Tautai

How Organizations Understand

The loop that comes before deciding, and how to build it in

Version 0.1 · 17 November 2026

Krishan Mathis · tautai · krishan@tautai.net

What this paper is for

This paper explains an argument, not a product. It sets out why understanding, rather than decision speed or data volume, is the bottleneck of adaptation; what understanding is made of; where it is held in an organization; and what it takes to build the whole loop into an operating model. The book behind it is an open draft, and the paper says where the argument is strong and where it is not.

The occasion is an observation about management systems. Most of them were built on one assumption: that the situation can be established first, by analysis, research or a planning cycle, and that decision and action follow from it. That assumption holds where conditions are stable relative to the response cycle. It fails where the environment changes faster than it can be known, and where acting on the situation is itself what changes it.

In five sentences. An organization cannot know its environment in advance, so it has to produce understanding while acting. Understanding is made from limited attention, early cues and plausible stories, by people together with their artifacts, forums and structures; it is distributed, not a leader's insight. From understanding, experienced people mostly arrive at action without a separate act of choosing. The design task is to build the whole loop — perceiving, understanding, choosing, acting, learning — into the operating model, at every level that faces its own environment. Data does not say what it means, and more of it does not help, because information consumes attention.

What in this paper is which

Element	Status
The five activities and the loop	the book's argument, How Organizations Understand, 2026-09-25 draft — no thesis accepted by the field
The internal model, in its loose and strong senses	draft; the strong sense is argued for, not claimed as proved
The six design grids and the 90-day shape	the book's design, built from its chapters; no unit is shown running them
The claims	37 claim records, many assessed `insufficient-evidence`; the book says so itself
Publication	open draft 0.3.0 on tautai.net since 2026-10-12, unreviewed as a whole; 0.3.1 is locked in the vault and unpublished

Nothing here is a research finding of this paper's. It restates the book's argument, and the paper's job is to make the argument usable in a room without the book in it.

1. The problem: the situation cannot be established first

Analysis, research and the planning cycle all assume a situation that exists, waits to be observed, and can be described before anyone acts on it. In a stable environment that is a good assumption and the planning cycle is the right instrument. The trouble begins when the environment changes faster than the cycle can close, and when the act of observing and acting changes what is being observed.

Two cases frame the book. One is a navigation exercise in which a team kept a workable picture of where it was even after an instrument failed, because the picture was distributed across the crew, the artifacts and the routine rather than held in one head. The other is a fire in which a crew's account of its situation collapsed and could not be rebuilt in time. Both are readings of published accounts, not experiments, and they are used to show a shape rather than to prove a point.

Four assumptions fail in that territory.

That the situation exists to be perceived before action. Much of the time an organization acts in order to find out. The act is a probe, and what it produces is information that did not exist beforehand.

That more data improves the picture. Data does not say what it means, and information consumes attention. A unit can raise its data volume and lower its capacity to notice, because attention is the scarce input and nothing in a dashboard decides what is worth looking at.

That one actor holds the picture. The unit that thinks is not a person. It is a system of people, artifacts, routines and roles, and the picture is held between them, in representations that are handed from one medium to the next.

That understanding can be replaced by speed. Speed is a consequence of the quality of the model a unit acts on. Set as a target, it invites fast decisions on poor models. The book's author makes this point against his own earlier book, which treated the loop as a speed instrument in three places and is corrected in the open draft.

2. The idea: five activities, one loop

Understanding is not a stage between noticing and deciding. It is a loop of five activities, and every one of them has a collective form.

