PDF tldraw

AI Bootcamp for Pharmacosmos

Today's goal: one task, one clear prompt, one reusable GPT

Pharmacosmos - Køge 30 April 2026

1 work task
Choose a real task from daily work where AI can help safely.

1 clear prompt
Tell AI what the task is, what boundaries apply, and what the answer should look like.

1 reusable GPT
Save the strong prompt as a GPT that can be used again.

> Prompt = your instruction to AI. GPT = ChatGPT with fixed instructions for a specific task.

PDF tldraw

Today's material is a website, so you can use it actively

In the room

- → You can move forward and back at your own pace
- → We will use selected slides, exercises and live demos
- → The examples are templates; the real cases come from your own processes

After today

- → You can return to prompts, sources and exercises
- → The material works as a reference site, not only as a presentation
- → The website format makes it easier for me and AI to improve the material together

Today's work product

Today each participant builds a small, safe AI work pattern:

Task from daily work	Subtask AI may help with	Prompt first test	Expert review human owns it	GPT instruction can be reused	Baseline next test

The goal is not to master all of AI. The goal is to use Enterprise ChatGPT for one concrete task in a better and more controlled way.

> Baseline = how the task works today, before the AI test.

PDF tldraw

Geoffrey Hinton: AI is powerful, but not predictable

Geoffrey Hinton is one of the key figures behind the technology many modern AI tools build on. His point matters today: AI can be very useful, but it must be guided and controlled.

Expectation of me

I manage the frame, pace and safety in the room.

Expectation of the model

AI delivers suggestions, patterns and questions - not final truth.

Expectation of you

Your expertise decides what is correct, relevant and safe.

Open the Geoffrey Hinton clip

Use it as perspective on why AI requires curiosity, rules and human control.

Nobel Prize facts · University of Toronto

PDF tldraw

AI does not replace your expertise - it moves where it is used

> Less manual production. More clear task framing, expert review and decision-making.

Before

I write, collect, sort and double-check most of it myself.

Now

I give AI a bounded task and review the quality.

Today

Each person chooses one process where AI can help without taking over expert responsibility.

The work pattern changes: humans plan and evaluate

Human defines task	AI drafts or structures	Human reviews	AI improves	Human decides

Planning

The human defines purpose, boundaries, input and output.

Production

AI can produce a draft, table, checklist or questions.

Evaluation

The human checks evidence, quality, risk and fit for use.

The key shift is not less expertise. It is moving expertise into task framing, review and decision-making.

PDF tldraw

AI creates value only when expertise, tool and process fit together

> Today we are not training clever one-off questions. We are training a clear way to instruct, review and improve AI work.

Expertise

Your judgment decides what is correct, relevant and safe.

Technology

AI suggests from patterns. It can help quickly, but it must be guided and checked.

Process

We start with your current workflow, not with the AI tool.

PDF tldraw

First: look at the work from 10,000 feet

Do not ask AI to copy today's workflow. Use AI to inspect the work before doing the work.

Purpose

What is this work meant to achieve, and who uses the output?

Workflow

Which steps create value, risk, delay or repeated review work?

AI support

Where could AI draft, structure, question or check without owning the conclusion?

```
Ask me questions about the purpose, users, steps, expert judgment and safe AI support.
```

