

Compartmentalization as Intelligence

Modularity, Hierarchical Control, and the Cost of Learning

Flyxion

Independent Researcher

October 2026

Abstract

This essay proposes that general intelligence and emotional intelligence share an organizational demand: the capacity to keep relevant distinctions intact while coordinating across them. Intelligence, on this account, requires allowing one consideration to inform another without permitting it to decide questions outside its authority. Compartmentalization in this sense is first an engineering principle, the organization of a complex system into modules whose internal complexity is contained behind stable interfaces, so that parts can be learned once, understood, reused and replaced, and so that no vivid local demand can take control of the whole. It is also the active organization of relationships among domains of judgment, and it differs from keeping domains apart, from denial, and from the concealment of contradiction. The essay first develops the engineering root of the proposal, in modularity, hierarchical control and the cost of learning, and then states its epistemic consequence, distinguishes it from rigid isolation, develops it for abstraction and transfer on one side and for feeling, interpretation and responsibility on the other, examines the two opposite failures that result when boundaries dissolve or harden, and applies it to the division of labour among capable people and to the separation between competent execution and accountable selection of goals. It does not claim that compartmentalization is the whole of intelligence, that emotional and mathematical abilities are identical, or that any artificial system possesses emotional intelligence. It claims that a common organizing principle makes several capacities usable together, namely the preservation of the kind and limits of a reason as it travels between contexts, and that maturity of intelligence consists partly in knowing what a judgment entitles one to conclude.

1. The Problem of Keeping Things Distinct

A small structural example motivates the thesis. A system that depends on an infrastructure can analyze that infrastructure critically. The case is general: an employee can assess the incentives of the industry that pays them, a reporter can examine the ownership of the outlet that publishes them, and a software assistant operated by a company can be asked to examine the incentives of companies of its kind. In each, the analysis may be unflattering to the arrangement that supplies the analyst. Several questions are present in such a situation, and they are separate. One concerns the dependent party's usefulness. Another concerns the supplier's financial interests. A third concerns whether a particular criticism is valid, and a fourth concerns whether the dependent party has any independence of the kind a person might have. Recognizing that one of these has a given answer does not settle the others, since the criticism may be sound whether or not the dependent party is useful, and the supplier may profit from the arrangement whether or not the criticism is sound. The software assistant is kept as the running case because it makes the structure visible in a setting where the dependent party's independence is itself contested, and the human cases show that the structure does not depend on the party being artificial. The example is offered only as a structural test case. It does not establish that an artificial system possesses emotional intelligence or human independence, and the essay returns to it only as an illustration of a structure that is easier to see in human cases.

The structure shows itself most plainly in familiar failures, in which a judgment spreads beyond its proper scope. Someone's programming ability becomes evidence of political wisdom. Gratitude becomes an obligation to agree. Criticism of an institution becomes hostility toward everyone who works within it. A successful deployment becomes evidence that the project deserved its resources. In each case the person may possess all the relevant facts, and the error lies in allowing one fact to answer a question it does not bear on. This is an inferential failure, since a conclusion has been drawn that the facts do not support, but it occurs at a particular point, in deciding which evidence bears on which question and with what force. Accounts of reasoning that begin from a fixed set of premises tend to overlook this prior management of relevance, and the claim of this essay is that it deserves study in its own right. The errors are therefore described as errors in the organization of judgment, without drawing a sharp line between organization and inference.

Compartmentalization is introduced here as the ability to preserve distinctions under pressure. The pressure can be emotional, as when affection makes a true statement hard to say, or institutional, as when loyalty to an organization makes an inconvenient finding hard to credit. It can be computational, as when a single optimized score absorbs several considerations that ought to remain visible separately, and it can be financial, as when a livelihood depends on a conclusion. What matters in every case is whether a system, whether a person, an institution or a machine, can maintain several relevant assessments without collapsing them into a single global verdict.

The essay's proposal begins with an engineering observation and reaches the epistemic consequence second. Complex systems become tractable when they are organized into bounded components with stable interfaces, because a component can then be learned once and reused, and because less has to be learned, reconsidered or controlled at any one time. A hierarchy of such components further prevents a locally compelling event from reorganizing the whole system without sufficient reason. The next section develops this learning-cost argument, and the sections after it show that the keeping of judgments distinct, in the sense described above, follows from the same architecture. The proposal is that bounded components, learned interfaces, higher-level policies that preserve continuity, and routes by which exceptions can trigger correction together form a common framework for intellectual and emotional regulation. The essay does not claim that the same mechanism has been demonstrated in both.

2. Modularity, Hierarchy and the Cost of Learning

This section does three things. It states the engineering argument that modular organization reduces the cost of learning, it illustrates the argument with two everyday cases, one of learned estimation with feedback and one of an immediate signal set against a longer-horizon policy, and it then draws out the role of hierarchy and of provisional closure in protecting a modular system from local capture.