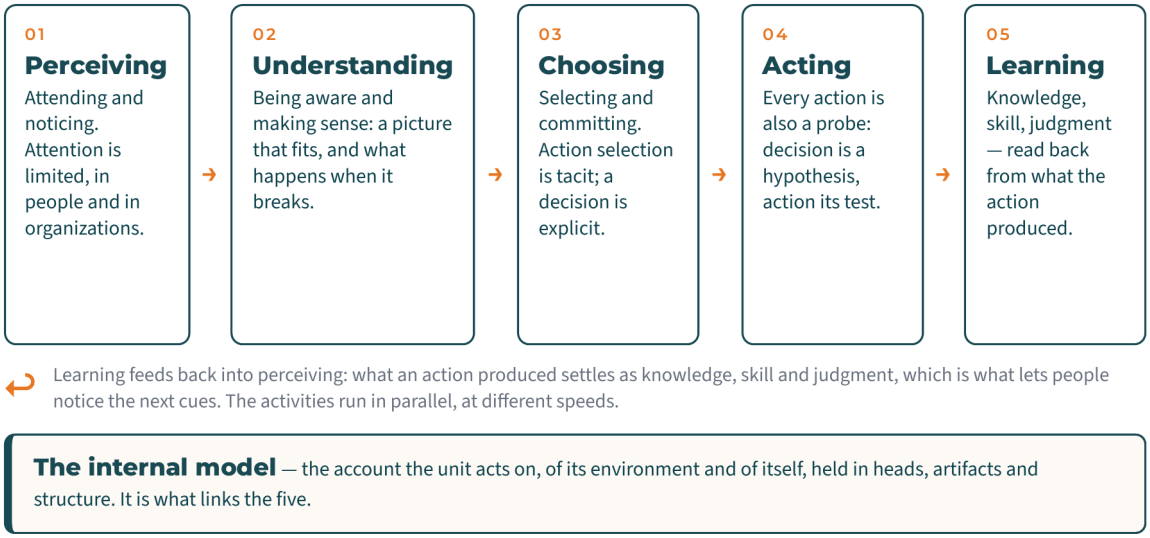


Figure 1 — The five activities, one loop, and the internal model at its centre. The activities run in parallel, at different speeds, and feed back into each other.

Perceiving — attending and noticing. Attention is limited in people and in organizations; what gets attended to is designed, usually by accident.

Understanding — being aware and making sense. Two literatures divide the work: situation awareness for holding a picture that fits, and sensemaking for what people do when the picture breaks.

Choosing — selecting and committing. Two terms are kept apart, and the distinction is one of the book's contributions. Action selection is continuous and mostly tacit: for an experienced person the situation is recognized as a type and a workable action arrives with it. Decision is the explicit, accountable commitment that others can rely on and later check.

Acting — every action is also a probe. In the full loop, decision is a hypothesis and action is its test, and some actions change what the actor knows rather than what the world is.

Learning — knowledge, skill and, over time, judgment. What an action produced is read back, and what is learned settles as the capacity to notice the next cues.

The internal model is what links the five. It is the account an organization acts on — of its environment and of itself — held in heads, artifacts and structure. The book keeps two senses apart. In the loose sense, every organization acts on some account, which is uncontroversial. In the strong sense, a regulator needs a model of what it regulates, which is argued for rather than proved: the cybernetic warrant is rigorous inside a scope and is not a proof of its own title.

One qualification belongs to the loop. For experienced people, understanding and the action arrive together; recognition is not a sequence of stages. "Understanding comes first" describes a dependency, not a timeline.

3. Where understanding is held

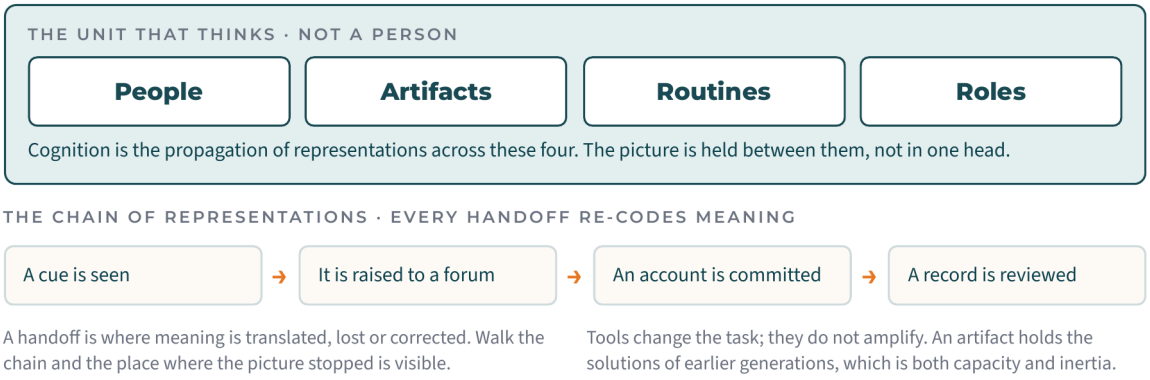


Figure 2 — Where the picture lives. Cognition is the propagation of representations across people, artifacts and media, which makes it traceable: one can see which step a person carries, which an artifact, and where the picture is lost.

Three consequences follow from taking distribution seriously.

Structure is cognitive architecture. Division of labor is division of cognition. Spans of control are spans of attention, and layers add latency, filtering and distortion. A handoff is a point at which a representation is re-coded for the next medium, and meaning is translated, lost or corrected on the way. In the language of the Viable System Model, the function that models the environment is System Four, and its test is severe and checkable: if nothing the unit decides depends on the model attributed to it, the model is trivial, and the unit is steering by something else.

