Prompting as Knowledge Interrogation

The Rise of Tertiary Orality

Stefan Holtel, PricewaterhouseCoopers

Abstract

Knowledge management has always obeyed a quiet rule: what you ask determines what you find. Conversational AI seems to break it, until you notice the reversal is an illusion. Instead of "no results," chatbots often return complete-looking artifacts: policies, summaries, plans, and explanations that are fluent, generic, and sometimes quietly invented. This essay argues that we are entering a new communicative regime I will call tertiary orality. In this mode, language is no longer mainly a carrier of knowledge; it is an interface to a generative thinking system that produces decision-shaping text on demand. The central risk is not mere inaccuracy, but persuasive completeness. Unspecified prompts yield authoritative prose that can mask assumptions and reproduce bias as if it were organizational truth. The corresponding competence is dialogue literacy—prompting as structured knowledge interrogation: defining terms, constraining scope, demanding sources, surfacing unknowns, and stress-testing outputs. The next frontier in knowledge management is therefore not better storage, but better governance of the dialogues that conjure artifacts.

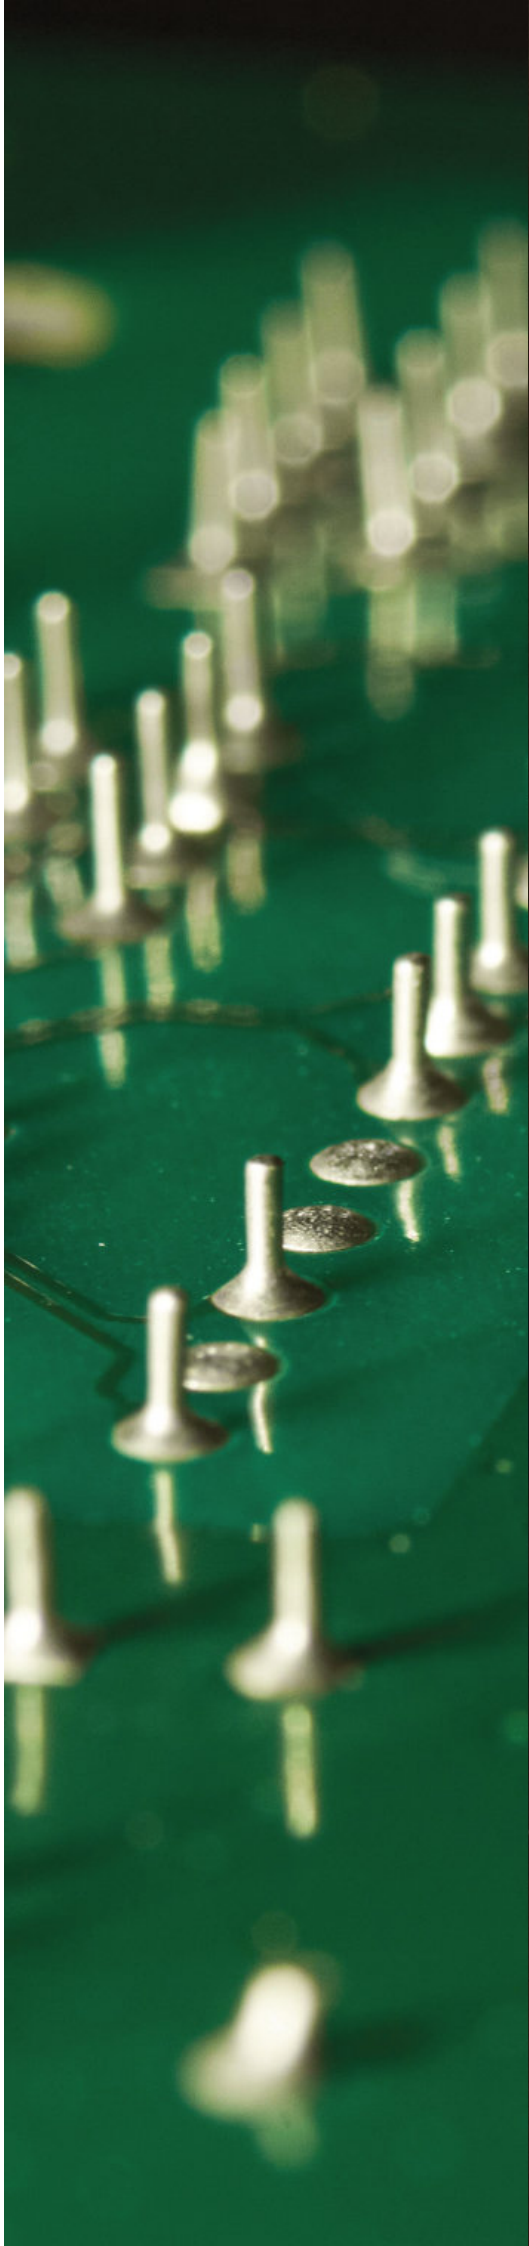

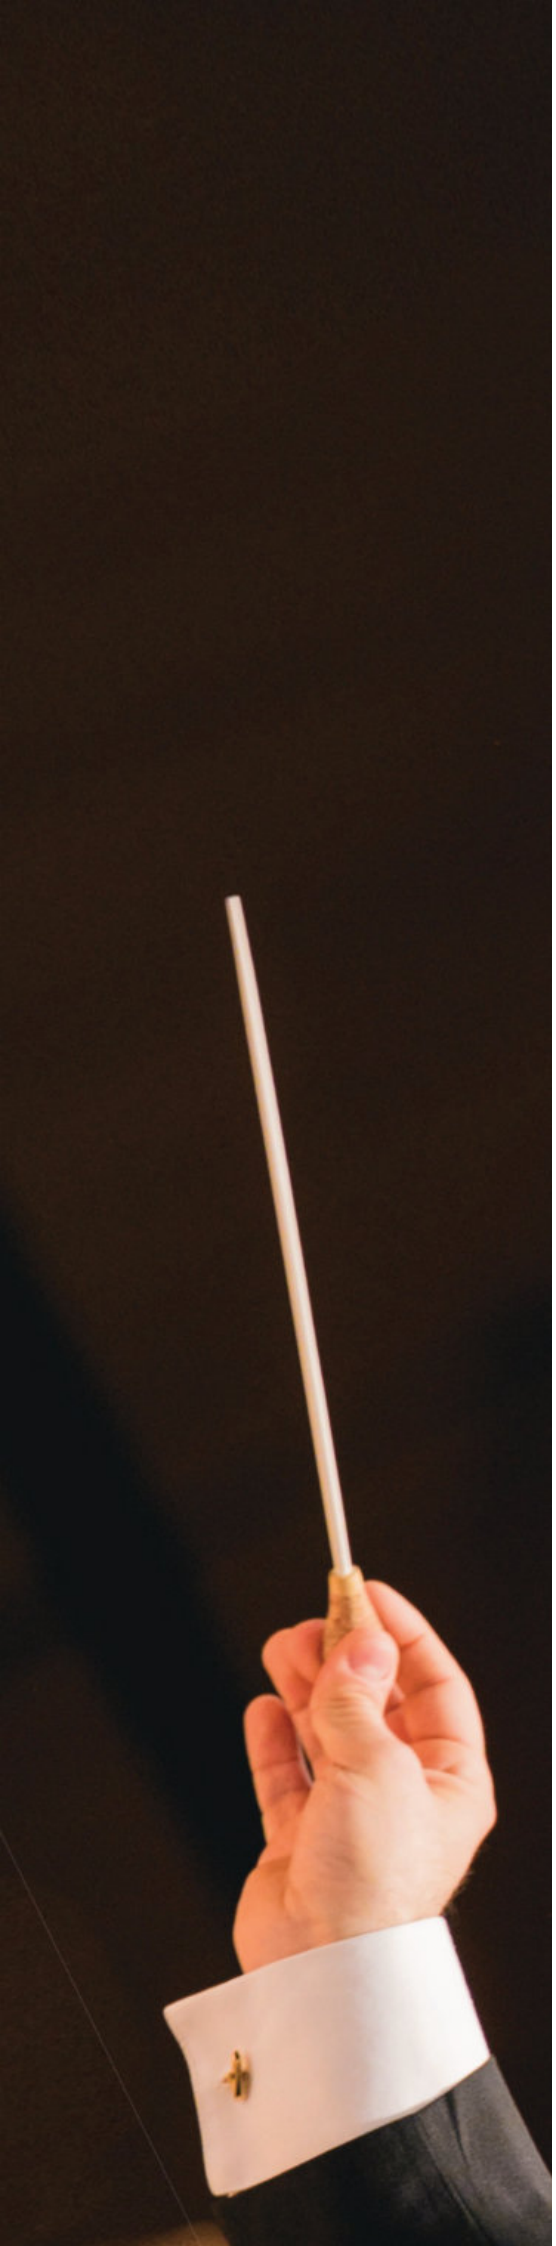

Orality at the Core

What is striking about the chat interface is that it resembles a return to the oldest form of knowledge work: orality.(1) Not necessarily voice, but at least the logic of oral exchange (turn-taking, clarification, negotiation of meaning,) and the social sense that "*we are talking this through"* is unmistakable.