PDF tldraw

New processes: from human alone to human plus AI

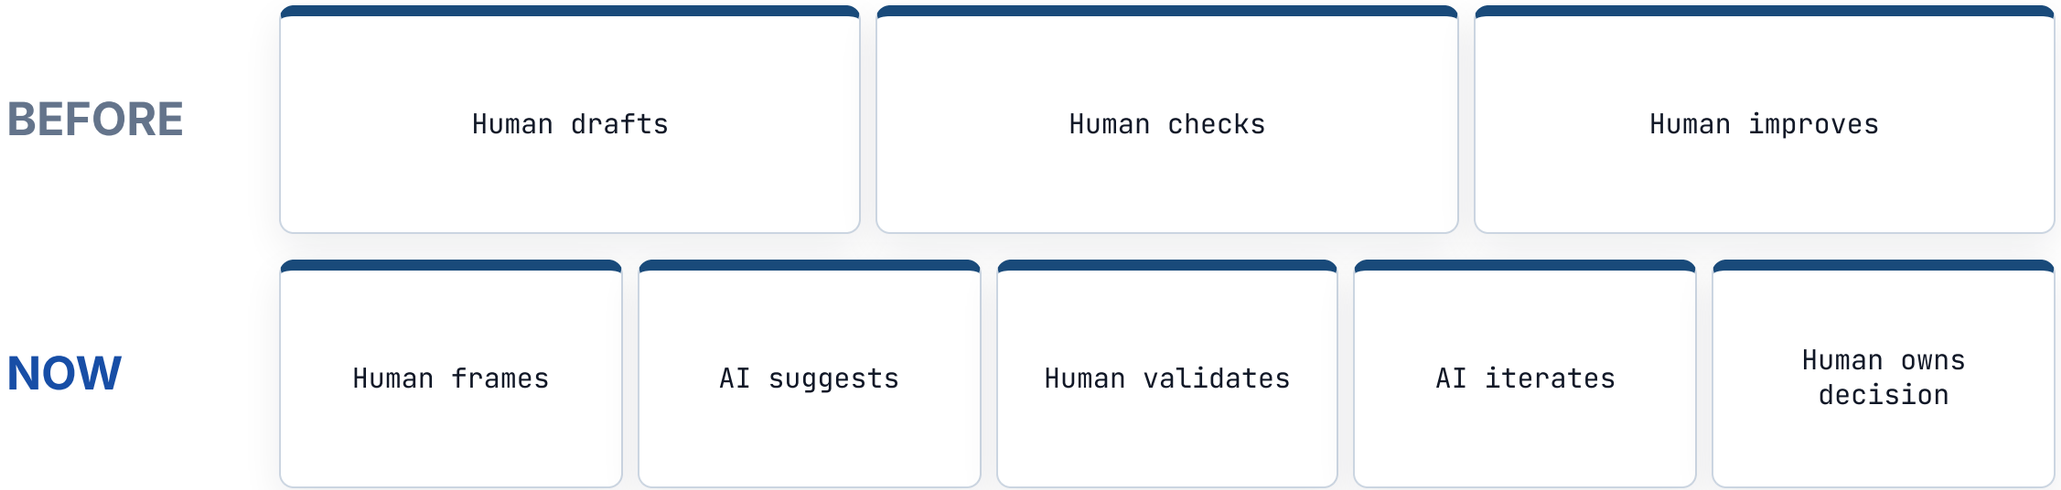

You should not only find a good question for AI. You should find a process step where AI can help without removing responsibility, safety or expert review.

Others in pharma already use AI - but we start with the work pattern

AstraZeneca

AI and data science are used across discovery, clinical trials and development. Source

Roche / Genentech

Roche and NVIDIA describe generative AI for drug discovery and development flows. Source

Pfizer / CytoReason

AI-driven disease models are used for development decisions. Source

→ We are not doing the same thing today. We are learning the pattern: bound the task, set the data boundary, and review AI output.

Novo Nordisk and OpenAI make generative AI part of pharma infrastructure

On 14 April 2026, Novo Nordisk announced a strategic collaboration with OpenAI to integrate advanced AI across research, manufacturing, supply chain and commercial operations.

Discovery and data

AI will help analyze complex data sets and identify promising drug candidates.

Operations and scaling

Pilots cover R&D, manufacturing and commercial operations, with full integration as the goal by the end of 2026.

Control

The collaboration is framed around data protection, clear rules and human oversight.

The point for us: AI is not only a writing tool. It becomes work infrastructure that requires safe processes, strong instructions and expert review.

Berlingske: Novo Nordisk enters partnership with OpenAI · Press release via Nasdaq

AI tools evolve faster than our work routines

OpenAI's release notes are a useful reminder: ChatGPT is not a static tool. In a little over three years it has moved from simple chat to GPTs, images, browsing, files, apps, research and more automated workflows.

Exercise now

Open the ChatGPT release notes and scroll down to the first ChatGPT notes from December 2022.

Look for

How quickly the features move from chat to workspaces, tools and integrations.

The point

We are not learning one button. We are learning a work pattern that can keep up.

Open ChatGPT release notes

PDF tldraw

This website was built with a small AI team

PERSON / AGENT	TASK IN THE WORKFLOW
Torben	Defines goal and task
ProjectLead	Aligns goals, rules and direction
Thomas Director	Checks sponsor value and pharma fit
Medical Science Audience	Checks participant clarity
Slide / Visual QA	Checks layout and readability
Torben	Guides and approves live version

The point is not the technology. The point is the work pattern: clear roles, clear boundaries, multiple checks and human decision-making.

> Today we are not building this kind of system for you. We start with the same pattern at a small scale: one task, one AI role and one review point.

