

Dark Context: What Your Organization Knows but Never Wrote Down

Why a general model is confident, generic, and slightly wrong without your context

DarkContext.ai · Ola Möller^a and Andriy Zhukov^b

^aola@darkcontext.ai ^bandriy@darkcontext.ai

ABSTRACT

Every organization runs on knowledge it never recorded: the decisions no one minuted, the reasons behind a rule, the sense of how things are actually done that lives only in people's heads. We name this mass *dark context*, by analogy with dark matter in physics: invisible, unmeasured, and yet the thing that holds the visible structure together. The analogy has lately become practical rather than poetic. A general language model arrives with all of the world's public knowledge and none of yours, so absent your dark context it answers in a way that is confident, generic, and slightly wrong, which is the most expensive way to be wrong. This paper makes three contributions. First, we define dark context and separate it into four kinds along two axes, capture failure and access failure, yielding *unarticulated*, *articulated*, *lost*, and *fenced* knowledge, each with a distinct remedy. Second, we trace how a unit of knowledge matures along two axes: a human ladder, by which dark context unfolds out of one person's head into the shared practice of an organization, and a seven level scale of machine access, from tacit to machine runnable; we then describe the lifecycle by which captured context decays from *green* (current and queryable) through *gray* (stale but still trusted) back to *dark*. Third, we argue that context is an asset few organizations price, that extracting it carries a cost to the person who holds it, and that some context should be kept dark by design. The central claim is that in an age of capable models the binding constraint is no longer capability; it is context.

KEYWORDS *tacit knowledge · organizational memory · knowledge management · large language models · context engineering · institutional knowledge · knowledge taxonomy · knowledge lifecycle*

1 Introduction: the mass you cannot see

Physicists do not observe dark matter directly. They infer it because galaxies rotate as though they carried far more mass than their visible stars and gas can account for. Something unseen is doing most of the gravitational work [6], [7]. Organizations have a structurally similar problem. The documents, the wiki, the recorded decisions: that is the visible mass. It is dwarfed by what no one wrote down. The reason a rule exists. The customer you quietly never chase, and why. The order in which a release has to happen or it breaks. The judgment a senior colleague applies in two seconds and could not fully reconstruct if you asked.

We call this *dark context*: what an organization knows but never wrote down. Its relatives carry older names, institutional memory, know-how [3], tacit knowledge [1], tribal knowledge, but the framing we want is the physical one, because it captures the central and uncomfortable fact. The invisible part is not a footnote to the visible part. It is most of the mass. The analogy has one deliberate limit, and it is the hopeful one: the physicist's dark matter is unobservable by construction, while your dark mass is recoverable. It can be surfaced, named, and walked into the light. The rest of this paper is about how, and about when not to.

<p>DARK CONTEXT (noun) /dark 'kan.tɛkst/ · What your organization knows but never wrote down. A general AI language model holds all of the general knowledge and none of yours. That gap is your dark context.</p>

1.1 Scope and claim

The literature on tacit and organizational knowledge is mature [1], [2], [4], [8], [9]. Our aim is not to displace it but to rephrase it for a moment in which the consumer of context changed. We offer a taxonomy (Section 3), a maturity ladder and a seven level scale (Sections 4 and 5), a lifecycle model (Section 6), and an argument about value, cost, and deliberate concealment (Sections 7 to 9). Throughout, the operative claim is that capability has become abundant while context has stayed scarce, and that this inversion is what makes the unwritten part suddenly legible as an asset.

2 Why now: capability became cheap, context became the constraint

Tacit knowledge is not a new concern; firms have worried about losing it to retirement and turnover for decades [4], [9]. What changed is the arrival of a worker with the opposite profile from any human hire. A capable model is pure capability with no situation. It has read more than any employee ever will and yet knows nothing about your Tuesday, your customer, your craft. Give it your dark context and it moves with you. Grounding a general model in your own retrieved text is by now a standard method [10]; the unsolved part is no longer how to condition the model but which context to supply, and even a model handed the right document can miss what sits buried in the middle of a long one [11]. Give it nothing and you get answers that are confident, generic, and slightly wrong: wrong in a manner that looks right, and therefore costly to catch.