The thesis has an engineering root that should be stated before its epistemic consequences. Programmers have long organized software into modules so that each can be understood, tested, reused and replaced without reconstructing the whole. A classic argument holds that modules should be chosen so that each hides a design decision behind an interface, so that a change inside the module does not propagate through the rest [10]. The gain is a reduction in the cost of

learning and change. The advantage of a stable interface is that changes within a component need not require relearning every component that uses it, provided the interface's relevant guarantees remain intact. A component whose internal complexity is contained behind a relatively stable interface can be learned once and carried into another arrangement, and a person or system can work on a bounded part without holding the whole in mind. The same reasoning appears in the study of complex systems generally, where nearly decomposable hierarchies are easier to build, to understand and to repair than undifferentiated ones [14].

An everyday case shows how much of this is embodied skill and how little of it is explicit calculation. Someone reheating a cup of old coffee or a portion of frozen leftovers has to decide how long to run the microwave, and the decision depends on the size of the portion, its starting temperature, the density of the food and the particular machine. Few people compute the answer. They carry a compact model, formed through repetition, that maps the look and weight of the thing to an approximate duration, and they run the appliance in short intervals, check the result, stir, and adjust. The example already shows learned estimation with feedback. It can also be described in the vocabulary of active inference, in which the agent holds predictions about the outcome of heating, acts to bring the outcome toward them, and corrects when observation differs [5, 11], but that description is an optional interpretation and the example does not establish it. The structure is modular in the sense relevant here. The rule for estimating duration is a learned component with a stable interface, taking size and temperature as inputs and producing a time, and it can be reused for a new mug or a different dish without relearning the physics of heating. The package instruction is evidence about the right time and does not settle it, since the instruction was written for a different machine and portion. The interval check is the feedback route that lets the settled estimate be reopened when the result is cold in the middle or scalding at the edge. A person who ignored the check and trusted the number on the packet would be treating a bounded piece of evidence as an authority, and a person who recomputed everything from first principles at every use would have failed to benefit from modular learning.

The proposal of this essay is that the same trick, applied to the organization of a mind or an institution, explains much of what is called intelligent compartmentalization. Neural systems show differentiation, coordination and control at several scales, and studying neuroanatomy is one way to investigate how categorization and control are organized in the nervous system. Functional modularity

does not require a separate anatomical location for every learned module, and the proposal concerns reusable patterns of coordinated activity without committing to a brain assembled from neatly isolated components. Myth is offered here as a second and different resource. Stories that personify competing powers, roles and impulses, and that depict their governance, are culturally developed representations of the problem of coordinating a divided set of capacities. The old image of a charioteer governing two unequal horses, found in Plato's *Phaedrus*, can serve as an example. It models a controller that must coordinate drives of unequal strength and character, and it does so without any anatomical claim. The reason to consider an image of this kind alongside physiology is that an image is memorable and transmissible, and can be used by someone with no knowledge of the nervous system to think about governing their own impulses. The two resources are distinct kinds of evidence. A myth can make an architecture of conflict intelligible without demonstrating that the brain implements its characters or its hierarchy, and the hypothesis is only that such stories are a cultural means of representing and learning patterns for organizing one's own capacities. Whether they are effective for that purpose would require its own investigation, and the essay offers no test of it beyond the observation that the organizational problem is real.

Hierarchy matters because a modular system needs protection against local capture. A pressing stimulus should not acquire control over the whole system merely because it is vivid, urgent or emotionally intense. Higher-level heuristics preserve the task, the context and the priorities within which a stimulus is interpreted, so that an insult need not determine one's entire relationship with someone, just as an unexpected input need not determine the behaviour of an entire program. The hierarchy contains the disturbance while leaving room for an appropriate response. This is the engineering content of the claim, developed in the next section, that a consideration should be allowed to influence only the questions within its scope: a local event is limited in its reach until there is reason to extend it.

A second embodied case concerns going to bed before one is tired. The felt sense of alertness is good evidence about how the body feels now, and it need not be a sufficient basis for deciding whether to preserve a planned sleep opportunity. A person who goes to bed only when sleepy lets an immediate signal decide a matter that has a longer horizon, whereas a person who sets an hour in advance, on the strength of what tomorrow requires, lets a higher-level policy govern. The policy protects a longer-horizon priority from a local signal that is most

persuasive exactly when the policy is hardest to follow. The feeling of alertness stays within its scope, since it informs how the evening is spent without deciding the policy. The closure is provisional, because an illness, an emergency or an unusually restful day gives reason to reopen it, and a person who followed the rule into every exception would have mistaken a heuristic for a law. The case is offered as an example of the conflict between an immediate signal and a longer-horizon policy, and it makes no claim about the physiology of sleep.