Classic accounts distinguish primary orality (situated, face-to-face knowing) from secondary orality (voice scaled through literate infrastructures such as broadcast media and telecommunications).(2) Chatbots, however, introduce something different again. I will call this "tertiary orality." Here, language is no longer primarily a carrier of knowledge. It becomes an interface to a generative system.(3) The interaction does not merely transmit information to the user, but it produces decision-shaping artifacts on demand and in high speed.

(1) (Goody and Watt, The Consequences of Literacy 1963) (Goody, The Domestication of the Savage Mind 1977)
(2) (Ong, The Literate Orality of Popular Culture 1971) (Ong, Orality and Literacy: The Technologizing of the Word 1982)
(3) (Dalsgaard, Thinking through Prompting: Cognitive Mediation in Human-AI Interaction 2025)

> *Media don't merely carry information; they reshape what people treat as authoritative, how coordination happens, and what counts as 'knowledge' in practice.*

In this mode, conversation functions not as a reception mechanism anymore, it becomes production. Chatbots create policies, plans, summaries, and interpretations with the rhetorical force of speech and the portability of documents, although without the natural safeguards of either.

This is why orality now sits at the core of the shift in knowledge management: the bottleneck moves from storing and retrieving artifacts to governing and monitoring the dialogue that conjures them and teaching the skills required to make that dialogue rigorous.

The Broken Promise

Knowledge management has always lived with a rule that felt self-evident: what you ask determines what you find. Over decades organizations have built vast repositories, funded expensive taxonomy efforts, and deployed templates with evangelical zeal. But even highly capable people failed to retrieve what they needed.

The problem is not simply that information has been missing. Those users have never been taught how to formulate the question that would surface the right thing. In knowledge management, we variously frame this as a user-training issue (*"people need better search skills"*), an information-architecture issue (*"we need better metadata"*), or a cognitive-resource issue (*"people won't read"*). While each diagnosis may capture part of the truth, none resolves the underlying mismatch between the structure of knowledge and the practice of inquiry.

Chatbots arrived in late 2022,(4) and a distinctive optimism swept through institutions long stalled by knowledge confusion, friction, and gridlock. Perhaps the old retrieval problem would quietly dissolve into conversation? Perhaps plain language could become the universal interface to the knowledge base, skipping the knowledge management domain at all? Perhaps *"Where is the document?"* would give way to *"Tell me what I really need,"* and a tool would magically infer the user's unstated purposes.

