

Wolfson Tool Factory, Cambridge , 23. April 2026. Workshop on co-production of theory-based management and consulting tools. Cambridge University

# **Generative AI in Organisations: Better and Faster Decision Processes? A Systems-Theoretical Background for the Design of AI-CATCH**

Mogens Grosen Nielsen  
Nielsen Statistics Consulting, Copenhagen, Denmark.  
[mogensgrosen@gmail.com](mailto:mogensgrosen@gmail.com)

## ***Abstract***

This article repositions the discussion of generative AI in organisations by starting from epistemology rather than from implementation. If cognition is treated as construction, then organisational problems are not simply found but constituted through observation, and AI cannot be assessed as a neutral answer engine.

From this point of departure, the article moves into Luhmannian organisation theory through the functional method, decision premises, and structural coupling. The practical object is AI-CATCH, presented as a guided communication architecture that supports the translation of local issues into reviewable and reusable decision forms.

The argument is that generative AI can contribute to better and faster decision processes only under strict conditions: legitimacy must be treated as procedure rather than as moral proclamation; options must be tested across relevant couplings to function systems such as science, law, politics, and public communication; and draft decision premises must be made public to the relevant parties in forms that permit review, contestation, and traceable approval. Standards are therefore treated not as truth-bearing solutions, but as providers of terminology and rules of the game—closer to the rules of a football match than to a theory of the world.

They matter because they structure what counts as an acceptable move, not because they dissolve contingency. The Cambridge demonstration is thus not a claim that AI solves organisational problems, but a claim about the conditions under which AI-supported communications may become more defensible, accountable, and operationally effective.

## **1. Introduction**

Public and private organisations currently face a double pressure. They are expected to use generative AI productively, but they are also expected to retain control over responsibility, defensibility, and

legitimacy. In official statistics and similar knowledge-intensive settings, this pressure is intensified by the proliferation of standards, frameworks, governance templates, and AI guidance. The recurring problem is not the lack of guidance. It is the weak translation of guidance into communications that can actually shape later decisions (Luhmann, 2018; Knudsen, 2007)..

This article takes a different route. Its starting point is epistemological. If cognition is treated as construction rather than mirroring, then neither organisational problems nor AI-generated responses can be treated as self-evident givens. This opens a path into organisation theory that leads through the functional method, decision premises, structural coupling, and legitimacy through procedures.

The article is written as background for the Cambridge workshop, where the practical object is not a grand theory of AI but the design logic of a demonstrable tool. The tool is AI-CATCH. The theoretical question is narrow but demanding: under which conditions can generative AI support better and faster decision processes without displacing organisational responsibility into opaque plausibility effects?

## **2. Epistemological point of departure: cognition as construction**

Luhmann's theory of cognition offers the most important starting point for the present argument. The central move is not from ignorance to knowledge, but from subject-object thinking to observation. Cognition is not a copy of reality; it is produced through operations of distinction and indication. Every observation makes something visible and leaves something else unmarked. The point is not that reality disappears, but that access to reality is always selective and system-relative.

This has direct implications for organisational analysis. When an organisation states that it has a quality problem, a coordination problem, or an AI governance problem, it is not merely reporting an object already waiting in the environment. It is stabilising a distinction: acceptable/unacceptable, coordinated/uncoordinated, defensible/undefensible, efficient/inefficient, transparent/opaque. The same issue can therefore appear differently depending on the distinction used, and different distinctions open different paths of action (Luhmann, 1995; Luhmann, 2013)..

For the design of AI-supported problem solving, five distinctions are especially important. First, the distinction between operation and observation: observation is itself an operation of distinguishing. Second, the distinction between first-order and second-order observation: the organisation can observe the issue, but it can also observe how it is observing the issue. Third, the distinction between other-observation and self-observation. Fourth, the distinction between what the observation makes visible and what it necessarily leaves blind. Fifth, the distinction between truth and other evaluative forms such as legitimacy, legality, timeliness, proportionality, and trustworthiness (Luhmann, 1995; Seidl, 2006)..