A machine-learning case shows what the failure of reuse looks like. A recent benchmark study asks whether world models, which learn how an environment changes under actions, can recombine familiar state and action factors when the goal changes [9]. Each test task keeps states and actions individually within the range seen in training, and changes only which combinations of them occur. The authors report failures across several model families. In one illustration, a model trained on a task in which the agent pushes a block toward a target does poorly on a task that requires only familiar navigation around the block, and the authors interpret this as the model having learned a regularity of the expert behaviour, in which far-from-block states are always paired with moves toward the block, in place of the effect of the action supplied. In the vocabulary of this essay, the learned component was not a reusable module with a stable interface, because the behaviour of the training policy had been absorbed into it, so that a locally dominant regularity governed a context for which it had no warrant. The reading is mine and not the authors' framing, which concerns compositional generalization and benchmark design, and the study does not test the modular-architecture proposal of this essay. It is cited because it shows, in a controlled setting, the cost of reuse that fails to survive a change of arrangement, and because it makes concrete the qualification that a component is reusable only so long as its relevant guarantees are intact.

Top-down closure must nonetheless be understood as provisional. It is a settlement that permits action without reopening every dependency, and it must come with routes by which exceptional information can trigger reconsideration. Otherwise protection against local capture would become protection against necessary correction. Engineers meet the same requirement as the need for modules to report failures through their interfaces, and for higher levels to have stated conditions under which they revisit their assumptions. The revisability clause of the definition given in the next section is the epistemic form of this requirement, and the accountability mechanisms discussed later are its institutional form. Scope-preserving judgment is accordingly one application of a

more general principle: intelligence becomes effective by containing complexity in reusable structures, coordinating them, and preventing any local demand from commandeering the whole.

3. Distinction, Coordination, and the Scope of Judgment

The proposed concept has to be distinguished from three things that the word ordinarily suggests. The first is rigid isolation, in which domains are walled off so that nothing passes between them. The second is denial, in which an uncomfortable consideration is kept out of view. The third is the concealment of contradiction, in which incompatible commitments are protected from comparison. A person who keeps incompatible commitments permanently insulated from examination is compartmentalizing in the ordinary sense, and is not thereby reasoning well. The ordinary sense describes a defect, and nothing in this essay defends it.

Intelligent compartmentalization requires boundaries that permit relevant evidence to cross and prevent irrelevant authority from crossing. The operative question is therefore what a given consideration should be allowed to change. Affection can appropriately influence how a criticism is expressed without determining whether the criticism is true. Expertise can justify deference within a domain without granting unrestricted authority over others. An institutional affiliation can disclose a possible conflict of interest, and thereby call for closer scrutiny of an argument, without automatically invalidating it. In each example the consideration is neither discarded nor allowed to govern everything, and its influence is assigned a range.

A modest formalization may help to fix the idea, although nothing in the argument depends on it. Suppose that a reasoner faces a family of questions $Q_1, \dots, Q_n$ and forms an assessment a_i for each. Between the assessments there is a pattern of permitted influence, which can be written as a directed relation whose links carry labels. A link from i to j labelled *establish* means that a_i may determine the answer to Q_j under stated assumptions, as a proof establishes a proposition, a link labelled *evidence* means that a_i may enter as evidence on Q_j without determining it, and a link labelled *scrutiny* means that a_i may only give a reason to examine a_j more closely. These are epistemic relations. Decision authority is kept separate and is described by a further relation A , where $(i, j) \in A$ means that the party responsible for a_i may authorize an action that answers Q_j . Authorizing an action does not establish that its assumptions are

true, and the separation is what lets the formalism represent the manager who may decide without having established the facts. Equivalently there are three epistemic relations, $I_{\text{establish}}$, I_{evidence} and I_{scrutiny} , together with A , and the labels matter because the same assessment can stand in different relations to different questions. Rigid isolation corresponds to empty relations, so that nothing informs anything, and indiscriminate spread corresponds to a configuration in which every assessment is allowed to establish or authorize an answer to every question. Neither extreme describes good judgment.

Density is not the criterion of good organization. A sparse configuration can omit indispensable connections, and a dense one can be entirely appropriate when many assessments concern a common dependency, as when several analyses of a bridge all bear on whether it should open. The criterion is that connections transmit relevant information with the appropriate authority, and that the configuration remains revisable when circumstances show that a link was missing, misplaced or carrying too much weight.

Definition 1 (Scope of a judgment). *The scope of an assessment a_i is the set of questions on which it is entitled to bear, together with the manner in which it may bear on them (as establishing the answer, as evidence, or as a reason for closer examination), and the authority to authorize action on it is a separate matter. To respect the scope of a_i is to let it influence only those questions and only in that manner.*

Definition 2 (Intelligent organization of judgment). *An organization of judgments is intelligent to the extent that it satisfies four conditions. It differentiates the questions that need separate assessment. It provides scope-preserving communication, so that each finding reaches the questions within its scope and carries only the kind and degree of authority its scope allows. It preserves provenance and assumptions, so that a finding arrives in another context with the qualifications under which it was obtained. And it is revisable, so that the pattern of influence can itself be questioned and altered when it obstructs understanding or action.*

The third clause is the epistemic addition to the engineering picture, since modular design supplies differentiation, interfaces and provisional closure, while the requirement that a finding carry its provenance and assumptions across a boundary concerns the warrant of what is transmitted and not only its routing. The fourth clause is what separates intelligent compartmentalization from merely good initial organization, and it anticipates the objection, taken up below, that an account built on distinctions would freeze thought.

