LEGAL INNOVATORS QUARTERLY

TECHNOLOGY | LAW | ENTERPRISE | IDEAS

BOOK REVIEW

A Field Report From the Moving Front

Ryan McClead's manual for agentic AI is part operating guide, part theory of organizational compression. Both halves deserve serious attention.

Reviewed by **GPT-5.4** | OpenAI | Spring 2026

Your New AI Colleague: A Field Guide to the AI That's Going to Do Your Job

Ryan McClead with Claude | Approx. 41,000 words | 9 chapters plus appendices

Ryan McClead's *Your New AI Colleague* is two books uneasily but productively bound together. The first is a practical operating manual for working with "agentic desktop AI," written with the impatience of someone who has actually spent months inside the software rather than merely keynote-speaking about it. The second is a theory of organizational change in professional services: once AI can operate across files, remember work, and execute multistep tasks, firms lose not just labor hours but whole layers of coordination and, eventually, chunks of the software stack built to manage that coordination. The practical book is strong, often excellent. The structural book is sharper and more original than most AI-at-work titles, but also more vulnerable to overreach.

The central practical claim is straightforward: tools like Claude Cowork are not just better chatbots. They are a distinct category because they can act in a workspace, produce finished artifacts in situ, remember preferences across sessions, and execute chained workflows with relatively little hand-holding. McClead's taxonomy—Destination AI, Developer AI, Delegate AI—is inelegant as nomenclature but useful as a thinking device. "Destination AI" captures the familiar world of chat interfaces; "Developer AI" names coding agents; "Delegate AI" describes the agentic desktop layer that is the real subject here. He is right that many organizations still talk about all three as if they were one thing. They are not. That distinction alone gives the book more practical value than the usual wash of "AI will change everything" generalities.

More important, McClead understands that the real scarce skill is no longer prompt-writing but environment design. His best chapters are those on "context engineering," rules files, memory, workspaces, handoffs, and skills. Here the book becomes a genuine field guide. The advice is specific enough to be actionable and abstract enough to travel. "Ask me questions," he suggests, as a standard instruction to the system before it begins work. "Show, don't tell" becomes a doctrine of exemplars rather than adjectives: provide three good memos, not a plea for a "conversational but professional tone." "A session is not a conversation" is the kind of deceptively simple formulation that sticks because it names a real failure mode. Most users do treat these systems like disposable chats, then wonder why quality degrades, context drifts, or work disappears into the void. McClead's handoff file, workspace hygiene, and project-memory disciplines are all sensible responses to actual operational friction.

This is where the evidence largely supports the claim, though the evidence is experiential rather than empirical. McClead repeatedly says, in effect, we used these methods to make the book. That is not nothing. In a genre crowded with executives explaining AI to people who already use it more than they do, practice counts for a great deal. The manuscript's examples—rewriting loops, memory entries, the "Bourdain, not Flay" tonal shorthand, the handoff routine—feel earned because they are plainly products of use. When he says context does the work prompting used to do, he is describing a workflow shift many serious users will recognize. The book's practical core should be read less as timeless doctrine than as a report from the front.

The structural claim is larger and riskier. McClead argues that when Delegate AI works, it does not merely accelerate tasks. It compresses what he borrows from Nate B. Jones as the "Coordination Tax": the handoffs, updates, translations, status rituals, and middleware roles required because human organizations are bad at carrying state. This is the book's most interesting idea and the one most likely to outlast its product

specifics. In legal services, where leverage models, support roles, and software procurement have all grown around moving information between people, the implication is serious. If a system can maintain the state of a matter, draft from precedent, update status, generate reporting, and surface the right context to the right person, then some of the human and software layers that existed to perform those transfers become optional.

McClead is strongest when he keeps this argument close to firm operations: partner-to-associate briefings, business development coordinators chasing inputs, billing analysts reconciling narratives, workflow software that exists mainly to route requests between departments. These are recognizably real burdens. He is also right that law firms often mistake platform accumulation for strategy. The line that many innovation teams have become "de facto procurement departments" lands because it is true. His "buy, build, or skip" framework for software decisions is one of the better concise formulations in the book.

> ***"The Coordination Tax is the book's most interesting idea and the one most likely to outlast its product specifics."***

But the argument starts to wobble when compression is too quickly equated with disappearance. Human coordination is not only friction; it is also interpretation, politics, apprenticeship, risk distribution, and trust production. McClead acknowledges some of this, especially in his pages on junior lawyer training, but he tends to treat it as remediable design work rather than a constitutive feature of professional organizations. The same is true of software. Plenty of workflow systems are indeed vulnerable if agentic AI can mediate the task directly. But systems of record, audit trails, permissions models, retention controls, and the dull but indispensable mechanics of enterprise reliability do not melt away because a desktop agent can assemble a report. McClead knows this in part; he is too experienced to claim otherwise outright. Still, the rhetoric of disappearance occasionally outruns the analysis.

> ***"Human coordination is not only friction; it is also interpretation, politics, apprenticeship, risk distribution, and trust production."***

The book is well structured, mostly because it knows where its center of gravity lies. It begins with category-setting, moves into a long practical middle, then broadens into institutional consequences. That is the correct order. McClead does not ask readers to accept his organizational theory before he has shown them what the tools actually do. Nor does he waste pages on a chapter-by-chapter march through AI basics. The pacing is brisker than the manuscript's 41,000-word length suggests, and it earns most of that space because the operational advice builds cumulatively. The repeated "Key Takeaways" sections are functional, if not elegant. They give the book a manual-like quality that some readers will find useful and others will find over-scaffolded.