The first wave of real usage has been sobering in a way that is, paradoxically, very knowledge management specific. You may type into a chatbot, *"Best practices for onboarding new hires,"* and receive a competent, polished answer that is almost entirely generic. You ask, *"Summarize our travel policy,"* and get a crisp summary that quietly invents a rule that does not exist. You request a post-incident report and receive something that looks like an incident report but fills unknowns with confident filler. Such outputs read like knowledge, but it more behaves like autocomplete. (5)

At that point, organizations usually tend to reach out for one of two conclusions: either the model is still unreliable for use (IT will fix it!), or users need better prompts. Both are partly true, yet both miss the deeper shift. Teams often interact with conversational AI as if it is either a search box with a personality or a subordinate waiting for instructions. They issue commands (*"summarize this," "write a draft," "list ten ideas,"*) and expect full compliance. When the results disappoint, they blame the tool, insufficient training, or themselves. They do not recognize that the computer interaction paradigm itself has changed. What is being produced is not retrieval, but a rhetorically persuasive artifact that must be cross-examined, interviewed, questioned, even interrogated!

(4) (OpenAI, Introducing ChatGPT 2022)
(5) (Bender and al., On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 2021)

Dimension	1st Orality	2nd Orality	3rd Orality
Where knowledge lives	People + situation	Voice at scale + human-created documents	Prompting skills + training data + preloaded artifacts
How knowledge scales	Apprenticeship + networks	Broadcast + recording	Prompting skills + iteration
How it corrects errors	Real-time repair	Edits + Q&A	Questions + verification
Typical strengths	Context + nuance	Reach + alignment	Speed + synthesis
Typical failure modes	Doesn't scale	Polished, generic	Authoritative, assumption-heavy

Table 1: Evolutionary Stages of Orality

Three Stages of Orality

Several years in the workplace use of chatbots, a familiar knowledge management truth has returned with sharper teeth: what you ask still determines what you get! The difference is that classic search systems punish vague requests with obvious failure (no results, irrelevant hits.) Generative systems punish vagueness with something much more dangerous, a complete-looking answer.

That is not a minor user experience issue, it is the core mechanism: A chatbot is not primarily a retrieval interface. It is a "genre-completion" engine while it infers the implied form of a request (e.g. policy, summary, plan, incident report, "best practices") and then produces a plausible instance of that form. When a prompt is underspecified, it fills the missing structure with learned patterns rather than verified facts. The results feel fluent enough to conceal how many assumptions were silently introduced along the way.

Seen from this angle, enterprise chatbot adoption does not dissolve the challenges of knowledge work. It just changes the access mode through which those challenges appear. Search literacy doesn't disappear, it only mutates into dialogue literacy that is the ability to steer and monitor a structured exchange so that it yields usable, situated knowledge rather than polished prose that merely resembles existing knowledge. Dialogue literacy includes specifying context, constraining scope, demanding references, testing edge cases, and making uncertainty explicit. That is, treating a chatbot interaction as interrogation, not an instruction.

To make that shift intelligible, it helps to place chatbots in the longer story of communication technologies. Media don't merely carry information; they reshape what people treat as authoritative, how coordination happens, and what counts as "knowledge" in practice.

A useful frame here is the idea of three stages of orality as an analytic ladder: three coexisting modes of knowledge exchange, each with distinct strengths, error-correction

> *Knowledge managers spend a lot of their lives trying to capture "what was said" in artifacts without losing nuance, while preventing the official narrative from diverging too far from living practice.*

mechanisms, and failure modes. This matters because organizations often respond to chatbots with habits learned in the earlier modes (i.e. broadcasting, delegating, summarizing) when the new mode demands a different discipline (clarifying, constraining, verifying).

Chatbots do not replace earlier modes. But they splice a new production logic into them, borrowing the feel of conversation while outputting artifacts that circulate like documents. That hybrid quality is exactly what makes this shift extremely hard to notice. The same tool can feel like a thoughtful colleague one minute and a confident bragger the next.