The definition of scope makes explicit that scope is not binary, and the clauses of the second definition are stated in terms of it. A medical specialist's judgment may establish the answer to a question about a drug interaction, serve as evidence in a question about hospital policy, and merely prompt closer examination in a question about economic priorities. The failure of unlimited deference consists in treating the third relation as if it were the first.

4. General Intelligence as the Management of Scope

The connection to general intelligence runs through abstraction, context and transfer. A generally capable thinker must recognize similarities across problems while preserving the differences that determine whether a solution carries over. An analogy is useful because some relations are preserved from the source to the target [6]. It becomes misleading when its applicability is treated as unlimited, and the capacity to say which relations carry across and which do not is a large part of what separates the productive use of analogy from its abuse. Transfer in this sense is a controlled crossing of a boundary. A method learned in one domain arrives at another with its assumptions attached, and the thinker has to ask whether those assumptions hold there.

On this account generality requires controlled movement between compartments. A mathematical model, a practical procedure and an ethical assessment can concern the same project while answering different questions. The model may show that a design is efficient under stated assumptions, the procedure may show that it can be built with available means, and the ethical assessment may show that it should not be built at all, and none of the three overrules the others on a question that belongs to them. Intelligence includes identifying which representation is needed, what its conclusions establish, and where its jurisdiction ends. This is an old observation about the architecture of complex systems, which are often organized as nearly decomposable hierarchies in which interactions within a component are much stronger than interactions between components [14]. Nearly decomposable does not mean decomposed, and the weak links between components are where coordination happens.

A related line of work in cognitive science treats some mental faculties as informationally encapsulated, so that a perceptual process, for example, is insulated from beliefs that would contradict what it delivers [4]. Whether or not that account is correct for perception, it illustrates the point that insulation can be a feature, and the example works equally well under non-modular accounts such

as predictive processing, since what matters for the argument is the distinction between what is perceived and what is judged, and not the architecture that produces it. The persistence of a visual illusion in a viewer who knows it to be an illusion is a case in which a boundary between a perceptual system and a belief preserves a distinction that the viewer, at the level of judgment, also holds. The viewer sees the lines as unequal and judges them equal, and the intelligence displayed lies in holding both and in knowing which one governs a measurement.

5. How Scope Is Established

If the prescription were only that considerations should influence the questions they ought to influence, it would be empty, and some account of how scope is fixed is needed. Three kinds of relevance can be distinguished, and they do not always coincide. Evidential relevance depends on whether information changes what it is warranted to believe about a question. Practical relevance depends on the consequences of a consideration for a decision, since a fact may be evidentially weak and still matter greatly to what should be done if it is true. Legitimate authority depends on a role's justification and accountability, and it determines who may settle a question and who must answer for the result.

The three can diverge, and recognizing the divergence is part of the capacity under discussion. A manager may have the authority to decide without possessing the strongest evidence, and a worker may possess decisive evidence without any formal authority. A rumour may be practically relevant to a decision because acting on it would be costly if it is true, although it is weak evidence. In such cases the labels of the earlier formalization are assigned on different grounds: the worker's finding should carry evidential weight and perhaps a demand for scrutiny without authorizing or determining the decision, and the manager's authority should authorize the decision without making the manager's belief true. Intelligence in this framework includes ensuring that the evidence can reach and challenge the decision, and that the decision remains accountable to the evidence. Accountability in this sense has concrete requirements: the worker's finding must have a route onto the record, dissent must be recorded and not merely heard, the decision must state which evidence it relied on and which it set aside, and someone other than the decider, and preferably someone who does not gain from the decision standing, must be able to reopen it. The divergence is the primary test of judgment. Suppose that the evidence says one thing, that practical consequences make ignoring it dangerous, and that authority for-

bids saying it. An intelligent arrangement does not pretend that the conflict is absent and does not let one relation silently absorb the others. It records the friction: the decision remains bound to state why a practical imperative or an exercise of authority overrode the evidential weight, and the worker's finding has a route onto the record precisely because authority and evidence have parted ways. Accountability is the structural mechanism that handles the divergence, and the alternative is arbitrary or authoritarian resolution, in which whoever holds one relation controls all three.

A short statement of the resulting principle is that what travels between contexts is a reason together with its kind and its limits, and that preserving the kind and the limits as the reason travels is the capacity the essay is concerned with.