For the first time the value of the unwritten part can be read off almost directly. It is the difference between those two outputs, the grounded one and the generic one. Capability is now abundant and roughly fungible across providers; one can be swapped for another in an afternoon. Context is scarce, private, and specifically yours. When the expensive input becomes cheap, the remaining scarce input becomes the constraint, and the constraint here is context.

A WORKED EXAMPLE: THE VALUE IS THE GAP

PROMPT. *Should we take this customer's rush order?*

GENERAL MODEL. Weigh the revenue against the operational strain; rush orders usually justify a premium and can deepen the relationship, so accept if the margin holds. Confident, reasonable, and generic.

WITH YOUR DARK CONTEXT. This is the account Operations flagged last quarter for changing the spec mid-run; the last two rush orders from them lost money and cost a weekend. Decline, or quote the rush premium that made the math work in March.

The distance between those two answers is the value of your context, and you can read it off almost directly. The general model is not less capable; it is less situated. Swap the provider and the first answer barely moves. Add the context and the second one appears.

3 A taxonomy: four kinds of dark

Not all dark context is dark for the same reason, and conflating the reasons is why "let us document everything" so reliably fails. We separate two independent failures. A piece of knowledge can fail to be *captured* at all, or, having been captured, fail to be *accessed* when needed. Crossing the two failures yields four kinds, set out in Table 1.

TABLE 1

The four kinds of dark context, by failure mode and remedy.

Kind	Failure	What it is	Remedy
Unarticulated	Capture	Never said. Silent knowledge that lives in habit and practice, the two second judgment no one has put into words.	Draw it out
Articulated	Capture	Said but not written. Spoken daily in standups and hallways, it never reaches a durable medium.	Catch it
Lost	Access	Written but undiscoverable. It exists in a dead wiki, a departed person's drive, a thread no search reaches.	Find it
Fenced	Access	Written and findable in principle, but behind a permission, a login, or a person who must be asked.	Permit it

The split is practical because each kind has its own fix. Unarticulated knowledge must be elicited; articulated knowledge must be captured before it evaporates; lost knowledge must be made findable; fenced knowledge must be deliberately permitted. A single "documentation initiative" attacks at most one of these and is then surprised by the other three.

4 The maturity ladder: how dark context unfolds in an organization

The taxonomy sorts dark context; the ladder follows it in motion, as a single piece of knowledge travels out of one person's head and into the organization, shown in Fig. 1. It begins *silent*, embodied in a person who acts on it without ever stating it. It becomes *articulated* the first time it is said aloud. (The word does double duty: among the four kinds it named a *failure*, said but not written; here it names *progress*, said at all. It is the same act, scored against a different baseline.) It becomes *documented* when written into a durable place. It becomes *teachable* when shaped so another person can absorb it without the original holder present. And it becomes *shared* when it has spread into the way the organization actually works, held by the group rather than by any one member. Each rung widens the circle that holds the knowledge and loosens its grip on the individual; this is dark context being unfolded, person by person, into the institution. Most organizational knowledge never climbs past the second rung. Surfacing it is the work of moving units up these rungs, deliberately, one at a time, and it is work done by and on people, which is why it carries the human cost we return to in Section 8. The ladder is one axis, human transfer; the next section measures a second and now decisive one, whether the same knowledge is reachable by a machine.




Fig. 1. The maturity ladder. As a piece of dark context climbs from silent to shared, the circle of people who hold it widens and its dependence on any single person falls. Surfacing dark context is the act of moving units up these rungs, out of individual heads and into the organization.

5 The seven levels: machine access to context

Where the ladder asks whether a colleague could take the knowledge over, this scale asks a different question: could a machine reach it, trust it, and act on it? The two share a footing. Level 1, *tacit*, is the ladder's silent rung; level 2 is the same *articulated*; level 3, *recorded*, is its documented. Above that the axes part. The ladder turns toward people, teachable then shared; the seven levels turn toward a machine and add the thresholds it actually needs: *findable* so it can retrieve the knowledge, *current* so it can trust it, *connected* so it can be asked, and *automated* so it can be run. A unit of context is scored on how many of these six thresholds it has crossed, is it *articulated*, *recorded*, *findable*, *current*, *connected*, and *automated*, and stacking them gives the seven levels of Table 2, from level 1, purely tacit, holding none, to level 7, fully machine runnable, holding all. The scale lets a team locate a body of knowledge precisely and name the next move, not the vague "we should document things" but "this sits at level 3, recorded yet not findable, so the next threshold is search."