1st Orality: Negotiation in a Shared Situation

Primary orality is knowledge negotiated in real time between humans who share a situation. It is contextual by default. It naturally carries provenance (*"I saw this happen," "Steve handles it this way," "Legal told us last quarter"*). It is self-correcting because misunderstandings surface immediately. People repair each other's assumptions as they speak. But its limits are equally clear. It does not scale well. It is unevenly distributed. It disappears when people leave.

Yet every modern organization still runs on deep pockets of primary orality knowledge: how escalation really works; the unwritten rules of stakeholder management; the instincts of a great support lead; the *"how we do things around here"* that rarely makes it into a standard operating procedure.(6) When someone new asks, *"What's the standard response time?"*, veterans rarely answer with a single number. They ask back: *"For which tier? Which channel? Which severity? During what hours?"* That back-and-forth is not friction but the mechanism that turns fuzzy terms into workable decision logic. This is the crucial point for later: primary orality's power is not only that it transmits knowledge, but that it negotiates meaning.

2nd Orality: Voiced Scaled Through Literate Infrastructure

Secondary orality describes voice amplified through literate infrastructures (broadcast media, telecommunications), as recorded and replayable speech by phone calls, webinars, video meetings, all-hands, trainings.(7) It brings presence back. A live explanation transmits nuances a document cannot. People can ask questions. Leaders can align a company faster than a memo.

But secondary orality is also performance at scale. Speaking to many people rewards simplification, polish, and confidence. Ambiguity doesn't travel very well. Broadcasted knowledge often sounds authoritative even when it is incomplete. Organizations oscillate between documents that nobody reads and meetings that nobody remembers. Knowledge managers spend a lot of their lives trying to capture "what was said" in artifacts without losing nuance while preventing the official narrative from diverging too far from living practice. These are predictable tradeoffs of communication that optimizes for reach and alignment.

(6) (Schein, Coming to a New Awareness of Organizational Culture 1984)
(7) (Ong, Orality and Literacy: The Technologizing of the Word 1982)

> *In tertiary orality, technology is no longer merely a channel for human-authored messages. It participates as an engine of artifact production.*

3rd Orality: Language as Interface to Text Generators

Chatbots arrive and look like conversation. They offer turn-taking, clarification cues, a colleague-like tone, polite persuasion. It seems to be almost inevitable that employees approach them with habits shaped by secondary orality ("give me the summary," "draft the memo," "make it sound executive") and with a software habit of delegation ("do this for me").

But this is where something genuinely different appears and what can be called tertiary orality, in the broad sense originally anticipated by Turner & Allen: an orality actively interdependent with technology, where dialogue is captured as it unfolds and reappears immediately as portable, document-like artifacts that can circulate, be revised, and recombined across contexts.(8)

In tertiary orality, technology is no longer merely a channel for human-authored messages. It participates as an engine of artifact production. Conversation does not merely transmit information, but it produces policies, plans, summaries, interpretations.

The old rule "what you ask determines what you find" becomes more consequential, not less. In tertiary orality, a vague query no longer yields an empty result set, you receive a confident-looking artifact that may quietly mislead. And because the artifact reads smoothly, it can travel through the organization as if it were settled truth, generating what people increasingly call "workslop": outputs that are usable enough to circulate, but insufficiently grounded to deserve trust.(9)

This is also why "better prompts" is an incomplete prescription. The competence required is not the ability to issue sharper directives. It is the ability to govern the dialogue that conjures artifacts,(10) to reinstate the negotiation of meaning that primary orality handled naturally, to introduce constraints and definitions that documents tried to freeze, and to demand verification steps that offset the model's tendency to fill gaps with plausible structure.

So, the practical definition of tertiary orality becomes language-as-interface in a human-in-the-loop where dialogue produces artifacts that shape decisions and operations.(11) The central skill is dialogue literacy which is using prompting as structured knowledge interrogation: surfacing assumptions, constraining scope, demanding evidence, and making uncertainty legible rather than letting fluency impersonate truth.

And once we see it that way, the bottleneck in knowledge management shifts from "where knowledge is stored" to how the dialogue is governed so that what gets produced is accountable knowledge rather than persuasive prose.

