SPOTLIGHT

The end of a document-centric world: why AI is forcing a rethink of medical writing

Artificial intelligence's greatest impact on medical writing may not be faster drafting. Instead, it may be the shift from treating documents as individual deliverables to managing a connected narrative – helping teams maintain consistency, traceability, and scientific integrity throughout the clinical development life cycle

Lisa Chamberlain James,
VP Medical Solutions
at Trilogy Writing & Consulting,
an Indegene Company

For decades, the pharmaceutical industry has thought about medical writing in terms of documents. Protocols. Clinical study reports. Investigator brochures. Common technical document (CTD) summaries. Safety reports. Plain language summaries. The list goes on. Each document has traditionally been viewed as a separate deliverable: owned by a particular team, developed at a particular point in the programme, reviewed through a particular process or by a particular review team. These review processes and/or review teams can differ across the deliverables.

That approach made sense when content was largely produced in sequence, systems were less connected and the main challenge was getting a high-quality document over the line. But clinical development no longer works in that way. Programmes are faster, evidence changes continually, stakeholders expect information in different formats and the same scientific messages are reused across an ever-growing number of outputs. The industry is however still organising work around documents, even though the knowledge inside them is increasingly interconnected. Yet patients, regulators and healthcare professionals do not experience those outputs as a collection of individual documents. They experience a story: what was studied, why it matters, what the data shows, where the uncertainties remain and what those findings mean in practice. If that story changes subtly from one document to another, trust can be weakened even when every individual sentence is technically defensible.

Increasingly, that end-to-end story is becoming difficult to manage. The true challenge facing modern medical writing is therefore not simply writing faster – artificial intelligence (AI) is rapidly helping us to do that. It is maintaining consistency, accuracy, context and scientific alignment across an expanding ecosystem of content, while still meeting the needs of very different audiences. Speed without alignment merely allows inconsistency to travel faster. AI is not simply changing how we write documents. More surprisingly, it has made us to question whether a document-centric model is still fit for purpose.

The hidden cost of document-based thinking

Medical writers understand better than anyone that a single clinical programme can generate thousands of pages of content. The same efficacy data may appear in a protocol, clinical study report, briefing document, publication, CTD module and patient-facing communication. The same safety narrative may be reviewed, revised and reproduced dozens of times throughout a product's life cycle. On paper, this can look like efficient reuse. In reality, it often involves copying, adapting, checking and rechecking content across disconnected files and systems. Every update creates an opportunity for inconsistency. Every version creates a risk of divergence. Every new document becomes another point of maintenance and every handover introduces the possibility that context will be lost. The cost is not limited to writing time. It appears in review cycles, reconciliation exercises, duplicated quality-control effort, delayed decisions and the cognitive burden placed on teams that must remember where a message has appeared and whether it is still current.

As programmes evolve, endpoints are refined, population analyses can change or be refined and regulatory feedback is incorporated. A sentence that was accurate six months ago may now need qualification. A table may be updated while the surrounding interpretation remains untouched. A preferred term may change in one



workstream but persist in another. The challenge is no longer a shortage of content; it is keeping content aligned with the current evidence and with itself. Many organisations unknowingly spend enormous effort manually reconciling information that already exists elsewhere. Highly skilled medical writers become human search engines, version historians and translators between teams. That work is essential, but much of it is also a symptom of a model that asks people to compensate for fragmented information architecture.

The problem of narrative drift

Perhaps the greatest hidden risk is the narrative drift. A programme begins with a relatively clear scientific story: the unmet need, the rationale for the intervention, the evidence being generated and the claims that evidence can reasonably support. Over time, that story is stretched across multiple authors, functions, teams, reviews and document iterations. One document describes an endpoint one way. Another uses slightly different terminology. A third introduces an interpretation that was once plausible but now no longer fully aligns with the current evidence. Meanwhile, audience needs legitimately shape the way information is presented, so variation is not always wrong. The difficulty lies in distinguishing purposeful adaptation from accidental inconsistency. Without a clear source of truth and strong narrative governance, that distinction becomes harder to see.

None of these shifts may be individually significant. Collectively, however, they can create confusion, inefficiency and regulatory risk. They can also make review unnecessarily difficult, because reviewers spend time debating wording rather than testing whether the scientific argument is coherent. Medical writers have always been the custodians of narrative consistency, but traditional tools provide limited support for maintaining alignment at scale and, increasingly importantly, at speed. We rely on experience, memory, and careful review to connect the dots. Those capabilities remain vital, but they should be strengthened by systems that can surface interconnectivity, flag divergence and show the provenance of a message before a document reaches final review.

