



A.T.L.A.S.

AI Tool Literacy, Adoption & Selection Guide.

For pharmaceutical teams

A practical framework for choosing AI tools with confidence. Steers you through data risk, evaluation criteria, team readiness, and how to build a credible business case for IT.



New adventures in medcomms

caminocomms.com

How to use this guide

AI tools can genuinely transform how pharma teams work. This A.T.L.A.S. helps you choose the right ones, adopt them well, and build confidence in the process.

It works for two common scenarios. Find yours and use it as your starting point.

SCENARIO A: Someone's pitched me a tool and I want to know if it's worth pursuing

A vendor demo, a recommendation at a conference, a LinkedIn post. Before you invest time or raise it with IT, you want a quick sanity check. Is this tool worth exploring further? Would it actually work in your environment?

Start with the 'Data risk spectrum' (page 2), then 'Evaluation criteria' (page 4), then 'Asking the right questions' (page 5) to look past the pitch. Finish with 'People and readiness' (page 7) for adoption planning.

SCENARIO B: I've got a use case and a shortlist. I need to compare before going to IT.

You've identified a genuine need and found a few options. Now you need to compare them properly so you can make a credible recommendation, not just "this one seems good."

Start with 'Evaluation criteria' (page 4) and score each tool. Use 'Workflow fit' (page 6) for integration, and 'People and readiness' (page 7) for adoption. Before any of that, check what you've already got (page 5).

The data risk spectrum

The good news: most of the AI use cases that pharma teams start with involve lower-risk data, like published literature, congress content, and product prescribing information. There's a lot you can do safely with general-purpose tools if you understand where your data sit.

Your internal IT and compliance policies are the final word. This spectrum gives you a starting framework for the conversation.

Lower risk Public / published	• Published literature • Congress abstracts and posters • Public prescribing information • Non-proprietary market data • General medical/scientific knowledge	<i>Wide range of tools appropriate. Most general-purpose AI can be used for summarisation, analysis and exploration. A great place to start building confidence. Check your company's approved list first.</i>
Moderate risk Internal / proprietary	• Internal strategy documents • Unpublished data or analyses • Competitive intelligence • Pre-launch planning materials • Internal meeting notes and insights	<i>Enterprise-grade tools with data agreements. Look for no-training clauses, data residency guarantees and your company's enterprise AI platform. Free-tier consumer tools are unlikely to be appropriate here.</i>
Higher risk Sensitive / regulated	• Patient-level clinical data • Adverse event reports (pre-submission) • Personal identifiable information (PII) • Commercially sensitive financials • Unblinded trial data	<i>Only tools explicitly approved by IT security and compliance for this data class. This is specialist territory, but the right tools do exist and are maturing fast.</i>

KEY POINT: The same tool might be fine for one data type and need more caution with another. Classify the data first, then match the tool. This makes every conversation with IT much easier.



Start with the use case

“We want to use AI” is not a use case. This is: “We want to reduce the time it takes to summarise literature for monthly therapeutic area updates.” The more specific you are, the easier it is to choose the right tool and make a case for it.

Where teams are getting real value

These are the kinds of tasks where AI tools are already proving themselves in pharma:

- **Drafting and summarisation:**
First drafts of emails, slide decks, meeting summaries. Not finished products, but solid starting points that save hours.
- **Literature work:**
Monitoring, synthesis, and summarisation of published content – turning a stack of abstracts into a structured overview.
- **Content repurposing:**
Adapting content for different audiences or channels – a congress report into an internal summary, for instance.
- **Data exploration:**
Asking questions of datasets in natural language. Useful for spotting patterns and generating hypotheses.
- **Review and quality assurance:**
Using AI as a second pair of eyes on draft content – checking consistency, tone and compliance alignment.

Higher-stakes tasks like HCP-facing content or regulatory submissions need more scrutiny, but that doesn't mean AI can't help. It means the review process needs to match the stakes.