Groups can understand more or less than their members. Whether a group does depends on conditions rather than on goodwill: diverse perspectives, equal participation, social sensitivity and structured interaction. More communication is not always better than less, and overlapping knowledge is a source of robustness rather than waste. Two designs carry most of the weight: forums that separate the competition between accounts from the moment one is committed to, and psychological safety, which is a precondition for reporting rather than a mood.

Whose account counts is a question of power. Leaders give sense as much as they make it. The design question is not which interpretation the leader holds but whether orientation is enabled: identity and purpose set, forums kept open where accounts can compete, people who report weak signals protected, decision rights placed where the relevant model sits. When a frame collapses, role structure goes with it, and rebuilding it is leadership work.

4. What a diagnosis examines

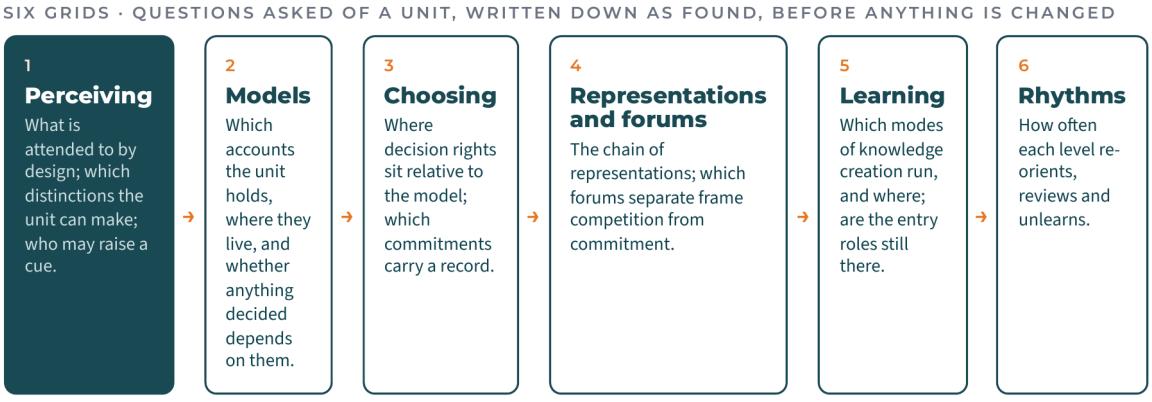


Figure 3 — Six grids, one per element of the design. Each is a set of questions asked of a unit, and the answers are written down as they are found before anything is changed.

Grid	The question it asks
Perceiving	What is attended to by design, which distinctions can the unit make, who may raise a cue and to whom, how is a cue amplified, is it safe to report, what do machines watch
Models	Which accounts does the unit hold, where do they live, of what, who drew their distinctions, does anything decided depend on them, how and how often are they revised
Choosing	Where do decision rights sit relative to the model, which commitments carry records, what releases an action and what stops it
Representations and forums	What is the chain of representations and where are its handoffs, which forums separate frame competition from commitment, who gives sense
Learning	Which of the four modes of knowledge creation run, and where; are the entry roles that grow judgment still there
Rhythms	How often does each level re-orient, review and unlearn

Two constraints govern the design. The first is recursion: every unit that faces an environment of its own needs its own understanding, so the grids are filled in per level, and one picture for the whole company hides what its parts need to see. The second is that knowledge sits in two places at once — created by individuals and amplified by the organization, and also held by the system of people and artifacts — so every row asks both questions: what does the system hold, and who among the people holds it. A design that builds only the individual's path leaves the system's learning unsaved; a design that builds only the system's records leaves nobody who can grow the judgment the system needs.

Measuring understanding invites a trap that looks like rigor. The loop has countable properties, above all how fast it turns, and once they are counted they become targets. The alternative is to measure the loop's conditions as traces rather than rates: follow one cue from the edge to wherever it stopped; count how many commitments carry a record with an expectation, an owner and a review date; ask when the unit last changed its account of its situation and what made it do so; ask where shared experience and articulation happen in the week, and whether the entry roles still exist. Speed and latency stay in the picture as outcomes to watch, never as targets to set.

None of these is a validated instrument. They are the questions the argument leads to.

5. What changes when models sit in the loop

AI makes content abundant and validated judgment scarce. The volume of candidates rises; the number of people who can tell a good one from a plausible one does not. The sharpest way to place a model in the loop is to ask which of the four modes of knowledge creation it can take over.

Combination is what language models do best: recombining explicit knowledge at a volume no team can match. Externalization they can support, at the risk of replacing people's own articulation of what they know. Socialization and internalization they cannot perform for people, because both require a person doing something with others or with their own practice over time.