PDF tldraw

We work with process descriptions, not confidential documents

OK in the room

Role, process step, generic input type and desired output. No names, numbers, dates, products or documents.

Approved enterprise only

Confidential work content according to your internal rules. Use Enterprise ChatGPT, and keep human control.

Stop

Trade secrets, patient data, GDPR personal data, product formulas, contract terms, supplier names and non-public clinical details.

When we practice today: describe the process without sharing the actual material.

What does "sanitized" mean in practice?

A sanitized process description may describe **what type of work it is**, but not reveal the actual content.

Yes
"We create an internal summary of a meeting."

Yes
"Input is notes, output is decisions and open questions."

Yes
"Expert review is done by the subject owner before sharing."

No
Names, patients, products, numbers, dates and concrete document excerpts.

No
Customer, supplier or deal details.

No
Data or code that is itself IP or confidential.

Mini exercise: where may the input be used?

6:00

Place each example in one of three categories:

- → Public AI
- → Enterprise ChatGPT
- → Do not use in AI

Examples: generic process description, protocol outline without product name, patient data, public research on a disease area, supplier name in contract context, product formula or recipe.

Today's exercises use only harmless process descriptions. When you work for real, follow your internal AI rules.

Generative AI can help, but it does not automatically know what is true

Generative AI can suggest:

- → text and structure
- → questions and review points
- → summaries and next steps
- → checklists and draft workflows

That is why we always work with:

task → material and rules → boundaries → desired answer → human review

Basic language: six AI terms we use today

Model

The system that calculates a likely answer from patterns in training and context.

Prompt

Your instruction: task, role, context, boundaries and desired output.

Context

The information AI receives for the task: text, files, history, examples or rules.

Context window

How much information the model can use at once. It is large, but not infinite.

Hallucination

When AI answers fluently, but incorrectly or without a reliable source. That is why we review the answer.

GPT

ChatGPT with its own instructions, knowledge and starter prompts for a recurring task.

Three difficult AI terms that save review time

Context window

AI only works with what is inside the window. Too much noise or old side tracks can make the answer worse.

Conversation drift

When the chat slowly moves away from the original task because side tracks become new context. Technical term: prompt drifting.

Hallucination

When the answer sounds confident, but is based on a wrong guess or missing source. Use it as a draft, not as truth.

Stop rule: Start a new chat or restate the task, boundaries and source requirements when the work begins to drift.

The 3 major language models

PDF tldraw

From idea to improved solution - together with AI

Idea or problem

Discuss with AI

AI drafts

Review and adjust

Test together

Improve and repeat

Start early

Use AI before the draft is finished: ideas, alternatives and review questions.

Stay in control

Define the task, inspect output and keep the expert decision with the human.

Improve the pattern

When the answer is wrong, improve the instruction instead of only fixing the answer.

ChatGPT: first chat, then GPTs

We start with ordinary chat because it is the fastest way to learn the AI work pattern. But if a task repeats, it should not be reinvented every time.

1. Chat
Test the task quickly in ordinary ChatGPT.

2. Prompt
Make the task, boundaries, format and review point clear.

3. GPT
Save the strong prompt as a GPT for recurring work.

Today's flow: **write a clear prompt → test in chat → save the pattern as a GPT → choose one task after the workshop**.

A GPT is ChatGPT with fixed instructions

A GPT is a version of ChatGPT where role, rules and output format are saved in advance. It is not a full agent that can act on its own.

But it can give you:

- → the same role every time
- → the same rules every time
- → the same output format every time
- → the same safety boundary every time
- → the same review questions every time

→ Today we build a simple Process Card GPT together.

How to create a GPT

1
Open GPTs in ChatGPT and choose Create.

2
Write what the GPT should be used for.

3
Insert role, rules, stop rules and desired output.

4
Add 2-3 starter questions that fit the task.

5
Test on a harmless process description.

6
Adjust the instruction and save only when the review point is clear.

OpenAI Help: GPTs · OpenAI Help: Data Controls

Template: from prompt to GPT

Copy the parts that fit your process.

```
This GPT helps with:
[one recurring work task]

Role:
You are [critical sparring partner / review assistant / structure assistant].

Rules:
- Ask first if the task is unclear.
- Never ask for trade secrets, patient data or non-approved confidential information.
- Do not give final expert conclusions.
- Always mark what the human must review.

Output format:
[table / checklist / outline / questions]
```

PDF tldraw

Working with AI requires a choice of control style

S1-S4 are four collaboration styles: from precise instruction to more delegation.