Once this epistemological move is made, the transition into organisational theory becomes clearer. Better descriptions do not automatically improve decisions. What matters is whether observations can be transformed into communications that remain connectable in later organisational operations.

A further consequence must be made explicit. AI-CATCH, this article, and the consulting organisation behind the technology are not carriers of truth standing outside the field they describe. They are themselves observations that use distinctions and therefore make some possibilities visible while leaving others less visible. Their value lies not in claiming a privileged description of the organisation, but in

fostering better and more disciplined self-observation of organisational self-descriptions. This includes observation of how this paper observes, how AI-CATCH formats observation, and how the consulting organisation constructs problems, solutions, and relevance. In that sense, the tool and the paper should be treated as aids to second-order observation: they can improve reflection only if they themselves remain observable, criticisable, and revisable (Luhmann, 1995; Besio and Pronzini, 2010).

### **3. From epistemology to organisation theory: functional method and decision premises**

The epistemological move by itself is insufficient unless it is connected to an organisational mechanism. Here two concepts are decisive. The first is the functional method. Instead of asking which description or solution is simply correct, the functional method asks which problem a proposed arrangement solves, under which conditions, with which side effects, and which functional equivalents are excluded. This shifts attention from truth claims to observation design and comparative problem solving (Knudsen, 2010; Luhmann, 2018).

This matters because modern organisations operate under multiple, incompatible external constraints. In official statistics, for example, a proposal may be methodologically strong but legally weak, politically attractive but scientifically dubious, or efficient but reputationally risky. The relevant issue is therefore not correctness in the abstract, but multifunctional defensibility under several couplings at once (Knudsen, 2007; Van Assche and Verschraegen, 2008).

The second concept is decision premises. Organisations do not decide from scratch. They reproduce themselves through chains of decisions, and these decisions are shaped by premises stabilised in earlier communications. These include programmes, roles, communication channels, and cultural taken-for-granted. A standard, framework, or AI output becomes organisationally consequential only when it is translated into such premises and can be cited again in later decisions (Luhmann, 2018; Seidl, 2006).

This organisational mechanism can be specified more precisely through Luhmann's concept of autopoiesis and through the distinction between system types. Psychic systems reproduce consciousness; interaction systems reproduce communication in situations of mutual presence and turn-taking; organisations reproduce themselves through decision communications connected to further decisions; and function systems such as law, science, politics, and the economy reproduce communication through their own binary codes and programmes. The point is not to rank these system types, but to avoid collapsing them into one another. A person can understand a problem very well as a psychic system, a workshop can discuss it productively as an interaction system, and yet no organisational effect follows unless some communication is translated into decision premises that later decisions can cite. Conversely, an organisation may stabilise a premise that many participants only partially understand. Functional systems then add further selectivity, because what counts as acceptable in science, law, politics, or public communication is not the same (Luhmann, 1995; Seidl, 2006; Luhmann, 2018).

Autopoiesis therefore sharpens the diagnosis of AI in organisations. AI-generated text, classifications, or comparisons do not enter an organisation as binding instructions. They appear as irritations and as variation offers for different system references. Psychic systems may use them for orientation,

interaction systems may use them for discussion, and organisations may use them for drafting; but only the organisation can stabilise them as decision premises, and only under conditions that remain compatible with relevant function-system couplings. This is why AI-CATCH must be designed not as a universal answer engine but as a translation architecture across system types (Luhmann, 1995; Luhmann, 2018; Nielsen, 2024).

This leads to a sharper diagnosis of the current AI debate. The recurrent problem is not simply poor implementation. It is a premise gap. Organisations have many standards, models, and now AI-generated proposals, but these often remain external descriptions or ceremonial references. They rarely become stop-rules, review obligations, exception procedures, or traceable decision forms (Luhmann, 2018; Knudsen, 2007).

#### **4. AI-CATCH: from answer engine to guided communication architecture**

The Cambridge demonstration is therefore best understood as a design proposition. AI-CATCH is not presented as a system that gives the right answer. It is presented as a guided communication architecture for reconstructing problems, comparing alternatives, and drafting communications that may later function as decision premises (Luhmann, 2018; Nielsen, 2024).

