SMS SECOND MIND SOLUTIONS · WHITE PAPER

AN URBANIST'S APPROACH TO MAPPING THE STRATEGIC AI FRONTIER

# Stop Asking AI for Answers. Ask It for Maps.

*A planner's discipline for mapping the known, finding the edge, and stepping beyond it.*

**Terence Milstead, PhD**

Founder & Principal · Second Mind Solutions, LLC

# The Map and the Answer

Most teams ask AI for answers. The best teams ask AI for maps.

That distinction is one I learned through my early training in urban planning. Planners do not start by asking what to build. They start by mapping what is there, what is possible given the constraints, and where the field has not yet built. That habit of mind, the disciplined survey of a possibility space before committing to any one move, turns out to be the most useful frame I know for thinking about how teams should be using AI.

Most teams are not yet using AI that way. They ask a question, they get an answer, they move on. The output is competent. The cycle is fast. The result, almost by definition, is a fluent restatement of what the field already says.

This is the answer-machine model, and it is the dominant pattern in professional services right now. It produces faster versions of work the team already knew how to do. For operational tasks, that is genuinely valuable. For knowledge-work firms whose entire value proposition rests on producing thinking their clients could not have generated without them, faster is not the same as different.

The structural problem with the answer-machine model is not the AI. It is the prompt. Teams ask the model to do what they would otherwise have done themselves, only quicker. The model obliges. Nothing in that loop produces anything new. But there is a different way of using AI, informed by the cartographic discipline urban planners bring to any problem space. Instead of asking the model for an answer, you use it to map the field of possibilities, identify which configurations are coherent and which rule themselves out, and surface the underbuilt corners where genuine innovation tends to live. That is the difference between a faster version of work that already exists and a real chance at producing something new.

And while planners may have internalized a cartographic way of thinking, it does not originate with them. In fact, one of the champions of this systematic mapping of a problem space was Swiss astrophysicist Fritz Zwicky.

# Zwicky's Morphological Box

Born in Bulgaria (where I happened to live for three years) and raised in Switzerland, Fritz Zwicky spent more than four decades at the California Institute of Technology. He predicted dark matter, neutron stars, and gravitational lensing decades before anyone else, and across the same career developed a structured thinking method he called morphological analysis. In 1971, he described the morphological outlook as his "philosopher's stone."[1]

The method is simpler than it sounds. Start by defining the problem precisely. Then identify its key parameters: the dimensions on which any solution will vary. For each parameter, list the possible values. The result is a structured field, often drawn as a grid, in which any candidate solution can be described as one combination of values, one drawn from each parameter.

That is the easy half. The harder half is what Zwicky called the principle of contradiction and reduction, and what Tom Ritchey, who runs the Swedish Morphological Society and has done the most to formalize the method since Zwicky's death, calls cross-consistency assessment.[2] You walk the field and eliminate combinations that are internally inconsistent, empirically implausible, or contextually impossible. What survives is a workable solution space, typically a manageable subset of an otherwise unmanageable field. That second stage is what separates morphological analysis from ordinary ideation.

But Zwicky's approach was not contained to astrophysics. By the late 1960s, the Swiss systems theorist Erich Jantsch had begun explicitly applying morphological analysis to technological forecasting and planning. Hasan Özbekhan, the Turkish-American planning theorist who would later spend two decades at the Wharton School, took the same logic into the planning of complex socio-technical systems: cities, institutions, and societal-scale change.[3] The migration was not accidental. The problem types are structurally similar. In both, you face a field of possibilities, most of which are ruled out by constraints, and a buildable space that only emerges after disciplined mapping. Other figures in the planning tradition gave related diagnoses. Horst Rittel and Melvin Webber called these ill-structured spaces "wicked problems,"[4] and Ian McHarg made the mapping method visual in *Design with Nature*, overlaying constraint surfaces of soils, slope, hydrology, and ecology to identify the parcels where development was simultaneously feasible and contextually appropriate.[5]

A parcel of land can technically support many configurations; most are ruled out by the soil, the zoning, the infrastructure, or the politics. What remains is the buildable design space. The habit of enumerating possibilities before committing to one is older than the field of AI by half a century, and it is exactly the habit most teams skip when they use AI.