Why AI changes the conversation

Many discussions about AI focus on drafting documents faster. That is understandable: drafting is visible, easy to measure and immediately attractive when timelines are under pressure. Yet speed is arguably the least transformative aspect of the technology. Producing a first draft in minutes rather than days is useful, but it does not solve the larger problem if the

underlying messages are inconsistent, the source evidence is unclear or the output creates more work for reviewers. The real opportunity lies in shifting from document management to narrative management: using AI not just to generate language, but to understand how claims, data, interpretations and audiences connect across a programme.

Instead of treating documents as isolated outputs, organisations can begin managing content as connected knowledge. That means identifying the core scientific messages that should remain stable, the evidence on which they depend, the approved language surrounding them and the contexts in which they may need to be adapted. A regulator, investigator, healthcare professional and patient will not receive identical wording, and nor should they. But each version should be traceable to the same evidence base and remain consistent with the same scientific position. Connected content does not mean bland uniformity; it means controlled, intentional variation where necessary and appropriate.

In this model, scientific messages exist independently of any single document. Evidence is linked directly to its source, key statements are associated with the data and decisions that support them, and changes can be assessed across connected content. If an analysis population changes, the system can help identify where that change matters. If a safety interpretation is refined, related narratives can be reviewed together rather than discovered one file at a time. Consistency becomes systematic rather than dependent on a heroic final-stage reconciliation. This is not about removing judgement from the process. It is about giving expert judgement better visibility and a more reliable foundation.

The role of AI, then, is not simply to generate text. It is to help manage the relationships between information, evidence and narrative. It can compare language across outputs, identify where a message has drifted, retrieve the material relevant to a specific decision and support impact assessment when evidence changes. Used well, it becomes part of the connective tissue of content development. Used poorly, it risks adding another layer of ungoverned text to an already overloaded environment. The difference will depend on the quality of the underlying content model, the strength of governance and the willingness of organisations to redesign processes rather than automate outdated ones.

Towards connected content ecosystems

Imagine a future where a key efficacy finding is updated once. The change automatically identifies every location where the message appears, from a CTD summary to a publication draft and a patient-facing explanation. Relevant content is flagged for expert review, associated documents are refreshed where appropriate and traceability to the source evidence remains intact. A medical writer can see not only the latest wording, but why it changed, who approved the scientific position and which outputs may be affected. Quality checks occur throughout the process rather than being concentrated at the end. Instead of recreating and revalidating information repeatedly, teams refine and govern a living body of scientific knowledge. That does not mean every update should flow through automatically. Some changes will require interpretation, discussion or deliberate differences between audiences. The value is that those choices become visible and controlled, rather than hidden in a maze of documents.

The concept may sound ambitious, but many industries have already moved towards structured and modular content models. Development in the pharma industry is increasingly positioned to do the same, although the transition will not be achieved by purchasing a single point solution. Organisations will need to agree what constitutes a core message, how evidence is linked, who owns narrative decisions and how exceptions are handled. Legacy content will need to be understood, not simply migrated. Review and approval processes will need to evolve. Above all, teams will need to resist the temptation to bolt AI onto the existing document factory and call it transformation. Technology can accelerate change, but only re-evaluating our operating-models will make that change meaningful.

What this means for medical writers

Far from diminishing the importance of medical writers, connected content ecosystems make their expertise even more valuable. The profession is evolving from document creation to narrative stewardship. Medical writers understand how evidence becomes an argument, how wording changes meaning, how uncertainty should be expressed and how the same science must be adapted responsibly for different audiences. Those are not peripheral skills; they are the skills that make technology safe and useful. Someone must define the narrative architecture, recognise when consistency is essential, know when variation is justified and challenge language that is not supported by the evidence. Medical writers are exceptionally well placed to do that.

Future medical writers are likely to spend less time constructing individual paragraphs from scratch and more time defining core scientific messages, curating reusable content, ensuring narrative integrity and managing content governance. They will continue to review AI-generated outputs, because it remains crucial to have the ‘human in the lead’, not merely somewhere in the loop at the end. But their remit will broaden. They may help to design rules for appropriate reuse, determine what evidence an AI system is allowed to draw upon, assess where automated comparisons are meaningful and identify where context makes a superficial match misleading. They will also be responsible for maintaining context consistency across programmes, as well as within them. In this environment, success will depend less on raw document production speed