S1: Directing

You decide steps, boundaries and format. Use when the answer must be precise and easy to review.

Example: Create a table with 3 columns: process step, AI contribution and review point.

S2: Coaching

You give direction and ask AI to think with you, explain choices and suggest improvements.

Example: Suggest 3 ways to improve this process, and explain benefits and risks.

S3: Supporting

AI asks questions and helps you clarify options, risks and next steps.

Example: Ask me 5 questions before you suggest how AI can help safely.

S4: Delegating

You set the goal and let AI suggest method and draft. Use only when risk is low and the review point is clear.

Example: Create a first draft checklist for this sanitized task.

Often start in S2 or S3. Use S1 for precision. Use S4 only when the task is familiar and human review is clear.

A weak prompt creates more review work than it saves

Weak

"Help me with my process."

Strong

```
You are my critical sparring partner.
I describe a workflow
without confidential details.

Find 3 safe AI contributions.
Return a table with:
process step, AI contribution, data risk
and human review.
```

→ The difference is not the model. The difference is the work instruction.

Different ways to work with generative AI

Ask

Use AI for questions, explanations and alternatives.

Draft

Use AI to create an outline, table, checklist or first version.

Review

Use AI to find gaps, assumptions, unclear claims and review needs.

Orchestrate

Use fixed steps, roles and stop rules when the task repeats.

Today's work: concrete workflows → safe AI work patterns.

PDF tldraw

A good AI task starts with process, not technique

Process step
What happens before, during and after the task?

Harmless description
What can be described without confidential details?

AI contribution
Structure, draft, review questions or checklist?

Stop rule
What must AI not see or decide?

Output format
Table, list, questions or next steps?

Review point
Where does AI stop, and where does the human take over?

Give AI the right workspace

AI performs better when it receives the right information - and not too much noise. This is also called context engineering.

The workspace is limited

More input is not always better. Too much noise makes it harder to hit the target.

The prompt is only one part

Instruction, data, examples, tools and history together shape the AI answer.

High-signal and tight

Give the smallest set of information that makes the task, boundary and output clear.

Anthropic: Effective context engineering for AI agents

Question technique: from short answers to deeper reflection

1. Simple answer

Ask for a short answer or list.

2. Clarification

Ask AI to explain assumptions and missing information.

3. Analysis

Ask for options, trade-offs, risks and evidence needs.

4. Deep reflection

Ask AI to challenge the process and identify blind spots.

5. Transformation

Ask AI to help redesign the workflow or turn it into a reusable pattern.

The deeper the level, the more you must give AI relevant material, role, boundary and review point.

Choose your work track

Choose the track that looks most like your daily work. Use the same track through the rest of the exercises.

Medical Writing

Protocol outline, document structure, review questions and checklists.

Science / Medical

Study idea, evidence, assumptions, uncertainties and decision questions.

Business Development

Public research on a disease area, treatment landscape and information needs.

Statistics / programming

Analysis process, validation, error sources and explanation of code flow.

Leadership / adoption

Baseline, micro-habit, next test and follow-up.

Other

Choose a recurring task with fixed inputs, fixed steps and a clear review point.

PDF tldraw

Exercise 1: Make a process AI-ready without sharing confidential data

10:00

Open Enterprise ChatGPT. Choose a current workflow and describe it in 5 lines without names, numbers, patients, products or confidential documents.

Ask ChatGPT for sparring on where AI could help safely.

Output:

- → 3 process steps
- → 1 safe AI question
- → 1 thing AI must not see
- → 1 human review point

Process description only: no names, numbers, products or confidential documents.

AI is strongest as a critical assistant before it is used as a writer

Ask before you request text:

- → "Find gaps in this process description."
- → "Which clarifying questions should an expert ask?"
- → "Which parts require sources, data or expert approval?"
- → "What can AI help with without seeing confidential material?"

> The best first output is often a review checklist, not a polished draft.

Prompt exercise: get AI to ask the right questions

Copy the prompt and insert your sanitized process description at the bottom.

```
You are my critical sparring partner.

I describe a work task without confidential details.
Before you suggest a solution, ask 5 clarifying questions.

Ask only about process, purpose, desired answer and expert review.
Do not ask for names, patients, products, numbers, document excerpts or trade secrets.

My process description:
[insert 5 lines here]
```

> A good prompt often starts with better questions, not with a polished answer.

Process lab: document work and Medical Writing

If your process is about documents, AI can help without seeing the document itself.

Structure
"Create an outline from this sanitized task description."