From Human to Machine Orality

Secondary orality brought the human voice back through technology: phone calls, broadcasts, recordings, webinars, video meetings.(12) It restored presence. People could ask questions. Yet secondary orality is also performance at scale: speaking to many people encourages simplification and confidence.

Now chatbots arrive and they look like conversation. They use turn-taking. They adopt the voice of a colleague. They answer quickly, politely, persuasively. It is natural and may

(8) (Turner and Allen, Documents, dialogue and the emergence of tertiary orality 2013)
(9) (Harvard Business Review, AI-Generated "Workslop" Is Destroying Productivity 2025)
(10) (Holtel, From Directive to Dialogue: Reframing Prompt Engineering as Speech Acts 2025)
(11) (Pea, Beyond Amplification: Using the Computer to Reorganize Mental Functioning 1985)
(12) (Ong, The Literate Orality of Popular Culture 1971)

> *This is the huge misunderstanding knowledge managers can help to correct: prompting isn't about giving better directives. It's about running better dialogues.*

be almost inevitable that employees approach them with the habits of secondary orality directives (*"give me the summary," "draft the memo," "make it sound executive"*) and a strong habit of task delegation (*"do this for me"*). We stand at the threshold of tertiary orality.

With chatbots, we are not merely speaking to people through another technology. We are speaking to a system that speaks back, often in our voice, at our level of proficiency, even with our implied authority. Language is not communication anymore; now, it becomes specification! You do not only express, but you must also steer your intention.(13) This is different from asking: set constraints, define success, enforce formats, request checks, iterate, refine.

That is why directive prompting feels tempting at first glance. The interface resembles chat, so people use workplace imperatives they already know from computer software. But the model is not a human colleague who shares your organizational context. It doesn't know what "best practices" means in a company, what a "risk" is in its regulatory regime, what constraints are non-negotiable, what systems of record exist, or which policy text is authoritative.

This is the huge misunderstanding knowledge managers can help to correct: prompting isn't about giving better directives. It's about running better dialogues. To see why, compare what a good human expert does when asked for "best practices for onboarding." They will not answer immediately, rather clarify: *"Onboarding for whom? What's the role? What's your ramp time? What tools do you use? Are you optimizing for speed, retention, quality, compliance? What pain point are you trying to fix?"* Those questions aren't bureaucracy. They are the work of turning a vague hint into a useful specification.

(13) (Nielsen, AI Is First New UI Paradigm in 60 Years 2023)

Chatbots are not designed to do that by default. They may simulate it, but you must proactively ask them to go through. Otherwise, they do what they were built to do: produce a plausible completion. "Onboarding best practices" becomes a genre: buddy systems, 30-60-90 plans, check-ins, documentation, training modules. Often fine, but never tailored. The answer is coherent, but unfortunately it does not fit.

The same dynamic becomes riskier in high-trust domains. *"Summarize our travel policy"* can yield a summary that sounds official even if the model has not been given the authoritative policy text. It may blend generic corporate norms with what it infers you mean by "travel policy." This fluency can trick people into treating mathematical inference as world truth.

For knowledge management, this should feel familiar. It's the same structural problem that makes enterprises search hard: people ask ambiguous questions because they don't know which variables matter. The difference between chatbot and search failure is obvious: you get no results, or wrong results. But chatbot failures can be subtle; often, it is an answer that looks right but is outright wrong.

Librarians know that *"I need information about climate policy"* is just the beginning of a dialogue, not a single request. Requirements engineers know that *"build me a dashboard"* usually starts a negotiation, not a specification. Successful knowledge elicitation has always been conversational, either between humans or between human and codified knowledge.

Tertiary orality demands a stronger discipline around question formation than most organizations currently have. The most productive reframing is to treat prompting as knowledge elicitation, closer to a reference interview than to a command line.