## An applied example

Let's take an example from industrial design. The parameters in a vehicle-design problem might include the powertrain, the body style, the seating capacity, the drive type, the price tier, and the target buyer profile. Each parameter has several possible values. Mapped out, the field looks like this:

**Morphological Box: Vehicle Design**

| Parameter | Value 1 | Value 2 | Value 3 | Value 4 | Value 5 |
|---|---|---|---|---|---|
| **Powertrain** | **Electric** | Hybrid | Gasoline | Hydrogen | — |
| **Body style** | Sedan | SUV | **Crossover** | Truck | Hatchback |
| **Seating** | 2 | 4 | **5** | 7 | — |
| **Drive type** | FWD | RWD | **AWD** | — | — |
| **Price tier** | Economy | Mid-range | **Premium** | Luxury | — |
| **Buyer profile** | Urban commuter | **Family** | Off-road | Performance | — |

*Figure 1. A simplified morphological box for vehicle design. The shaded cells trace one configuration: an electric, crossover-body, five-seat, all-wheel-drive, premium vehicle aimed at the family buyer. That configuration describes, more or less, a Tesla Model Y.*

The grid above describes dozens if not hundreds of formally possible vehicle configurations. The highlighted path is one of them. Change one value, swap "electric" for "hybrid" and "family" for "performance," and you have described a different vehicle, with different competitors and different economics.

The cross-consistency step is where the real work happens. Some combinations rule themselves out. A two-seat family vehicle is incoherent. A hydrogen-powered economy vehicle is impractical given today's refueling infrastructure. Other combinations are technically possible but commercially unattractive. A luxury hatchback aimed at the off-road buyer is unlikely to find an audience. After pruning, what remains is the workable design space.

Zwicky's argument was that human minds, working alone, cannot hold all the possible configurations of a complex problem in view at once. The discipline of the box forces the enumeration. The discipline of the consistency check forces the pruning. Without both steps, teams converge prematurely on whichever configuration came to mind first and call it the answer.

# The Frontier

Morphological analysis is usually framed as a way to surface options you would have missed. My extension of the logic, and the one I find most useful in professional services work, is what I think of as the frontier: the point at which a disciplined search of the known space still fails to produce a satisfactory fit.

Imagine you complete the vehicle exercise. You map every parameter, list every value, eliminate every inconsistent combination. And the configuration the market actually needs, the one your competitors have not built and the one your customer is asking for, is not in any cell of the grid.

You have not run out of ideas. You have systematically demonstrated that the answer you need requires something the grid did not contain. That something is usually one of three things. It may be a parameter the industry has not yet learned to think in. The early-2000s grid for cars did not have a "level of autonomy" row, because that parameter did not exist as a meaningful dimension. It may be a value range no one has yet considered, like a price tier below "economy" that only becomes possible with a different ownership model. Or it may be a redefinition of the problem itself: you are no longer designing a vehicle; you are designing a service that contains one.

That is the frontier. It is what planners call an edge condition: the position where the field's existing vocabulary has run out, and where genuinely differentiated work tends to come from. Crucially, it is a position you can only reach by exhausting the known first. Teams that skip the mapping and try to leap straight to the frontier mistake novelty for insight. They confuse "no one is doing this" with "this is worth doing."

# Asking for Maps

This is where large language models become useful in a way the answer-machine model misses.

LLMs are well-suited to one specific task: rapidly surfacing, restating, and recombining large amounts of already-articulated knowledge. That is, in effect, the work of the mapping stage. A team using an LLM well is not asking it "what should our strategy be?" The team is asking it "what are the known approaches to building credibility for a B2B technology brand in a crowded category? Map them. List the variations within each approach. Identify which combinations have been tried and which have been left untried. And tell me which combinations are inconsistent with our budget, our brand, and our timeline."

That is a different conversation from the one most teams are having with AI. It treats the model as an instrument for cartography rather than as an oracle.

My own research has shown that some people, when simply given access to an AI tool and instructed to use it as they best see fit, arrive at a simplified version of this use case all on their own. The 12-month enterprise AI pilot I led the user-experience research for at the Commonwealth of Pennsylvania, conducted in partnership with OpenAI across 175 employees and 14 state agencies, found that the most effective users treated ChatGPT as what one participant called "a bare bones starting point that we could use to expand upon," not as the source of finished work.[6] In the findings report we grouped this kind of use under what we called "Innovation Engines": using the tool to generate ideas, explore solutions, and accelerate problem-solving while leaving the judgment that follows to the human.