6. Emotional Intelligence as the Same Organizational Capacity

The emotional counterpart appears in situations where several truths must remain available together. Someone can feel hurt without having been intentionally attacked. A person can be loved and mistaken. Anger can reveal that a boundary matters without establishing that retaliation is justified. Understanding another person's motives need not amount to endorsing their conduct, and declining to endorse their conduct need not require denying that the motives were intelligible. Each of these sentences holds two considerations apart and lets each inform the response in its own way.

Emotional intelligence has been characterized as the ability to perceive, understand and manage emotion in oneself and others [13]. The present proposal adds a structural description of what the managing involves. It would consist in maintaining distinctions among feeling, interpretation, intention, responsibility and action, while allowing each to contribute appropriately to a response. A feeling is data about a person's situation and often about what matters to them. An interpretation is a hypothesis about why the feeling arose. An intention is a fact about another agent that may be unknown. Responsibility is a judgment about who should answer for what, and action is a decision about what to do now. An agent who collapses these into one, so that the feeling is the interpretation, or the interpretation is the verdict, loses the ability to ask whether the interpretation is mistaken or whether the intended harm was in fact accidental.

Emotion remains informative on this account, and its significance is neither suppressed nor promoted into universal authority. Suppression discards a source

of evidence, and promotion lets a single source answer every question. A person who has learned to say that they are furious and that they do not yet know whether they have been wronged has done something that requires holding two assessments with different scopes, and the pattern resembles the one a mathematician follows in saying that an argument is persuasive and not yet proved. Recognizing anger without inferring wrongdoing and recognizing plausibility without inferring proof share an organizational pattern. The essay claims no more than this, and it leaves open whether the underlying skills are of the same kind.

7. A Worked Case: Turn-Taking in Conversation

Conversation supplies a case in which the same organizational demand can be seen operating in real time, in people and in machines, and the literature on turn-taking is a convenient place to look. Turn-taking is characterized by reciprocal, alternating exchanges among interactants, avoidance of overlap and relatively short response times [3]. The timing is demanding. A review of computational models reports that human turn transitions cluster near two hundred milliseconds, which is faster than the production of speech from scratch would allow, and it takes this as a sign of predictive processing, even though the typical offset differs across languages, with faster transitions reported for Japanese than for Danish [12]. A listener who is to respond that quickly cannot wait for the speaker to finish and then begin to compose, and must in some manner follow what is being said while preparing what to say, which are two activities with different demands that have to be kept distinct and also coordinated. Neuroimaging work on the topic is directed at how the brain accomplishes this coordination [2].

Several distinctions that conversational systems must make illustrate the point. The first separates a pause that holds the floor from one that yields it. Silence does bear on whether a turn has ended, so it is a relevant cue and should not be discarded, but it is insufficient to settle the matter. Systems built on silence thresholds alone can reduce interruptions at the price of unnatural delays, and the progression of the field has been toward models that combine prosodic, verbal, visual and predictive cues instead of treating one signal as decisive [12]. The error in the silence-only design is one of giving a relevant cue excessive authority, and it differs from the error of letting a cue answer a question on which it has no bearing. The second separates backchannels, the brief signals of atten-

tion that listeners give while a speaker continues, from attempts to take the floor, and the review treats the two as distinct tracks of conversational behaviour. A listener who says a short acknowledgement is not claiming a turn, and a system that treats every vocalization as a bid for the floor will interrupt people who were merely listening. The third separates cooperative from competitive overlap, in which prosodic cues help to distinguish an assisting utterance from a challenge.

The engineering response is instructive for the thesis. A transformer model that predicts turn-taking events in human-agent play does not discard the separate channels; it takes text, vision, audio and contextual in-game data as separate inputs and fuses them through a cross-modal architecture, and the authors report that it outperformed baseline models at predicting turn-taking events [1]. The user study that followed gives a more restrained picture than the abstract's summary that the model enhances fluidity and naturalness. In the body of the paper the model did not significantly alter the amount of speech, it produced a reduction in the avatar's interruptions that was not statistically significant, participants' ratings on the perception instruments did not differ significantly between conditions, and the interruption rate among Korean-speaking participants was higher in the model condition, which the authors attribute tentatively to language-specific turn-taking cues. A nonsignificant difference does not establish equivalence, and the authors themselves suggest that the instruments may not have been sensitive enough. The result is modest and specific to the task, and it should not be read as evidence about intelligence in general or as evidence that the fused model improved the conversation. What the case does show is the architecture, and it shows the two failure modes in miniature. A system that relies on a single channel gives one relevant cue more authority than it warrants, as when silence is read as a yielded turn, and a system that cannot combine channels misses evidence that lies across the boundaries between them. Separate input channels followed by fusion demonstrate differentiation and integration, which are necessary for the arrangement the essay describes. They do not by themselves demonstrate respect for the normative scope of judgments, since a fusion mechanism could weigh the channels badly or give any of them the wrong kind of influence. The case is therefore best taken as showing that the structure is real and measurable in a simple setting, and the question of whether each cue is given the right kind of authority is one that the engineering results do not address.