An organization that lets machines take over combination and the entry-level tasks, and does not protect socialization and internalization, becomes better informed and less capable. The mechanism is stated as a limit rather than a policy: delegation is legitimate only as deep as the capacity to evaluate what comes back, and the depth is set per task, not per organization. The practical test is one question, asked of any role that receives model output: could the person in it tell if the output were wrong, and how would they know?

6. What this does not promise

No design in this paper delivers faster decisions, a score for organizational understanding, a maturity level, a dashboard that tells you what your situation means, or the assurance that a redesign will improve adaptation.

It also does not promise that a leader's clearer vision will do the work. The argument is explicitly against the idea that understanding is concentrated at the top and delivered downward: that is the simplification the book spends its third part removing.

What it does claim is narrower. If understanding is produced in a loop by a distributed system, then the shape of that system can be chosen, and choosing it is a design task rather than a communication task.

7. How you would know it is wrong

The book's convention is that every chapter states its claims and what would count against them. Three of those conditions carry the argument this paper restates.

If investing in shared understanding did not help. Organizations that invest in forums, representations and time for orientation adapting no better than organizations that invest only in faster decision paths or more data, in comparable complex settings. Where to look: the high-reliability literature, the studies of psychological safety, and the accident analyses in which cues crossed layers and were lost.

If delegation to models cost nothing. Organizations that delegate combination and interpretation without protecting socialization and practice keeping their judgment capacity as well as those that protect it. Where to look: the expertise literature, and studies of how experts and novices use the same tool differently.

If the design could be installed without changing anything. The forums, records, channels, trajectories and rhythms going in without any change in what the organization notices, chooses and learns. The traces described above would show it, because the cues would still stop where they stopped and the records would go unread.

The evidence position is stated plainly in the book and should not be softened here: it is a synthesis with 122 evidence records behind it, and most of its central claims are assessed `insufficient-evidence` rather than supported. The book is published as an open draft for exactly that reason, and it carries a list of the places where it departs from the author's own earlier book.

8. What this means for an organization that has data and no picture

Start with a trace, not a redesign. Four questions can be asked of any unit this week, and the answers are found by walking the chain rather than by surveying.

What does this unit attend to by design? Take its agendas, dashboards and standing forums, and ask which of them were chosen and which happened. Then ask which categories its reports use, and who drew them. A cue that fits no distinction is not weak; it is invisible.

Which account does it act on? Take the last three commitments the unit made, and ask what account of its situation each rested on. If there is no written account, or if nothing decided depends on the one that exists, the unit has a plan and a calendar rather than a model.

Who may raise a cue, to whom? Follow one signal from the person at the edge who saw it to whoever could act on it. Count the handoffs, and name the place where it was dropped. Then ask what happened to the last person who reported something unwelcome.

When did the unit last change its account? If the answer is "at the annual planning round", the model is being revised by the calendar rather than by what the unit's actions produced.

What follows is not a program. The book's own starting shape is three months for one unit at one level: a month of tracing the loop as it is found and having the unit's own people correct the description, a month installing the two cheapest elements — a commitment forum with decision records, and a channel for cues with a named receiver who tells the sender what happened — and a month in which the first review dates fall due. Then the next level of recursion, with the interface between the two designed on purpose. What travels from the first unit to the second is people who have run the loop once, not a manual.

9. Lineage

The argument stands on a flat list of theories, each answering one question, and it names where they disagree instead of settling the disagreements early. Weick's sensemaking answers how understanding is built and why it collapses. Endsley's situation awareness answers what knowing what is going on consists of. Boyd's orientation, with Craik, Beer and Klein, answers how the picture is kept current. Klein's recognition-primed decision answers how understanding becomes action. Hutchins' distributed cognition answers where understanding is held and how it travels. Woolley, Edmondson and the transactive-memory literature answer when a group understands more than its members. Beer's Viable System Model and Ashby's requisite variety answer which structures could carry what a unit faces. Gioia, Chittipeddi and Maitlis answer whose account counts. Nonaka and Takeuchi, with Argyris and Schön, answer how knowledge is created, spread and lost. Dreyfus, Ericsson and Polanyi answer how learning becomes judgment. Snowden's Cynefin places which situations need which kind of understanding.

The disagreements are part of the argument rather than a gap in it: Weick against the internal model, Hutchins against Nonaka on where knowledge sits, Klein against Cynefin on what recognition can do in the complex domain. A reader who wants one settled theory will not find it here. A practitioner who wants the questions that expose a broken loop will.