This pairing of structured methods with AI is not novel. A 2023 Harvard Business School field experiment with Boston Consulting Group consultants, titled "Navigating the Jagged Technological Frontier," found that AI substantially improved consultant performance on tasks within its capability boundary but degraded performance on tasks just outside it.[7] That finding maps directly onto the morphological argument. AI is at its most reliable when it is helping you exhaust the known. It is at its least reliable when it is being asked to substitute for the judgment that defines the frontier. The broader literature on human-AI collaboration treats this kind of pairing as an instance of hybrid intelligence, in which AI extends human cognition rather than replacing it.[8]

Three caveats are worth holding in mind, all well-documented in current research. First, LLMs hallucinate as a structural property of how they generate text;[9] a confident-sounding "map" can include cells that do not exist. Second, model performance degrades on information placed in the middle of long contexts, contradicting the common assumption that more context produces better recall.[10] Third, fluent output makes a shallow search feel complete. The cure for all three is the same: the mapping is heuristic, not authoritative. Human judgment has to decide what counts as a real parameter, which combinations are coherent, and whether the apparent frontier is genuine or just an artifact of weak framing.

## The Habit, and the Risks

Used this way, repeatedly, AI changes how a team thinks. Separating what is known from what is assumed. Distinguishing fluency from coherence. Asking whether the frame is adequate, not just whether the options inside it are exhausted. Over months, that becomes a cognitive habit, and a competitive one. Teams that have built it can tell, much faster than their competitors, when a strategy conversation has actually exhausted the available options and when it has only exhausted the team's patience.

That outcome is not automatic. The same workflow, used carelessly, produces the opposite result. A Microsoft Research and Carnegie Mellon survey of 319 knowledge workers found that higher confidence in AI was associated with less critical thinking during the task itself.[11] A 2024 study found that AI-assisted creative work was more individually novel but more similar across users than unaided work, which is the convergence problem in a single sentence.[12] A 2025 study found that frequent AI use was correlated with cognitive offloading and reduced engagement in critical thinking, particularly among younger users.[13]

The teams that get the upside are the ones that use the model as a provisional partner rather than an authority, and that treat every mapping exercise as something to be audited, not accepted. The scaffold becomes a crutch the moment the team stops auditing it.

# When You Hit the Frontier

What do you actually do when the mapping is complete and the frontier is in view? Three moves are legitimate.

First, **expand the parameter set**. The most common reason a known map fails to produce a fit is that the field is missing a dimension. The vehicle example is concrete: the grid that did not include "level of autonomy" in 2005 was not a wrong grid; it was an incomplete one. Adding the parameter does not solve the problem, but it relocates the conversation onto ground where the answer can exist.

Second, **redefine the problem complex**. Sometimes the issue is not the parameters but the question. If you are mapping vehicle configurations and the answer keeps not being in the grid, the right move may be to stop designing a vehicle and start designing a service that contains one. The morphological field shifts entirely.

Third, **commit to genuine invention**. Once you have exhausted the known and confirmed that neither a new parameter nor a redefined problem produces the answer, what remains is to make something new and accept the risk of being wrong about it. This is the only place where creative confidence is actually warranted, because it is the only place where the alternative, defaulting to the known, has been systematically ruled out.

For professional services teams, the value of the discipline is that it makes the difference between these three moves legible. Most strategy conversations conflate "we don't know yet" with "we need to invent something." A morphological pass separates them. It also separates teams that have actually done the work from teams that have only generated faster output.

One practical caveat applies to all three moves. Any team using LLMs this way on real client problems should think carefully about what gets pushed into the conversation. Strategic problem definitions, client variables, and competitive context all carry confidentiality obligations. The mapping is more useful, not less, when it is paired with explicit ground rules about what the team is willing to expose, to which model, in which environment.

# A Note on the Work

The way I work with professional services teams is not about teaching tools. The work is about building the organizational habit of using AI as a mapping partner: exhausting the known with discipline before claiming the next idea is new. The capability that emerges, the ability to find the frontier and the confidence to step beyond it, is what sustained competitive advantage looks like for teams whose value rests on differentiated thinking.

If your team is moving faster but not yet thinking differently, that is the conversation worth having. Visit **www.secondmindsolutions.com** to book a no-cost discovery conversation.