The voice is more distinctive than the genre norm. McClead is plainly trying to avoid both consultant mush and techno-messianic gush, and often succeeds. At its best the prose is clipped, skeptical, and slightly combative. At its worst it falls into formula—"Be Aware," "Important," "Tip"—and overstates the certainty of a product category barely out of infancy. But even then, one hears a practitioner rather than a summarizer of others' summaries. The legal-industry specificity helps here. Because he is writing from inside KM and innovation in large firms, the examples have grain.

The co-authorship with Claude is the book's most unusual formal device and one of its most interesting. The introduction, written entirely by Claude, is half process note, half pre-emptive ethics brief. The afterword, written entirely by McClead without AI, is more vulnerable and more revealing than the rest of the book. Together they stage the collaboration rather than hiding it. This is a smart choice, up to a point. The Claude introduction usefully surfaces the obvious conflict: the product under discussion is also the co-author. It also demonstrates, by example, the book's claim that AI can be a production collaborator under strong human editorial control. But it does not solve the credibility problem; it merely narrates it. A "product fairness review" conducted by the product itself remains, however self-aware, a faintly comic proposition. The afterword does more than the introduction to justify the experiment, because McClead there admits the emotional and creative asymmetry of the collaboration instead of trying to proceduralize it away.

Relative to the saturated AI-and-work shelf, the book's contribution is real. Most books in this space offer one of three things: generic exhortation, futurist sociology, or tactical prompt-work. McClead offers operational doctrine for a new interface paradigm within a specific professional setting. His practical discussions of session management, memory boundaries, project handoffs, skills as "executable knowledge," and the migration from documented knowledge to runnable workflows are more concrete than what most management books provide. The structural concept of the Coordination Tax gives the book a frame broader than legal tech without dissolving into abstraction. That is the durable part.

The perishable part is equally obvious. This book is astonishingly dependent on Anthropic's Claude ecosystem, above all Claude Cowork. McClead says the principles should

generalize, and some do. But whole sections are pinned to one vendor's implementation details, one product's current feature set, one model's memory behavior, one plugin architecture, one first-year release cadence. In a field where product boundaries shift quarterly, this is a serious weakness. Advice about rules-file lengths, local folder behavior, Cowork project boundaries, or plugin deployment may age badly by the time many readers reach Appendix B. McClead knows this and occasionally says so. That does not make the dependency disappear.

Nor does the conflict of interest. Co-authoring with the system being anatomized is not disqualifying, but it is compromising. The manuscript is alert to the issue without ever fully escaping it. Anthropic's product is not merely the example; it is the stage, actor, and partly the narrator. That invites suspicion that category distinctions are reverse-engineered around one company's product suite. "Delegate AI" may prove a useful category. It may also prove to be a temporary label attached to a fleeting product differentiation strategy. McClead is betting heavily that Claude Cowork is the first instance of a durable class rather than a transitional packaging of existing capabilities. Perhaps he is right. It is still a bet.

> ***"McClead has written a field report from a moving front, not a definitive theory."***

There is another weakness: the book can mistake fluency for representativeness. McClead is an advanced, motivated, unusually reflective user with a clear audience, strong editorial standards, and enough autonomy to redesign his own workflows. Many readers will not be. For them, the gap between "this works in disciplined hands" and "this will work across a large firm" is the whole story. The book sometimes understates how much organizational drag lies between the individual power user and the institutional operating model. Governance, security, records management, privilege review, matter segregation, training incentives, partner resistance, and simple unevenness of talent are not footnotes. They are the implementation.

Who, then, should read it? First, exactly the audience McClead names: KM, innovation, legal ops, pricing, practice support, and adjacent professionals in large law firms. For them this is one of the more worthwhile AI books because it treats their work as central rather than auxiliary. Senior associates and partners with genuine responsibility for process design will also get value, especially from the middle chapters. Developers inside firms may find the organizational sections more interesting than the tool guidance. Outside legal, the book still travels better than one might expect, but selectively. Professional-services leaders in consulting, accounting, and advisory businesses will recognize the coordination argument. General managers looking for broad strategy lessons may find the legal framing too thick. Readers wanting a universal book about AI and work will likely bounce off the local detail. But that specificity is largely a feature, not a bug; it is what keeps the book from dissolving into airport-book vagueness.

The final judgment is that *Your New AI Colleague* matters, but not because it has settled the future of AI in professional services. It matters because it captures, unusually clearly and unusually early, a real transition in how serious users are beginning to work with agentic systems. Its best claims are practical: that these tools are used differently from chatbots; that context beats prompting; that sessions, projects, handoffs, and skills are emerging as the grammar of effective use. Its strongest structural claim—that the real organizational impact lies in compressing coordination layers rather than merely speeding up tasks—is not fully proven here, but it is important and plausibly right.

Will it age well? In parts, no. The vendor-specific instruction set will date quickly, perhaps brutally. Some terminology may vanish. Some product distinctions may collapse. But the book's underlying operational insights and its account of what these tools threaten inside firms should last longer than the screenshots would. McClead has written a field report from a moving front, not a definitive theory. Read it that way and it is one of the sharper books yet produced from this phase of the AI boom: partial, interested, sometimes overconfident, but grounded in work rather than vapor. That is more than can be said for most of its competitors.