ASK YOURSELF: If this AI output is wrong, what are the consequences – a wasted afternoon or a compliance incident? That answer shapes how much review you build in.



The 80:20 principle

AI doesn't need to be perfect to be useful. If a tool gets you 80% of the way to a good first draft in 20% of the time, that's valuable. The remaining 20% is where your expertise and review process come in. That's not a flaw in AI. That's how it's supposed to work.

Don't let the pursuit of a perfect AI strategy stop you from getting started. Start small, learn fast, build from there.

Evaluation criteria

Whether you're assessing a single tool or comparing several, run each one through these criteria. For Scenario B, this works as a scorecard for side-by-side comparison.

Criterion	What to look for	Watch for
Data processing	Where is data processed? Stored? Used to train models? Look for no-training clauses and DPAs.	No clarity on data handling. Free-tier tools that train on inputs. Vague privacy policies.
Data residency	Where are servers? Compliant with your data governance? EU-hosted for EU data?	Conflicting jurisdictions. No option to choose region.
Security and access	SSO/SAML? Role-based access? Audit logs? SOC 2 or ISO 27001?	Shared logins. No audit trail. No certifications.
What's under the bonnet?	What foundation model powers it? Is the vendor transparent? (See next section)	Claims of 'bespoke AI' without evidence. Won't disclose the underlying model.
Pharma-specific fit	Does the vendor understand pharma? Compliance features? Existing pharma clients?	No pharma awareness. Marketing pitched only at tech audiences.
Output quality	Test on real tasks. Accurate? Handles medical terminology? Consistent across runs?	Impressive demo but inconsistent on real tasks. See 'jagged frontier', next section.
Integration	Works with Microsoft 365, Veeva, SharePoint? API available?	Standalone with no integration path. Would create a new silo.
Cost and scalability	Per-user pricing? Enterprise tiers? Hidden costs (tokens, API)?	Opaque pricing. Per-query charges. No enterprise tier.
Vendor viability	Established or startup? Funding? Support? Roadmap?	Very early-stage. Limited support. May not exist in 12 months.

API, application programming interface; DPA, data processing agreement; ISO, International Organization for Standardization; SAML, Security Assertion Markup Language; SOC, System and Organisation Controls; SSO, single sign-on.

FOR SCENARIO B: Print this table and score each tool on your shortlist. It makes the recommendation to IT much more credible than "I saw a demo and it looked good."



Asking the right questions

Check what you've already got

Before you go looking for a new tool, check what's already available. Many pharma companies have Microsoft 365 Copilot licences, and most are barely being used. Copilot can summarise documents, draft emails, pull insights from Excel data, catch up on Teams meetings, and more.

It won't do everything. But if your team hasn't properly explored what your existing enterprise AI platform can do, you might be shopping for something you've already paid for. The gap isn't always the tool. Sometimes it's the training and support around it.

Understand what's under the bonnet

The AI landscape is full of companies that have built a useful interface on top of someone else's foundation model (usually GPT from OpenAI, or Claude from Anthropic) and are packaging it for a specific audience. That's a perfectly valid approach and plenty of these tools are genuinely useful.

Where it gets tricky is when a vendor implies they've 'trained their own AI model' or built something 'bespoke'. Training a foundation model from scratch costs hundreds of millions and requires world-class infrastructure. If a small company is claiming bespoke AI, it's worth understanding what they actually mean. They may have done smart work with fine-tuning or retrieval-augmented generation (RAG), and that's genuinely valuable. Or they may have wrapped an application programming interface (API) call in a branded interface. Knowing the difference helps you assess what you're paying for.

Good questions for any vendor

- What foundation model powers your product?
- What have you built on top of it (fine-tuning, RAG, integrations, custom workflows)?
- Where do my data go? Do they pass through a third-party provider?
- If the underlying model provider changes their terms, how does that affect your product?
- Can I trial this on my own data and tasks before committing?