TABLE 2

The seven levels of context. Each level adds one threshold to the one below it.

Lvl	Name	Threshold gained	Zone	Example
1	Tacit	none (embodied)	Dark	A veteran's feel for a risky deal
2	Articulated	+ said aloud	Dark	Advice given in the hallway
3	Recorded	+ written down	Dark / Gray	A note in someone's private drive
4	Findable	+ discoverable	Gray	A page returned by search
5	Current	+ kept fresh	Green	A swept, dated, trusted playbook
6	Connected	+ queryable	Green	You, or an AI, can ask and get an answer
7	Automated	+ machine runnable	Green	The knowledge executes itself

One row repays a second look. Level 3, *recorded*, still sits in the dark band: writing something down is capture, not access, and a note in a private drive has crossed none of the thresholds that would let anyone else reach it. Recording is necessary for the light but does not by itself bring it, the capture-and-access split of Section 3 reappearing as a single cell.

6 The lifecycle: dark, gray, green

Capturing context is not the end of the story, because context does not stay put. We mark three states, traced as a cycle in Fig. 2. *Green* context is current, findable, and queryable: fresh enough to trust, reachable when you need it, and shaped so a person or a machine can ask it a question and receive a straight answer. *Dark* context is what never reached the light at all. Between them sits *gray* context: knowledge you did capture and once did trust, that has quietly gone out of date. The chart still on the wall, slowly becoming wrong. Gray has a precise meaning here: context that was once green and *decayed*, not context that was captured but never made current in the first place. A page that is findable yet was never swept is gray only by omission, and it is the milder case, because no one has come to trust it, so no one is misled. The dangerous gray is the decayed kind, green that went stale without anyone noticing, and it is the one the rest of this section is about.




Fig. 2. The context lifecycle. Left to itself, green context decays to gray and gray to dark (solid arrows). Tending and re-capture move it back the other way (dashed). Gray is the only state that misrepresents itself: it looks current, so it is trusted.

Gray is the dangerous state, and the claim is sharper than it first sounds: gray is worse than dark. Dark context keeps you honest, because you know that you do not know. Gray context does the opposite. It looks current, so you trust it, and you act on it: a price list six months stale, a playbook from two reorgs ago, an onboarding guide no one has swept. Left alone, green drifts to gray and gray to dark. Keeping context green is therefore not a task that is ever finished but a cadence that is kept, swept at the speed of the operation it describes.

TABLE 3

The three states, and what each returns when you ask it a question.

State	Definition	Ask it a question	Risk
Dark	Never reached the light.	Cannot answer; you know it cannot.	Absence, but honest
Gray	Captured and trusted, quietly out of date.	Answers, wrongly, with confidence.	Silent error
Green	Current, findable, queryable.	Answers in the present tense.	Upkeep cost

7 The new asset: context is priced by no one yet

Context is the situated sense that makes a general system specific to your work: to your Tuesday, your customer, your craft. As capability commoditizes, that situated sense becomes the scarce and differentiating input, and yet it sits on no balance sheet and in no model card. An organization that has surfaced and tended its context owns something a competitor with identical tools cannot quickly copy, because the tools are shared and the context is not.

The unit is concrete: a current, queryable body of context, the green end of the lifecycle, the thing a model is pointed at to become yours. Pricing tends to appear first where that unit already changes hands unpriced. A share of what an acquirer pays for a company is its context, the relationships, the judgment, the way the work is really done, now separable from the tooling because the tooling is everyone's. What lets a team switch AI providers in an afternoon is that capability is fungible; what lets the work survive the switch is that the context moved with them, which is to say the context was the asset all along. A domain that holds context others lack, legal, clinical, industrial, has something to license rather than give away. None of this needs a new marketplace, only someone willing to charge for what is already transferred for free. Today most of it is given away, embedded silently in the outputs people generate, with no one tracking the transfer.