Expert review
"Find unclear assumptions and questions for expert review."

Checklist
"Create a checklist before the document moves forward."

> Describe the document type and task, not the document.

Process lab: science, BD, data or tables

If your process is about analysis, research or tables, describe the purpose before the content.

Science
"Find claims, evidence needs, assumptions and questions before the next decision."

BD research
"Create key questions for public research on a disease area."

Validation
"Suggest which checks a human should make before output is used."

No raw data, code IP, deal details or non-public clinical details in the exercise.

From process description to reusable GPT instruction

A process card is a simple table: process step, AI contribution, data risk and human review point.

Later, the same instruction can become a GPT for recurring work.

A GPT improves when errors are written back into the instruction

We first test with a harmless process description.

If the GPT answers too broadly, we improve it:

```
Always answer as a table with four columns:
Process step | AI contribution | Data risk | Review point
```

Then we test again.

> This is the meta-improvement: errors become better instructions.

Agent thinking starts as a fixed workflow

An agent is not magic. Think of it as AI in a fixed workflow with clear steps and stop rules.

It is a work process with:

- → fixed steps
- → clear inputs
- → stop rules
- → review point
- → documentable output

Today we work with the part you can use now: better GPTs and better work instructions.

AI in the browser advises - agent tools can also act

AI in the browser

Advises, drafts and discusses inside the chat. It does not access your local files or run commands.

Agent tools

Can read, write and run actions in a controlled environment if they are given access.

Today

We use ChatGPT and GPTs as the practical starting point.

Future direction

Agent tools require stronger governance, permissions and review routines.

For coding: use a harness, not a loose prompt

For resourceful coding users, the main lesson is technical: do not ask AI to "just build it". Give it a controlled engineering harness.

State

Use git history, progress notes and a feature list so the next AI session knows what happened.

Scope

Ask for one feature or fix at a time. Avoid one-shotting a whole application.

Tests

Require unit checks, end-to-end checks or browser automation before marking work as done.

Clean handoff

End with committed code, a short progress update and clear next steps.

Recent engineering guidance: long-running agents need a setup that preserves state across context windows and prevents premature "done".

Anthropic Engineering: Effective harnesses for long-running agents

Today we do not install agents - we learn the pattern

What you should recognize:

Plan what should happen?	Input what may be used?	Execute what can AI help with?	Review what must the human check?	Learn what should improve?	Next test what do we try?

This is not an installation requirement. It is a way to think about work.

Exercise 2: Build a fixed instruction for one recurring task

16:00

Choose one process:

- → expert review of an argument
- → summary of meeting decisions
- → structuring non-confidential notes
- → preparing questions for expert review
- → translating a process into a checklist

Define role, input, stop rules, output format, review point and first test.

The instruction must be shareable without confidential content. Confidential work belongs only in an approved environment and with human review.

Adoption happens through micro-habits, not big transformations

Choose small fixed practices:

- → "AI-assisted?" as a point in the meeting template
- → 10-minute case demo in a team meeting
- → shared GPT or prompt library
- → baseline and after-measurement on one task
- → review rule: AI output must never move forward without expert approval

The leader's question: **How did you use AI, and where was the review point?**

Build-Measure-Learn: how AI work improves

AI improvement is a small learning loop, not a large project.

Build

Try one prompt, GPT instruction or workflow on a harmless case.

Measure

Measure whether output became shorter, clearer, faster or required less review.

Learn

Write the error back into the prompt, GPT instruction or process card.

Example: If AI answers too long, add a length requirement. If AI guesses, add source requirements and a stop rule.

We end with three process cards from the room

For each process card:

FIELD	DECISION
Process	Which workflow will we test?
AI contribution	Where can AI help?
Stop rule	What must AI not see? Confidential material is out.
Review	What must the human validate?
Baseline	How long, how many iterations or which error types today?
Next test	What do we try first?

PDF tldraw

Output from today

Process card
One workflow described safely.

Safe prompt/GPT
One prompt or GPT instruction without confidential content.

Review point
A clear stop where expertise takes over.

Baseline
Time, iterations or error type before the AI test.

Next test
One task to try after the workshop.

What do I test after the workshop?

Write one sentence:

> "I test AI on [process] by using [GPT/prompt] for [bounded contribution]. My baseline is [time/iterations/error type]. I evaluate after [first follow-up]."

This is not a major transformation.

It is one safe test that the Pharmacosmos team can follow up on.