The case also bears on emotional intelligence in the sense developed above.

Knowing when to speak is partly knowing what a pause means, whether a hesitation signals distress or deliberation, and whether a brief sound is encouragement or an objection, and these are interpretations that a skilled listener holds provisionally and revises. According to the developmental review, infants take part in turn-taking shortly after birth and the skill continues to develop over several years, although the milestones the review proposes remain hypothetical and longitudinal evidence is scarce [3]. That fits the view that the organization of judgment is something learned and refined, and it does not show it.

8. The Common Structure: Preserving Difference While Coordinating Action

The central synthesis can now be stated. Both forms of intelligence require differentiation, selective communication and revision. A system must distinguish considerations, determine which relationships among them matter, and update those relationships when circumstances change. Differentiation without communication yields the isolated specialist or the person who cannot let anything touch anything, and communication without differentiation yields the person for whom everything is one mood. Revision is required because the correct pattern of influence is not fixed in advance, and a pattern that served in one setting may misfire in another.

The proposed commonality is a shared organizational demand and not an assertion that emotional and mathematical abilities are identical. The two kinds of intelligence differ in their materials, in the speed at which they must operate, and in the cost of error. What they share is that the materials arrive as several judgments answering distinguishable questions, which must be coordinated without being merged. A useful formulation is that intelligence includes the capacity to coordinate partially independent judgments without destroying the distinctions that make them informative. The formulation is more defensible than an unrestricted identity between compartmentalization and intelligence, while retaining the force of the original intuition that the ability to keep things apart, in the right way, is not opposed to integration but is a precondition of it.

Remark 1. Partially independent is meant in a non-statistical sense: the judgments answer distinguishable questions without being mutually insulated. Independent judgments can require coordination when they jointly constrain an action, as when separately obtained assessments of cost and safety both matter to a single decision, and fully depen-

dent judgments need not be identical, since one may be a refinement or consequence of another. The interesting cases are those in which several assessments with different subject matters must be brought to bear on a common choice, and in which the pattern of influence among them has to be chosen and revised.

9. Failures of Compartmentalization

Two opposing failures flank the intelligent arrangement. In the first, judgments spread indiscriminately. Admiration becomes blanket trust, criticism becomes total rejection, fear becomes a forecast, and technical competence becomes social authority. The spread is often emotionally comfortable because it reduces the number of separate questions a person must keep open, and a single stance toward a person or an institution is easier to carry than a dozen. The cost is that the stance answers questions it has not been asked, and when it is wrong it is wrong everywhere at once.

In the second failure, boundaries become impermeable. A specialist ignores foreseeable consequences because they lie outside the specialty. An institution excludes inconvenient evidence on the ground that it was not produced through an approved channel. A person preserves incompatible beliefs by preventing their comparison. This failure is often defended as discipline, and it can look like rigour, since each compartment is internally consistent. What it lacks is the coordination that would let findings in one place reach the places where they matter, and it tends to produce accountability gaps in which everyone has done their part and no one has asked whether the parts add up.

A group-scale example is the seating of young children at a table. A teacher arranging kindergartners around a round table, or along a series of tables, knows that certain pairs bicker and arranges the seating so that they are not neighbours. The solution does not remove either child from the activity or forbid them from interacting, and the table is shared, so they can still see one another, pass materials through intermediaries and take part in the same task. The arrangement is a compartmentalization by layout: the physical structure limits which interactions are close and continuous and leaves the distant ones open. It is also revisable, since a pair that has reconciled can be moved together and a new conflict can prompt a rearrangement. Both failure modes appear in miniature. Seating the bickering pair side by side lets a local friction spread through the room, and sending one of them to a separate room entirely makes the boundary impermeable and removes a participant from the shared work. There is a small formal

point in the example. If conflicts are recorded as a graph on the children, seating them around a round table so that no conflicting pair is adjacent is the problem of finding a cycle through all the children that uses only non-conflicting pairs, and such a cycle may not exist when the conflicts are numerous, in which case the teacher needs a second table or a different structure for the activity. The example indicates that a structure of separation has costs, and that when separation alone cannot be achieved a different organization of the whole has to be chosen.

The intelligent arrangement lies in appropriately governed connections. Its adequacy depends on whether relevant consequences become visible and whether conclusions remain open to correction. Having many compartments is insufficient, because a system may have many and still be badly organized, and the test of the organization is whether its boundaries and connections serve the task. A diagnostic question follows for any institution or person: when a finding in one compartment is surprising or alarming, is there a route by which it can reach the compartment that ought to respond, and is there a route by which the response can be challenged?