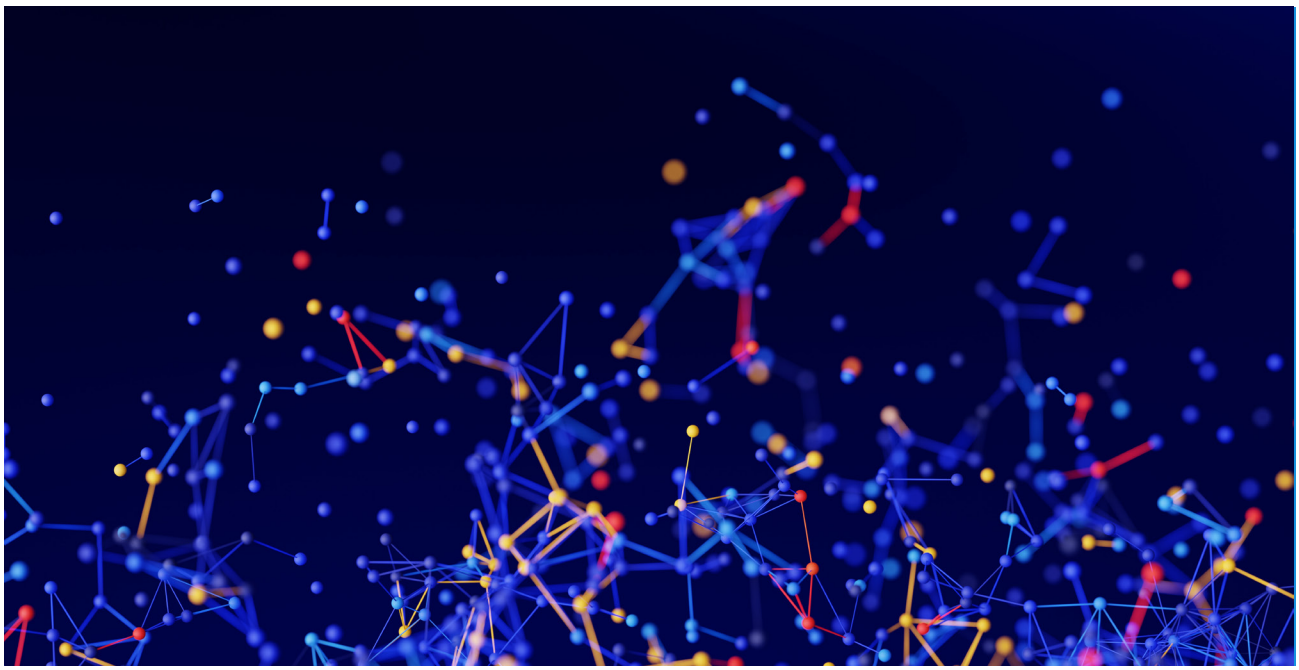

and more on strategic communication expertise, scientific judgement, and the confidence and skill to govern how a story evolves.

The future is not faster documents

The temptation is to view AI as a productivity tool. There is nothing wrong with productivity: medical writing teams are under genuine pressure and removing repetitive work matters. But productivity alone is not transformation. If AI is used solely to produce more documents at a faster pace, it risks exacerbating the existing fragmentation that organisations already struggle to manage. The greater opportunity is to rethink how knowledge is organised, reused, challenged and governed throughout the clinical development life cycle. That means starting with the scientific narrative and the evidence that sustains it, rather than starting with a blank document and rebuilding the story each time.

For years, the pharma industry has focused on creating better documents: clearer templates, faster authoring, more efficient review and stronger quality control. All of that remains important. But the next decade will be defined by our ability to create better narratives – narratives that remain coherent as evidence develops, that can be adapted without losing their scientific integrity and that are traceable from source to audience. This will require new technology, certainly, but also new habits, roles and measures of quality. Evaluating the health of the entire content ecosystem, rather than focusing solely on the quality of the final document, will become increasingly important. As this shift takes place, medical writers will play a central role in one of the most significant transformations the medical writing profession has experienced. Their value will lie not in producing the greatest volume of content in the shortest time frame, but in understanding how that content must connect evidence, meaning and people while maintaining scientific accuracy and integrity.

The future is not faster documents. The future is safeguarding regulatory narratives.



Dr **Lisa Chamberlain James** is VP Medical Solutions at **Trilogy Writing & Consulting, an Indegene Company**. Aside from management activities, she leads client projects, with extensive experience in a variety of documents. Lisa has a special interest in writing for the public, and in the effective and ethical use of AI in medical writing. Following a PhD and postdoc in Pathology at Cambridge, UK, Lisa began her medical writing career in 2000. Since then, she has been involved in the European Medical Writers Association (EMWA) as a member of the Educational Committee, mentor, leader, and assessor of workshops, and is a member of the American Medical Writers Association (AMWA)'s Executive Forum, AI Task Force, and teaches and reviews workshops for AMWA. Lisa regularly takes part in podcasts, articles, and is a member of several global, pan-industry panels and committees, holds an EMWA professional development certificate, is a visiting lecturer for King's College London and the University of Lisbon, and is the current chair of the Communicating with the Public SIG, and the Geoff Hall Scholarship Committee. She is a section editor of *Medical Writing*, an EMWA Nick Thompson Fellow, and a Fellow of the Royal Society of Medicine.