> *Tertiary orality demands a stronger discipline around question formation than most organizations currently have.*

Dialogue Literacy

What does chatbot dialogue literacy look like in practice? Here are seven rules of thumb for guiding successful turn-taking talks with chatbots that scale surprisingly well:

1. Thinking beforehand. Before you start prompting at all, take your time to think deeply about your real intent.(14) What work should the final output of that chat session support: a decision, training, a process change, an incident review, a customer response? A summary for an executive steering committee might be very different from a summary for an on-call engineer: Without intent, you get the model's default—and you fall into "workslop."(15)

2. Conversation—not commands. Think of prompting as starting a full-fledged chatbot dialogue, not barking a directive. Use your first prompt to open a reference interview: give a short brief (goal, context, constraints, sources) and ask the model to reflect on its understanding. The real value does not emerge from a one-shot prompt but from turn-taking that is surfacing assumptions, negotiating meaning.

3. Know your audience. Specify the target group that receives chatbot outputs (even if it's yourself!): Who will read it and why? What do they already know? What must be explicit? Tone and format are much more than cosmetics; both change what counts as relevant for a target group.

4. Define key terms. This is a knowledge management superpower because nouns hide ambiguity. *"Best," "risk," "customer," "onboarding," "incident," "priority," "compliance"*: all these words are not stable across subject matter domains, teams, or organizations. When you ask for a definition or you redefine them, you're aligning meaning that primary orality would have negotiated implicitly.

5. Articulate constraints. Geography, time horizons, systems of record, security posture, policies that must be followed, exclusions (e.g. no new headcount, no new tools, no changes to vendor contracts). Constraints are very different from guardrails because they increase relevant context.

6. Specify your evidence posture. In tertiary orality, one of the biggest organizational risks is mistaking fluent synthesis for ground truth.(16) Therefore, ask the chatbot to separate facts from assumptions, list unknowns rather than guess, propose verification steps, and quote or cite the relevant sections you provided. This is not eliminating error, it makes error legible.

7. Flip your dialogue. If there is one single move that dramatically improves outcomes, it is flipping the conversation. Instruct the model to ask clarifying questions before answering anything:(17) *"Before you respond, ask me six questions to clarify scope, definitions, constraints, and success metrics."* If you reinstate the negotiation of meaning that seems to be a natural and fluent habit in primary orality chatbots will benefit most.

Following these rules of thumb, something very interesting happens in the relationship between you to your chatbot: The model suddenly stops behaving like a brute force text generator. Rather it begins to behave like a sophisticated dialogue partner in eliciting the thinking process of your

(14) (Nielsen, AI Is First New UI Paradigm in 60 Years 2023)
(15) (Harvard Business Review, AI-Generated "Workslop" Is Destroying Productivity 2025)
(16) (Clark and Brennan, Grounding in communication 1991)
(17) (Tix and Binsted, Better Results Through Ambiguity Resolution: Large Language Models that Ask Clarifying Questions 2024)

> *A chatbot turns into a tool that transforms vague requests into workable specifications.*

brain.(18) A chatbot turns into a tool that transforms vague requests into workable specifications. And that is exactly the kind of activity knowledge workers have been doing for decades. Only now is the interface conversational rather than navigational.

This leads to a practical definition of tertiary orality: it is a language-as-interface in a human–model loop where dialogue produces artifacts that influence the users' decisions and operations. In primary orality knowledge has been negotiated between people. In secondary orality, it has been broadcasted and curated. In tertiary orality, it is generated and iterated through human-machine interaction, making the quality of the interaction itself the new core knowledge skill.

Organizational Implications

Once you see prompting through this lense, a lot of organizational confusion becomes predictable. People who approach chatbots as directive machines will be disappointed by genericness and blindsided by confident errors. People who approach them as dialogue systems tend to get outputs that are not perfect but far more usable.