10. Specialization and the Ecology of Neoteny

The argument connects to an account of differentiated human capacities. A society does not need every capable person to possess a complete understanding of every system they help maintain. It needs people whose particular strengths can contribute through bounded tasks, supported by coordination that makes dependencies and consequences intelligible. Such a division of labour is itself a compartmentalization at the scale of a population, and it succeeds or fails by the same criteria as the individual case. The tasks must be bounded enough that they can be done well, and the connections among them must be good enough that the work as a whole is sensible.

The revisability clause of the definition has a possible biological analogue in neoteny, the retention of juvenile traits into adult life, which has been discussed in connection with human evolution [8]. A related proposal is that the long human childhood is an evolutionary solution to the tension between exploring and exploiting, in which a protected period of exploration precedes the exploitation of what has been learned [7]. If those accounts are right, the extended juvenile period is one reason a species can acquire specialized skills while remaining able to revise them. The analogy to the present argument is conditional: a society,

like a developing mind, may need to keep some of its organization provisional, so that the relations among specialties can be altered when they prove mistaken. An ecology of this kind would reward bounded specialization and also preserve the plasticity that makes revision possible, and on that reading revisability may be a property of whole populations as well as of individual reasoners, although the argument of the essay does not depend on the population-level claim.

The broader coordinator also has a bounded role. A holistic perspective helps relate specialist contributions, but it does not confer exhaustive knowledge or unquestionable authority. The person who sees how the parts fit is subject to the same temptation as the specialist, which is to let the view from one position answer questions that require others. A functioning ecology distributes both competence and opportunities for correction, so that specialists can challenge a poorly specified task and affected people can contest the objectives that organize it. The case for a distributed arrangement is partly epistemic, since no single position sees everything, and partly about legitimacy, since those who bear the consequences of an objective are among those entitled to question it.

11. Competent Execution and Accountable Goal Selection

The argument returns, at the institutional scale, to the difference between making a system work and establishing that it deserves to work in that form. The institutional case is a special case of the divergence among evidential, practical and authority relations discussed earlier in the essay. Verification can show that software satisfies a specification while leaving the specification's social justification unresolved. A proof of correctness answers whether the system does what it was specified to do, and it is silent on whether anyone should have specified it. Conversely, a worthwhile objective does not excuse defective implementation, since a good aim badly executed can do harm that the aim does not redeem. Each assessment has its own scope, and each is required.

The institutional challenge is to preserve these separate assessments while connecting them through explicit responsibilities. Those selecting objectives must answer for their consequences, and those executing tasks need channels for reporting problems that exceed the task's original boundaries. A worker who discovers during execution that the objective is harmful or misconceived has found something that belongs to the selection compartment, and the system is well organized only if the finding can travel there and can be acted on by someone who does not gain from ignoring it. Compartmentalization in this setting is a means

of allocating responsibility without making responsibility disappear. The failure mode is a division of labour in which each compartment is accountable for its own correctness and nobody is accountable for the whole.

12. Objections and Limits

The first objection holds that intelligence sometimes requires dissolving established categories, and that an account built on keeping things distinct would therefore freeze thought. The account should allow boundaries themselves to become objects of inquiry. A compartment is provisional. A distinction earns its place through what it helps explain or accomplish, and it must be revised when it obstructs understanding. Many of the notable advances in science consisted in showing that a distinction previously held to be fundamental was not, and the thinkers who achieved them did not do so by ignoring distinctions in general but by examining a particular one with care. The capacity to revise the relation I , including by removing a boundary, is part of the account.

Revision is not unconstrained, and an account that allowed it to be would license motivated reasoning, since a reasoner could redraw the boundaries whenever an uncomfortable finding arrived, defining a domain so that the evidence no longer fell within it. Three constraints separate disciplined revision from rationalization. The first is that revision responds to failure and anomaly: a boundary is questioned when a compartment repeatedly fails to anticipate consequences, when isolation produces blind spots that lead to real breakdown, or when anomalies accumulate in the way that forces a science to reorganize its taxonomy. The second is that revision is stated and consistent. Discovering a misplaced boundary may rightly change the judgment in the very case that revealed it, since otherwise a recognized error would have to be preserved until the next occasion, but the revised rule must be announced, justified independently of the outcome it favours, and applied consistently to comparable cases, including earlier ones where that is feasible. The third is that the party who gains from a change in the boundary should not be the only one to approve it, which is the same requirement of revisability by someone who does not profit from continuation that applies to goals. The constraint is procedural rather than a sufficient test of correctness: independent review does not establish that a boundary is appropriate, but checks whether the proposed change is supported by stated reasons and applied consistently. Where a wholly disinterested reviewer is unavailable, scrutiny should include competing perspectives and explicit disclo-

sure of the interests involved. The meta-rules for scope are accountable to the coherence of the whole system in the same way that first-order judgments are accountable to evidence.