8 The stakes: take the knowledge, lose the knower

There is a cost to extraction that the enthusiasm for "capturing knowledge" tends to skip. When a model, or an organization, learns from a person, something in that person's standing can leave with the knowledge. Part of the expert's leverage was precisely that the knowledge lived only in them.

Surfaced carelessly, the knowledge is transferred and the knower is diminished: the same act that makes the organization smarter can make the individual more replaceable, and they know it, which is exactly why the knowledge was never volunteered in the first place. Any honest program for surfacing dark context has to hold this tension rather than wish it away. The knowledge is not free to take. The real question is whether the person can see a future for themselves on the far side of giving it.

9 When to keep it dark

Not all dark context is a deficiency to be removed. Sometimes keeping something out of the record, away from the next hire, the vendor, the model, is the feature and not the gap. A journalist protects a source. A clinician holds a confidence. A negotiator keeps a reservation price off the page. The discipline is to decide deliberately what stays dark, and never to let a tool make that call by default: a system that ingests everything it can reach will capture exactly the things that were dark on purpose. Surfacing context and protecting context turn out to be the same skill, exercised in opposite directions.

10 Method: getting a group onto the same page

Dark context is surfaced socially before it is surfaced technically [5]. The bottleneck is rarely storage; it is getting a group to agree, in words and at the same time, on how things really work. The instrument we use is MethodKit, a card based method that puts the facets of a subject in front of a group so the unwritten parts get said, compared, and recorded. That it has kept working for more than ten years and across more than one hundred and twenty countries, in governments, companies, and schools, is itself the point: the tools for storing and querying context are rewritten every year, while the human act of articulation stays put, because it is the part that does not automate. It is where the surfacing has to start.

11 Conclusion: where to look next

Capability has become abundant; context is the constraint. We have given dark context a name and a physics, divided it four ways by failure mode, laid it on a seven level scale, and traced its decay from green through gray to dark. The practical reading is simple. Find your highest value dark context, decide deliberately what should stay dark, and move the rest up the ladder before it goes gray. We close with the open question the project itself asks its readers: in your field, where does the unwritten knowledge dominate, and what would it be worth to make it green?

REFERENCES

1. M. Polanyi, *The Tacit Dimension*. New York: Doubleday, 1966.
2. I. Nonaka and H. Takeuchi, *The Knowledge-Creating Company*. New York: Oxford Univ. Press, 1995.
3. G. Ryle, *The Concept of Mind*. London: Hutchinson, 1949.
4. T. H. Davenport and L. Prusak, *Working Knowledge: How Organizations Manage What They Know*. Boston: Harvard Business School Press, 1998.
5. J. S. Brown and P. Duguid, "Organizing knowledge," *California Management Review*, vol. 40, no. 3, pp. 90–111, 1998.
6. V. C. Rubin and W. K. Ford, "Rotation of the Andromeda nebula from a spectroscopic survey of emission regions," *Astrophysical Journal*, vol. 159, p. 379, 1970.
7. F. Zwicky, "Die Rotverschiebung von extragalaktischen Nebeln," *Helvetica Physica Acta*, vol. 6, pp. 110–127, 1933.
8. E. Wenger, *Communities of Practice: Learning, Meaning, and Identity*. Cambridge: Cambridge Univ. Press, 1998.
9. J. P. Walsh and G. R. Ungson, "Organizational memory," *Academy of Management Review*, vol. 16, no. 1, pp. 57–91, 1991.

10. P. Lewis, E. Perez, A. Piktus, *et al.*, "Retrieval-augmented generation for knowledge-intensive NLP tasks," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
11. N. F. Liu, K. Lin, J. Hewitt, A. Paranjape, M. Bevilacqua, F. Petroni, and P. Liang, "Lost in the middle: How language models use long contexts," *Trans. Assoc. Comput. Linguistics (TACL)*, vol. 12, pp. 157–173, 2024.