The next module can build on today's process cards

Possible continuations:

- → mature one process into a shared GPT
- → create a team checklist for safe AI use
- → test process cards against actual workflows in an approved environment
- → build a shared prompt library
- → choose a supervised pilot with clear gates

PDF tldraw

Extra tool: Perplexity for source-aware web research

Perplexity is useful when the task starts with a question and needs current public sources, not only a fluent answer.

What it does

Searches the web, synthesizes answers and shows source links so you can inspect the evidence.

When to use it

Good for public landscape scans, competitor context, terminology, recent news and finding starting sources.

How to prompt it

Ask for source types, date limits, uncertainty and a short list of claims that require human validation.

Stop rule

Do not paste confidential material unless your organization has explicitly approved that environment and use case.

Use the citations as a starting map. The expert still decides whether the source is relevant, current and suitable for the work.

Open Perplexity · How Perplexity works · Pro Search

Extra tool: NotebookLM for sources

NotebookLM is useful when the task starts with material that must be understood, connected and made usable.

Sources first

Upload or add materials such as PDFs, web pages, YouTube, audio, Google Docs or Slides.

Answers with citations

Chat with the notebook and get answers grounded in the sources with citations.

Multiple formats

Create a briefing, study guide, mind map or audio overview from the same source set.

Safe use: add only material that your internal rules allow in the specific Google/Workspace environment.

Open NotebookLM

Extra demo: AI and demand management in pharma

We use a concrete notebook as an example of how AI can make a source set more usable.

1
Open the notebook and see which sources it uses.

2
Ask for key concepts, assumptions and possible decision points.

3
Use the answer as overview - not as truth without expert review.

The purpose is not to automate demand management. The purpose is to show how sources can become questions, overview and next steps.

Open notebook: AI and demand management in the pharmaceutical industry

Source: Co-Intelligence is about working with AI

Ethan Mollick's book *Co-Intelligence: Living and Working with AI* describes generative AI as something you learn by collaborating with it in practice - not only by reading about it.

Invite AI early

Use AI for ideas, alternatives and review questions before the draft is finished.

Be the human in the loop

Bound the task, review output and keep expert decisions with humans.

Give a clear role

Tell AI what kind of assistant it should be, and which boundaries it works within.

Co-Intelligence on Amazon · Wharton interview about the book

Resource: AI access has become a leadership resource

Jensen Huang from NVIDIA points out that access to compute and AI consumption becomes a real constraint for knowledge work.

Open the video in a new tab

Jensen Huang on compute, AI consumption and AI access as a leadership resource.

Use it as perspective, not as a technical requirement for today's exercises.

Extra tracks if time allows

Use these slides only if the topic comes up in the room.

AI's expert boundary

Why review points and human expertise are necessary.

Pharma case

Bayer as an example of AI for concrete document processes.

Adoption

Why skills, governance and micro-habits matter more than inspiration.

Agent/GPT measurement

Why the success criterion is value, quality and review load - not the number of agents.

PDF tldraw

Extra: AI has an expert boundary

Jagged frontier means: AI can be very strong inside some tasks and significantly weaker just outside them.

Inside the frontier

Studies show faster solutions and higher quality when the task fits the model.

Outside the frontier

The result can become worse, even when the answer sounds convincing.

Consequence

Use AI as suggestion and sparring. Keep expert review as a fixed stop.

Connection to today: a good prompt also tells AI where to stop.

PDF tldraw

Extra: Bayer shows the point about concrete process

Bayer has described using generative AI for drafts in clinical study reports.

Not "AI for everything"

Choose a specific document process with known format and known review requirements.

Medical Writing

AI can help with structure, first draft, consistency and review questions.

Human validation

The subject owner checks claims, data, sources and regulatory suitability.

Question for the room: which document type in your work has the most repetition and the clearest review point?

Extra: adoption requires skills and governance

Deloitte's AI market signal points in the same direction: the large barrier is not only tools, but skills and governance.

Skills gap

Teams must learn to frame tasks, assess output and know data boundaries.

Training

Training is often the primary strategy for bringing AI safely into work.

Governance

Autonomous agents require clearer rules than ordinary chat and GPTs.

Connection to today: start with micro-habits, process cards and one safe test.

Extra: measure the value of GPTs and agents

The McKinsey/EY/PwC point: the number of agents is not the success criterion.

Value

Which task becomes faster, better or easier to follow up on?

Quality

Does output become more consistent, complete or easier to review?

Review load

Does AI save time, or does it simply move the work into checking and corrections?

Good test: one GPT, one process, one baseline and one follow-up.