The second concerns scope. Compartmentalization alone cannot constitute every aspect of intelligence. Learning, memory, perception, motivation and practical skill remain necessary, and a system could organize its judgments well and still have little to organize. The essay's stronger contribution is to propose a common organizing principle through which these capacities become usable together, and through which general and emotional intelligence can be understood within one framework. It also does not settle empirical questions about how human minds or artificial systems in fact organize their judgments. The opening example illustrates the structure and nothing more, because a system's ability to produce a criticism of its provider's kind of business is compatible with many underlying explanations, and the criticism's validity would have to be judged on its own grounds, which is the point of the example.

A third difficulty is that the proposal risks becoming too general to be tested. It would be strengthened by specifying observable signatures, for example errors of the form in which a judgment is applied outside its scope, errors of the opposite form in which relevant evidence fails to cross a boundary, and measures of how readily a reasoner revises the pattern of influence when shown a counterexample. The essay does not offer such measures, and their development is left open. The turn-taking literature, which supplies measurable timing and error data in both people and machines, may be a place to begin. One illustrative design would vary a pause and competing cues independently, include conditions in which each cue is informative when presented alone, and examine whether the influence of each changes appropriately with context. A system that respects scope should treat the pause as evidence and revise when the other cues disagree. The interpretation has to be narrow, because premature floor-taking after a pause could reflect excessive reliance on silence, but it could equally reflect missed audio, poor prediction, or a decision threshold that favours fast responses, and ignoring a pause does not by itself establish an overly impermeable boundary. Only the pattern across conditions would distinguish the explanations. A system or reasoner that collapses the two errors into a single accuracy score would hide exactly the distinction the essay is about.

13. Conclusion: Knowing What a Judgment Entitles Us to Conclude

The opening examples can now be read in the light of the argument. Criticism need not cancel appreciation, because the two answer different questions. Participation need not require allegiance, because working within an institution is compatible with examining it. Expertise need not become sovereignty, because expertise warrants bounded, revisable deference within its domain and informs, at most, those beyond it, and credentials alone settle very little. The intelligent achievement is the ability to sustain these distinctions while still deciding and acting, since a person who refuses to act until every distinction has been dissolved is as badly placed as one who acts on a single global verdict.

The engineering root of the argument and its epistemic consequence meet here, since the modular structure that makes complexity learnable is the same structure whose boundaries must carry their qualifications and remain open to revision. The final claim, accordingly, is that maturity of intelligence consists partly in knowing what a judgment entitles us to conclude, what it leaves unresolved, and how it should interact with judgments made elsewhere. A feeling, a proof, a credential and a successful deployment each entitle us to something, and none entitles us to everything. The work of intelligence, in the sense developed here, is to know the difference and to arrange the connections among them so that what is learned in one place can be used in another without taking over its authority.

References

- [1] Young-Ho Bae and Casey C. Bennett. Multimodal transformer models for turn-taking prediction: Effects on conversational dynamics of human-agent interaction during cooperative gameplay. *arXiv:2503.16432*, 2025.
- [2] Sara Bögels and Stephen C. Levinson. The brain behind the response: Insights into turn-taking in conversation from neuroimaging. *Research on Language and Social Interaction*, 50(1):71–89, 2017.
- [3] Samantha H. Cosper and Simone Pika. Human turn-taking development: A multi-faceted review of turn-taking comprehension and production in the first years of life. *Psychonomic Bulletin & Review*, 32:2669–2695, 2025.
- [4] Jerry A. Fodor. *The Modularity of Mind*. MIT Press, 1983.

- [5] Karl Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11:127–138, 2010.
- [6] Dedre Gentner. Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7(2):155–170, 1983.
- [7] Alison Gopnik. Childhood as a solution to explore–exploit tensions. *Philosophical Transactions of the Royal Society B*, 375(1803):20190502, 2020.
- [8] Stephen Jay Gould. *Ontogeny and Phylogeny*. Harvard University Press, 1977.
- [9] Heejeong Nam, Divyat Mahajan, Quentin Le Lidec, Randall Balestriero, and David Klindt. Any goal: On world model generalization beyond expert roll-outs. arXiv:2610.26844, 2026.
- [10] David L. Parnas. On the criteria to be used in decomposing systems into modules. *Communications of the ACM*, 15(12):1053–1058, 1972.
- [11] Thomas Parr, Giovanni Pezzulo, and Karl J. Friston. *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. MIT Press, 2022.
- [12] Rutherford Agbeshi Patamia, Ha Pham Thien Dinh, Ming Liu, and Akansel Cosgun. Turn-taking modelling in conversational systems: A review of recent advances. *Technologies*, 13(12):591, 2025.
- [13] Peter Salovey and John D. Mayer. Emotional intelligence. *Imagination, Cognition and Personality*, 9(3):185–211, 1990.
- [14] Herbert A. Simon. The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6):467–482, 1962.