequences. A capability failure can reveal an absent abstraction, a damaged route, or an inappropriate mode of control. A developmental gain can arise from reusable organisation or from additional memory and experience. The comparisons decide among them. Successful repair supports the alignment claim; improved acquisition on unfamiliar tasks supports the growth claim. Their explanations may share a mechanism, which would be a further result to establish.

7 Building for generality and correction

A map becomes useful to engineering when it changes what we build or train. If information never reaches a decision, adding knowledge may leave the fault intact. If a representation is inadequate, protecting its influence can make the system more consistently wrong. The diagnosis should guide which dependency to change.

Architectures shape the cost of useful encounters. Dense layers allow learned dependencies; convolution makes spatial reuse inexpensive; attention enables content-dependent exchange [6, 17, 18]. The ecological design question concerns which arrangements can be discovered, retained, and revised through these mechanisms. In a transformer, one possibility is to give

recurrently useful coalitions persistent, addressable representations. Their value would depend on what they enable across new tasks and on the costs they incur.

An *Occam Transformer* names this proposed separation of fast interaction from slower decisions about retention. Begin with a bounded bank and records of recurring arrangements. Test whether making a group reusable improves upon invoking its constituents separately. Additional shared state, memory, or autonomy must justify their costs. Permit dormancy, splitting, and revision as well as consolidation. A frequently used routine can be a shortcut; an infrequently used distinction can prevent a consequential failure.

Selection should therefore assess contributions across declared contexts, including counterexamples to each candidate's habitual use. Confidence in a hypothesis, computation allocated to it, and authority over an action are distinct quantities. An unlikely possibility can deserve attention because investigating it would expose an expensive mistake. An efficient ecology preserves some alternatives and redundancy for that reason. Its practical economy concerns the cost of reliable performance, including detection and recovery when expectations fail.

Consolidation should carry the conditions that made a coordination useful. A procedure for moving attached objects needs to preserve the distinction between attachment and proximity. A disclosure routine needs a way to consult changed permissions. Distillation can make an acquired behaviour cheaper to reproduce [19]; tests of new bindings and corrections must establish what survived. Reuse is most valuable when later evidence can still reach the dependency it calls into question.

This is where generality and alignment meet. Greater reach makes a sound abstraction more useful and a brittle one more consequential. A system can also coordinate effectively around a harmful objective. The objectives, authorities, and limits governing an application therefore require independent justification. Existing approaches such as Constitutional AI already use principles, critique, and revision in training [20]. The ecological contribution would be to show which internal relationships change and whether those changes explain a benefit under unfamiliar conditions.

An aligned organisation must remain responsive to the right differences. Prestige should leave a permission unchanged; an authenticated amendment should alter it. An unresolved ambiguity may warrant another observation. Designing for these distinctions favours explicit tests of evidence use, accessible dependencies, and recovery over an undifferentiated demand for consistency. The same questions apply to a team whose shared account survives member turnover, or to a research programme whose familiar methods encounter an awkward result. Human consent and accountability remain constraints on interventions involving people; attributing a larger mind confers no authority over its members.

The useful map leads to a decision: which interaction to enable, which distinction to preserve, which influence to interrupt. Its success would make an organisation easier to improve and its failures easier to diagnose. That is the practical claim on which this ecological account should stand.

A thought can spend years on a page. Its next reader may supply a missing distinction, carry it into an experiment, or give an old error another generation. Everything depends on what happens when the result returns. Minds grow through achievements they can use again. Their ability to discover and repair mistakes depends on keeping those achievements within reach of experience.

To spot a mind, follow the work. To build a better one, make its useful discoveries easier to carry forward and its mistakes easier to reach.

References

- [1] E. Hutchins. *Cognition in the Wild*. MIT Press, 1995. doi:10.7551/mitpress/1881.001.0001.
- [2] A. Clark and D. J. Chalmers. The extended mind. *Analysis*, 58(1):7–19, 1998. doi:10.1093/analys/58.1.7.
- [3] S. C. Pratt, E. B. Mallon, D. J. T. Sumpter, and N. R. Franks. Quorum sensing, recruitment, and collective decision-making during colony emigration by the ant *Leptothorax albipennis*. *Behavioral Ecology and Sociobiology*, 52:117–127, 2002. doi:10.1007/s00265-002-0487-x.
- [4] H. K. Inagaki, L. Fontolan, S. Romani, and K. Svoboda. Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature*, 566:212–217, 2019. doi:10.1038/s41586-019-0919-7.
- [5] F. A. Hayek. The use of knowledge in society. *American Economic Review*, 35(4):519–530, 1945. JSTOR:1809376.
- [6] A. Vaswani et al. Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 2017. arXiv:1706.03762.
- [7] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafran, K. Narasimhan, and Y. Cao. ReAct: Synergizing reasoning and acting in language models. *International Conference on Learning Representations*, 2023. arXiv:2210.03629.
- [8] S. Blackmore. *The Meme Machine*. Oxford University Press, 1999. Author’s synopsis.
- [9] S. Spratley. *On the Use of Progressive Matrix Problems in Understanding Abstraction and Generalisation in Vision Systems*. PhD thesis, The University of Melbourne, 2024. Especially Chapter 2, §§2.2–2.3, pp. 14–21.
- [10] R. A. Watson and J. B. Pollack. A computational model of symbiotic composition in evolutionary transitions. *BioSystems*, 69(2–3):187–209, 2003. doi:10.1016/S0303-2647(02)00135-1.
- [11] K. Ellis et al. DreamCoder: Bootstrapping inductive program synthesis with wake-sleep library learning. *Proceedings of PLDI*, pp. 835–850, 2021. doi:10.1145/3453483.3454080.
- [12] C. Lu, J. Gallagher, J. Michala, K. Fish, and J. Lindsey. The Assistant Axis: Situating and stabilizing the default persona of language models. Preprint, 2026. arXiv:2601.10387.
- [13] Y. Chen et al. Reasoning models don’t always say what they think. Preprint, 2025. arXiv:2505.05410.
- [14] A. Geiger, H. Lu, T. Icard, and C. Potts. Causal abstractions of neural networks. *Advances in Neural Information Processing Systems*, 34, 2021. arXiv:2106.02997.
- [15] F. Zhang and N. Nanda. Towards best practices of activation patching in language models: Metrics and methods. *International Conference on Learning Representations*, 2024. arXiv:2309.16042.
- [16] F. Chollet. On the measure of intelligence. Preprint, 2019. arXiv:1911.01547.
- [17] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning representations by back-propagating errors. *Nature*, 323:533–536, 1986. doi:10.1038/323533a0.
- [18] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi:10.1109/5.726791.
- [19] G. Hinton, O. Vinyals, and J. Dean. Distilling the knowledge in a neural network. Preprint, 2015. arXiv:1503.02531.
- [20] Y. Bai et al. Constitutional AI: Harmlessness from AI feedback. Preprint, 2022. arXiv:2212.08073.

Provenance and status. This standalone position paper revises *When Ecologies Become Minds* and draws on the landscape and causal-cartography sections of Steven Spratley’s *Ecologies of Ecologies, Minds of Minds: On the Integration and Growth of General Intelligence* (working manuscripts, 2026). The cited ant, neural, and machine-learning studies provide empirical precedents. The integration criterion, memplex protominds, ecological interpretations, and proposed workshop and archive studies are research proposals; no new experiments are reported. Chapter 2 of the author’s doctoral thesis supplies the conceptual-boundary and concept-vampire material. Language models assisted drafting and reference checking.