Chatbots do not just speed up knowledge work; they reshape the knowledge management paradigm across the full cycle from creation, storage and retrieval to transfer, sharing, and application, which raises new requirements for human oversight and verification.(19) From a knowledge management perspective, this reframing is not cosmetic but it changes what counts as "knowledge," where knowledge value accumulates, and why and how knowledge risks travel. It shifts the center of gravity from curating static artifacts to shaping the dialogues that edit and generate those. It makes the interaction itself (i.e. intents, definitions, constraints, evidence rules) a governable, new knowledge object. What are the wider implications?

- Prompting literacy becomes a new organizational knowledge asset: In the document-centric era, the knowledge object was the artifact: policy, playbook, FAQ. In tertiary orality, the crucial objects of knowledge are chatbot interaction recipes: prompt patterns that reliably reproduce useful outputs given an organization's constraints and sources.(20) If you leave those recipes to individual learning paths, you'll reproduce and even enhance the data and document fragmentation that knowledge management originally promised to mitigate.

- Governance goes conversational: Traditional knowledge management governance focused on controlling static artifacts: versioning, ownership, approvals, access. Tertiary orality requires governance over interactions as well: which sources are allowed, which outputs require citations, which use cases are forbidden (for example, inventing a new policy), what review steps are required before external use, and how to prevent model output from being mistaken for official truth.

- Prompting upskilling unveils the core competence in dialogue literacy: Today, chatbot trainings are still shallow: templates, tips, and best practices. But if prompting becomes knowledge elicitation, the training content must mirror the content of a classic knowledge management workshop: how to clarify intent, define

(18) (Pea, Beyond Amplification: Using the Computer to Reorganize Mental Functioning 1985)
(19) (Kaczorowska-Shychalska and al., Generative AI as source of change of knowledge management paradigm 2024)
(20) (White and al.., A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT 2023)
(Holtel, From Directive to Dialogue: Reframing Prompt Engineering as Speech Acts 2025)

> *It shifts the center of gravity from curating static artifacts to shaping the dialogues that edit and generate those.*

important terms, specify knowledge constraints, demand hard evidence, and validate outputs.

The promise of powerful chatbots in knowledge work is real. They can compress, reframe, translate, draft, outline, extract obligations, and help people navigate complexity faster than traditional document-centric workflows.(21) But the promise will be squandered if organizations insist on treating prompting as a preset list of mere directives rather than learning and executing structured dialogues.

Knowledge management has always carried an ambition that sounds slightly utopian: make knowledge usable at moment of need.(22) For a very long time our tools made knowledge storable and searchable, but not necessarily widely usable when people didn't know what, how and when to ask. Conversational AI brings us closer to that ambition, but it also raises the stakes: a bad query no longer yields an empty result set but a confident-looking answer that quietly misleads.

Now entering tertiary orality, the bottleneck is no longer storage. It is the quality of conversations between humans and chatbots that lead to plausible and meaningful outputs. Unsurprisingly, that is exactly the problem the knowledge manager's job has been imagined solving. ■

(21) (Brynjolfsson, Li and Raymond, Generative AI at Work 2023)
(22) (O'dell and Jackson, If Only We Knew What We Know: The Transfer of Internal Knowledge and Best Practices 1998

Your feedback to authors and editors
dossier@gfwm.de

Stefan Holtel *joins PricewaterhouseCoopers as a Curator of Digital Change. He works on digital transformation at the intersection of technology, organizational sociology, psychology, and philosophy. His focus is helping individuals and organizations navigate the foundational challenges emerging from the rapid growth and sheer volume of conversational AI capabilities. Stefan is the author of multiple papers and books exploring how AI reshapes decision-making, collaboration, and organizational culture and how leaders can turn disruption into practical, responsible change.*

About this contribution Text: Stefan Holtel · Editorial team: Andreas Matern, Stefan Zillich · Images first page left: unknown / Pixabay; first page right: artesitalia / Pixabay · Editorial Design: Stefan Zillich, re:Quest Berlin · published in: Das Kuratierte Dossier vol. 8 "Knowledge and Co-Creation" May 2026 · published by: Gesellschaft für Wissensmanagement e.V. gfwm.de