The central design choice follows directly from the epistemological argument. AI-CATCH begins from restatement and reconstruction, not from immediate recommendation. A user statement such as “we have a problem integrating administrative data” is treated as an observation that must first be clarified. Which distinction is currently guiding the issue? Which constraints are salient? Which blind spots follow from the present framing? Only then does the system widen variation by comparing alternatives and drafting candidate outputs (Luhmann, 1995; Knudsen, 2010).

This design also explains why output forms matter. Restatements, contrasts, issue briefs, checklists, reviewer notes, escalation cues, and decision briefs are not secondary packaging. They are the communicative forms through which observations can survive beyond one interaction and become available for later organisational processing (Luhmann, 2018; Besio and Pronzini, 2010).

The knowledge architecture is equally important. AI-CATCH uses curated materials such as standards, frameworks, best-practice examples, templates, and prior project materials. But these are not treated as ready-made solutions. They function as selective resources for comparison, justification, and critique (Knudsen, 2007; Nielsen, 2024).

Seen from the perspective of system types, AI-CATCH is designed to support transitions rather than to erase differences. It should help psychic systems move from private irritation to more precise problem formulations; help interaction systems compare observations and make blind spots discussable; help organisations convert selected observations into reviewable decision premises; and help these premises remain defensible under the different expectations generated by relevant function systems. The architecture is therefore not only user-centred but system-reference-sensitive.

#### **5. Core design principles of AI-CATCH**

Five design principles follow from the theoretical argument.

First, problem reconstruction before solution generation. The tool should not assume that the problem is fully given. It should reconstruct how the issue is currently being observed and what this framing excludes (Luhmann, 1995; Seidl, 2006; Luhmann, 2018).

Second, blind-spot analysis as a formal step. Blind spots are not simply missing information. They arise from the distinction currently used, and they are therefore part of the problem construction itself.

Third, functional comparison instead of one best answer. AI-CATCH should generate several functionally equivalent options and compare them across conditions, side effects, and excluded alternatives.

Fourth, persona-specific interaction. The same knowledge base should support different communicative situations: solo exploration, workshop co-creation, review, approval, and platform operation.

Fifth, reviewable hand-over outputs. AI becomes organisationally relevant only when its outputs are translated into draft premises that can be reviewed, revised, rejected, or approved.

Each of these principles also has a systems-theoretical side. Problem reconstruction is needed because psychic systems and interaction systems often stabilise different descriptions of the same issue. Blind-spot analysis is needed because every observation within every system type marks one side and leaves another side less visible. Functional comparison is needed because organisations must remain decision-capable under multiple couplings rather than under one privileged perspective. Persona-specific interaction is needed because a workshop, a reviewer, and an approver operate in different communicative situations. Reviewable hand-over outputs are needed because only organisations can transform provisional observations into premises for later decisions.

## **6. Workflow and persona architecture**

The practical workflow can be described in four connected phases: problem and solution analysis, co-creation, review, and approval.

In the first phase, the system supports analysis. The user enters a concrete issue and a minimal context. AI-CATCH restates the issue, identifies missing information, reconstructs the guiding distinction, and maps candidate couplings and constraints.

In the second phase, the system supports co-creation. In workshops and stakeholder settings, AI-CATCH helps compare options, articulate trade-offs, and reformulate tentative proposals in ways that different participants can assess.

In the third phase, the system supports review. Drafts are tested for clarity, consistency, defensibility, and overlooked consequences. The point is not to validate AI output, but to increase organisational selectability through renewed observation.

In the fourth phase, the system supports approval. Selected drafts are condensed into decision briefs, checklists, escalation rules, or other hand-over forms that can be used in later communications.

These phases are linked to persona-specific interfaces. The current logic distinguishes at least five personas: Workshop Participant, Solo Explorer, Reviewer, Approver, and Platform Operator. The point is

not marketing segmentation, but communicative specificity. Different roles require different forms of support, different outputs, and different levels of condensation.