OUR TAKE: There's nothing wrong with a tool built on GPT or Claude. What matters is transparency and genuine added value. If a vendor is charging a premium for what is essentially a prompt template over someone else's API, you can probably do that yourself.



Test on your terms

Ethan Mollick at Wharton coined the term 'the jagged frontier' to describe the unpredictable, uneven capabilities of AI. A tool can perform brilliantly on one task and stumble on something that seems similar. This matters because **a polished demo tells you very little about day-to-day performance.**

The best way to evaluate a tool is to test it on your own tasks, with your own data, across multiple runs. Ask the vendor for a proper trial. If they're confident in their product, they'll welcome it.

How does it fit your workflow?

A tool can tick every technical box and still fall flat if it doesn't fit how your team actually works. These questions help you think through the full journey of an AI output.

1. How do outputs enter your review process?

AI-generated content should enter your existing medical/legal review (MLR) or content review workflow at the same stage as any other draft. If your current workflow doesn't account for AI-assisted content, that's a straightforward gap to close before scaling up.

2. Who reviews the output, and against what?

Define who reviews AI outputs, what they check them against (source documents, approved claims, prescribing information), and how they document that review. AI-generated references need to be verified individually. This isn't unique to AI; it's good practice for any content.

3. Is AI use visible in the approval chain?

Do your approvers know when content has been AI-assisted? Transparency builds trust and helps your team learn what's working. It also means fewer surprises down the line.

4. Are you keeping a record?

A simple log of tools used, prompts given, and review steps gives you traceability. If anyone asks how content was produced, you can give a clear answer. That's a strength, not an overhead.

5. Could you manage without it?

If the tool disappeared tomorrow, would your team be stuck? Building in a fallback isn't pessimism; it's good planning.

What good looks like

When a team is using AI well, you can tell. The work is better, the process is faster and nobody's cutting corners.

- **AI outputs are starting points, never endpoints.**

They go through the same review as any other draft. Nobody's copying and pasting raw output into a deliverable.

- **References are verified.**

Every citation is checked against the actual source. AI models can fabricate plausible-looking references with real-sounding journal names and DOIs. Good teams check every time.

- **The human voice is still there.**

The final product reflects the team's expertise, not the default tone of an AI model. If everything starts to sound the same, AI is replacing thinking rather than supporting it.

- **Speed comes with scrutiny.**

The time saved on drafting gets reinvested in review and refinement. Faster doesn't mean less rigorous.

- **People know their boundaries.**

Users are clear on what data they can and can't input, which tasks are appropriate, and when to escalate. That clarity comes from good training, not from locking everything down.

A word on slop

Slop is AI-generated content that takes more effort to read than it took to produce. When someone sends you raw AI output without editing it, they're not saving time. They're shifting the work onto you.

A tool that generates rough output isn't necessarily a bad one. But a team that ships rough output has a process problem. The fix is usually better review habits, not a different tool.

THE TEST IS SIMPLE: would you put your name on this? If yes, send it with confidence. If not, it's not finished.



People and readiness

Tools don't adopt themselves. The difference between a tool that delivers value and one that sits unused is almost always about the people around it.

Who's using it, and for what?

Be specific. A medical director summarising published literature has a different profile from a junior team member drafting external content. Mapping this out makes both training and governance much simpler.

What training do they actually need?

You don't need 'prompt engineers'. You need people who are good at briefing clearly, checking work critically and knowing when something doesn't look right. That's not a new skill. It's the same thing you do when you delegate to a colleague or brief an agency; you're just applying it to a digital tool instead.

A few specific areas are worth covering:

- **AI literacy:**
What the tool can and can't do, roughly how it works. Not a computer science degree. Just enough to set realistic expectations.
- **Prompting:**
How to give effective instructions and provide context. Think of it as writing a good brief. Specific, clear, with examples where helpful.
- **Critical evaluation:**
How to spot hallucinations, bias and gaps. The most important skill and the one most often skipped.
- **Compliance awareness:**
What data are appropriate to input, and how to handle outputs going external.
- **Tool-specific training:**
The features and quirks of your chosen tool. Every tool has its strengths and blind spots.