---

*Terence Milstead, PhD came to AI consulting from urban planning and strategic communications, and the planner's discipline still shapes how he approaches strategy work. He is the founder of Second Mind Solutions, an AI workflow performance consultancy that helps knowledge-work firms turn AI into measurable workflow improvements, and an Adjunct Lecturer in the Undergraduate Communications Program at The Wharton School at The University of Pennsylvania.*

# Sources


1. Zwicky, F. (1969). *Discovery, Invention, Research – Through the Morphological Approach*. New York: Macmillan. Zwicky's book-length treatment of the morphological method, in which he describes it as his "philosopher's stone."

2. Ritchey, T. (2006). "Problem structuring using computer-aided morphological analysis." *Journal of the Operational Research Society*, 57(7), 792–801. The canonical methodological clarification of Zwicky's approach, introducing cross-consistency assessment as the modern term for what Zwicky called contradiction and reduction.

3. Özbekhan, H. (1969). "Toward a General Theory of Planning." In E. Jantsch (Ed.), *Perspectives of Planning*. Paris: OECD. Özbekhan extended morphological-style systems thinking into the planning of complex socio-technical problems including cities, institutions, and societal-scale change. The OECD volume, edited by Erich Jantsch, is one of the clearest intellectual bridges between Zwicky's astrophysics-rooted method and the planning fields.

4. Rittel, H. W. J., & Webber, M. M. (1973). "Dilemmas in a General Theory of Planning." *Policy Sciences*, 4(2), 155–169. The seminal paper introducing the concept of "wicked problems," ill-structured problem spaces typical of urban planning where the question itself is unstable and the set of possible moves cannot be enumerated in advance. The intellectual context within which planning developed its mapping-first discipline.

5. McHarg, I. L. (1969). *Design with Nature*. American Museum of Natural History / Natural History Press. Landmark text in landscape and regional planning that introduced the layered suitability-mapping method: overlaying constraint surfaces (soil, slope, hydrology, ecology) to identify where development is simultaneously feasible and contextually appropriate. Methodologically a close cousin of Zwicky's morphological analysis.

6. Commonwealth of Pennsylvania (2025). *Lessons from Pennsylvania's Generative AI Pilot with ChatGPT*. Office of Administration, March 2025. A twelve-month enterprise AI pilot conducted in partnership with OpenAI across 175 employees in 14 state agencies. The report identifies three primary use patterns ("Innovation Engines," "Bureaucracy Hackers," and "Strategic Communicators") and documents the brainstorming and starting-point use case quoted here. Available at pa.gov.

7. Dell'Acqua, F., McFowland III, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., & Lakhani, K. R. (2023). "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality." *Harvard Business School Working Paper 24-013*. A randomized field experiment with 758 Boston Consulting Group consultants showing that AI raises performance inside its capability boundary and lowers it just outside.

8. Dellermann, D., Ebel, P., Söllner, M., & Leimeister, J. M. (2019). "Hybrid Intelligence." *Business & Information Systems Engineering*, 61(5), 637–643. The foundational framing of AI systems as complements to human judgment rather than substitutes, and the source most commonly cited for "hybrid intelligence" in management research.

9. Ji, Z., Lee, N., Frieske, R., et al. (2023). "Survey of Hallucination in Natural Language Generation." *ACM Computing Surveys*, 55(12), Article 248. The most-cited technical survey establishing hallucination as a structural property of current language models.

10. Liu, N. F., Lin, K., Hewitt, J., et al. (2024). "Lost in the Middle: How Language Models Use Long Contexts." *Transactions of the Association for Computational Linguistics*, 12, 157–173. Empirical evidence that LLM performance degrades on information placed in the middle of long contexts, contradicting common assumptions about reliable long-context recall.

11. Lee, H.-P., Sarkar, A., Tankelevitch, L., et al. (2025). "The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects from a Survey of Knowledge Workers." *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI 2025)*. Microsoft Research and Carnegie Mellon survey of 319 knowledge workers finding that higher confidence in AI tools is associated with reduced critical-thinking effort.

12. Doshi, A. R., & Hauser, O. P. (2024). "Generative AI enhances individual creativity but reduces the collective diversity of novel content." *Science Advances*, 10(28), eadn5290. Experimental evidence that AI-assisted creative outputs are more individually novel but more similar to each other across users.

13. Gerlich, M. (2025). "AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking." *Societies*, 15(1), 6. Survey evidence that frequent AI tool use is associated with cognitive offloading and reduced engagement in critical thinking, particularly among younger users.