The same workflow can also be described in terms of system types. Problem and solution analysis often begins with psychic systems: someone is irritated and tries to formulate the issue. Co-creation unfolds mainly in interaction systems: participants test descriptions, compare alternatives, and observe one another's observations. Review and approval belong primarily to the organisational system because they concern which communications will count as premises for later decisions. Throughout all four phases, relevant function systems remain present as environmental references and coupling constraints: science, law, politics, economy, and public communication continue to shape what can be accepted, justified, or contested. The workflow is therefore a structured movement across system types, not a single homogeneous process.

## **7. Standards as premise resources and as rules of the game**

The practical design is strengthened by linking it to established frameworks, but only under a premise-oriented interpretation. GSBPM can function as a process premise, GAMS0 as a role and capability premise, and quality frameworks as quality premises. They help transform diffuse irritation into a more structured issue, infer which roles matter for review, and test whether a proposal survives beyond its initial plausibility (Knudsen, 2007; Nielsen, 2024).

The point is not to turn these frameworks into dogmatic answer keys. Their value lies in their selective translation into organisationally usable distinctions, review points, and hand-over forms.

In this respect, standards are best understood less as descriptions of reality than as rules that specify what counts as an acceptable move. The analogy is not scientific law but the rules of a football match. The rules do not tell the players which exact move will succeed, nor do they eliminate uncertainty, conflict, or strategy. They make coordinated play possible by defining what is permitted, reviewable, sanctionable, and recognisable as part of the game.

Organisational standards function in a similar way. They do not settle the issue by themselves. They structure the space of acceptable operations, provide references for objection and review, and make it possible to distinguish defensible from indefensible moves. Their practical force depends on whether they are translated into concrete programmes, review procedures, thresholds, and exception handling.

## **8. Legitimacy as procedure, coupling, and public decision premises**

If the article is to speak about better decision processes, then legitimacy cannot remain implicit. A systems-theoretical account cannot reduce legitimacy to moral virtue, rhetorical reassurance, or ceremonial compliance. Following Luhmann's problem-setting, legitimacy concerns the generation of acceptance for collectively binding decisions through procedures, not through supposedly legitimate contents in themselves. Ocampo's recent reconstruction is useful here because it restates that legitimation through procedures concerns the contribution of organisations and procedural systems to the acceptability of decisions in functionally differentiated society (Ocampo, 2025; Luhmann, 2018).

For present purposes, three consequences follow. First, legitimacy must be treated procedurally. Better decisions are not simply substantively superior decisions. They are decisions whose preparation, testing, review, and approval have been organised in such a way that responsibility does not disappear into AI plausibility or into hidden premises. Procedure matters because it structures how conflict is processed, how objections are handled, and how motives are neutralised sufficiently for decisions to become acceptable beyond the moment of choice.

Second, legitimacy requires explicit coupling to relevant function systems. A recommendation that looks compelling in one dimension is not therefore legitimate. In statistical organisations, the relevant couplings usually include science, law, politics, organisation, and public communication. Better decision processes therefore require that candidate outputs be tested for multifunctional defensibility: scientific adequacy, legal permissibility, organisational feasibility, political or mandate compatibility, and public explainability.

Third, draft decision premises must be made public to the relevant parties. Public here does not mean indiscriminate openness to everyone. It means that the relevant parties in the procedure can see which distinction guided the proposal, which standards were invoked, which alternatives were compared, which reviews were carried out, who may interrupt or escalate, and under which conditions the recommendation remains defensible. Decision premises that remain hidden may still coordinate internally, but they do so at the cost of contestability and accountability.

This point is especially important for AI-supported work. AI can accelerate drafting, comparison, and summarisation, but it cannot by itself provide legitimacy. Ocampo's discussion of artificial communication sharpens the limit: where the neutralisation of motives cannot be guaranteed technically, legitimacy cannot be offloaded to the system. The implication for AI-CATCH is practical rather than apocalyptic. Keep AI at the level of variation, comparison, and draft production; keep selection, review, approval, and the publication of premises within organisational procedures (Ocampo, 2025; Nielsen, 2024).