OUR TAKE: Training isn't a one-off event. The teams getting the most from AI treat it as an ongoing capability, not a workshop they attended 6 months ago. As tools evolve and usage matures, keep coming back to it.



Roles and responsibilities

Clarity on roles prevents both duplication and gaps. For any AI-assisted workflow, someone needs to be covering these functions:

Role	Responsibility
Tool champion	Owns the relationship with the tool. Stays current on updates and best practices. First port of call for questions. Feeds back on what's working.
Day-to-day users	Using the tool in their workflow. Need training, clear scope, and a route to escalate. Should feel confident saying "I'm not sure about this" without being shut down.
Output reviewer	Named individual who reviews AI outputs before they progress. Has subject-matter expertise to catch errors. Should understand AI-specific risks (hallucinated references, confident inaccuracies).
Compliance link	Connection to IT, legal or compliance for escalation. Ensures usage stays within approved boundaries. Doesn't need daily involvement, but needs to know what's happening.

Not every role needs to be a different person. In a small team, one person might cover several. The point is that someone is consciously filling each function.

You own the output

This is the thread running through everything in this guide. Regardless of the tool, you are responsible for everything that goes out with your name, your team's name or your company's name on it.

AI doesn't change accountability. It changes speed.

A document generated by AI and reviewed by a human is a human-produced document. That's a good thing. It means AI fits into your existing quality standards rather than requiring a whole new framework. Your expertise is what makes the output valuable. The tool just helps you get there faster.

In practice

- Review every AI output before it leaves your desk
- Verify all references, data points and claims independently
- Apply the same approval processes you'd use for any content
- Keep a record of tools used, prompts and review steps
- If you'd put your name on it, send it with confidence

Working with agencies

When you commission work from an external agency, the same applies. A responsible agency stands behind everything it delivers, regardless of the tools used. AI is a tool in the kit alongside research, writing, design and strategy.

If your agency uses AI, expect transparency, robust quality control and full ownership of the deliverable. The tool doesn't diminish accountability. If anything, it raises the bar for quality control.

At Camino, we stand behind every piece of work we produce. The tools evolve. The responsibility doesn't.

Quick reference checklist

Run through this before adopting any new AI tool or applying an existing one to a new task.



Data and security

- ☐ I've classified the data sensitivity level for this task
- ☐ The tool is on my company's approved list (or I'm building a case for approval)
- ☐ The tool is appropriate for this data sensitivity level
- ☐ I know where my data will be processed and stored
- ☐ Data terms are acceptable (no training on my data)
- ☐ Data residency requirements are met
- ☐ I know what foundation model powers this tool

People and training

- ☐ I know who will use this tool and what for
- ☐ Users have had proper training (not just "had a play with it")
- ☐ There's a named tool champion or go-to person
- ☐ Users know what data they can and can't input
- ☐ There's a clear route to escalate questions or concerns

Use case and fit

- ☐ I can describe the specific use case in one sentence
- ☐ I've considered the consequences if the output is wrong
- ☐ The tool's capabilities genuinely match the task
- ☐ I've tested it on my own real-world tasks (not just the vendor demo)
- ☐ I've checked whether an existing tool could already do this

Workflow and review

- ☐ AI outputs enter our standard content approval workflow
- ☐ A named person reviews every output before it progresses
- ☐ All references and factual claims are verified independently
- ☐ We're keeping a record of tools, prompts and review steps
- ☐ I can explain how this content was produced if asked
- ☐ I'm confident putting my name to this work

TIP: Save this checklist as a reusable template. Revisit whenever you adopt a new tool or apply one to a new use case.





New adventures in medcomms

james@caminocomms.com
caminocomms.com

Produced for Adventures in Pharma • 2026