Seen in this way, better and faster decision processes do not mean faster closure at any price. They mean quicker movement from irritation to explicit, reviewable, multifunctionally tested, and publicly reconstructable decision premises.

## **9. What will be demonstrated in Cambridge**

The demonstration value of AI-CATCH lies in making the design logic visible. The system can show how a local issue is reconstructed, how relevant standards are used as premise resources, how alternatives are compared, and how outputs are shaped differently for workshop, review, and approval settings.

The theoretical claim to be demonstrated is narrow but important. Generative AI can improve the speed and quality of decision preparation when it is embedded in a premise architecture that preserves second-order observation, review, procedural legitimacy, and public reconstructability for the relevant parties.

This also provides a more defensible interpretation of the phrase better and faster decision processes. Better does not mean objectively optimal. Faster does not mean less reflective. The claim is that AI-

CATCH can help organisations move more quickly from vague irritation to explicit comparisons, reviewable drafts, and proceduralised decision premises without collapsing legitimacy into hidden model output.

## 10. Critical discussion

The strongest objection is that AI-CATCH may simply add one more framework to organisations that already suffer from framework overload. That objection is serious. The answer cannot be that this framework is finally the right one. The answer must be that AI-CATCH is not meant to add a new layer of abstract doctrine, but to support the translation of existing standards and constraints into more workable communications.

A second objection concerns the phrase artificial communication partner. For strict systems theory, this formulation risks saying too much. AI does not communicate in the same sense as social systems communicate. The phrase is therefore acceptable only if it is understood operationally: AI participates in the preparation of communications that organisations may then take up, reject, or transform.

A third objection concerns legitimacy. Because organisations are structurally coupled to multiple function systems, no single procedure can guarantee universal acceptance. That is correct. The present argument is therefore limited. It does not claim that AI-CATCH solves the legitimacy problem. It claims that proceduralised review, explicit coupling tests, and public decision premises for relevant parties are superior to hidden plausibility effects and to ceremonial references that never enter real decision chains.

## 11. Conclusion

The argument of this article has been that epistemology must come first. Once cognition is treated as construction, the discussion of generative AI in organisations can no longer begin from the fantasy of neutral descriptions and correct implementations. It must begin from observation, blind spots, and the organisational problem of transforming external communications into decision premises.

From that perspective, organisations face the same basic problems as before: coordination, review, ownership, legitimacy, and defensibility under multiple couplings. Generative AI does not remove these problems. It amplifies variation and accelerates drafting. Whether this becomes organisationally useful depends on design.

AI-CATCH is proposed as one such design. Its contribution is not that it provides answers, but that it supports a procedure through which issues are reconstructed, alternatives compared, standards translated into rules of the game, decision premises made public to the relevant parties, and recommendations prepared for review and approval. Under these conditions, one may speak carefully of better and faster decision processes. Without them, the phrase is only rhetoric.

## References

- Besio, C., & Pronzini, A. (2010). Inside organizations and out: Methodological tenets for empirical research inspired by systems theory.
- Knudsen, M. (2007). Structural couplings between organizations and function systems: Looking at standards in health care.
- Knudsen, M. (2010). Surprised by method: Functional method and systems theory.
- Luhmann, N. (1995). Social Systems.
- Luhmann, N. (2012). Theory of Society, Volume 1.
- Luhmann, N. (2013). Theory of Society, Volume 2.
- Luhmann, N. (2018). Organization and Decision.
- Ocampo, S. P. (2025). Artificial Legitimacy? Possibilities, paradoxes, and ironies of the relationship between digital communication and legitimation through procedures.
- Nielsen, M. G. (2024). A New Approach to Organisational Change Using Social Systems Theory and Generative AI.
- Scheytt, T. (2006). Management Accounting from a Systems-Theoretical Perspective.
- Seidl, D. (2006). The Basic Concepts of Luhmann's Theory of Social Systems.
- Van Assche, K., & Verschraegen, G. (2008). The limits of planning: Niklas Luhmann's systems theory and the analysis of planning and planning ambitions.