DATA ON DECK

How to Use AI in Your Project

A practical guide for project teams.

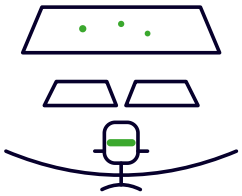

Dr Gregor Lichtfuss

Version 1.7 · June 2026

lichtfuss.earth/data-on-deck

Suggestions and marked-up copies welcome: openchannel@lichtfuss.earth

INTRODUCTION

Picture yourself in the captain's chair. You have a crew of officers who never tire, know a little, or a lot, about everything, and answer the instant you ask. Not one Data — several. They can do all the work for you. Splendid.

But they share one flaw. They'll answer with the same calm confidence whether they're right or wrong; they have no memory of what being wrong costs; they don't care about the consequence. “I calculate that you have a 67% chance of survival.” (Note: there is no “we” on this ship!)

And sooner or later someone — an examination marker, a supervisor, a conference room — will ask how you got to the conclusions of your work. “The AI said so” is not a strong answer. In fact, it's a lot worse than “I have absolutely no clue.” A captain who can't evaluate a response to an order should never have given the order.

The job was never to have the crew that can do all the work for you — everyone has a crew like this now. The job is to command it: to know what to ask, what to check, what to own, and what to do with the output. And to keep the log, so you can show your work when you're asked — what you requested, what you checked, what you decided.

This is what this guide is for. Starfleet officer training.

1 Where we stand on AI

AI is a tool. Use it.

That sentence is the policy. Everything else in this guide is about how — because using AI without method is not the same as not using it. It is worse.

When you submit work that AI generated without your direction, your verification, and your documented decision-making, you are not demonstrating that you can use AI. You are demonstrating that you cannot. That is the distinction that matters here.

This guide is not a restriction on AI use. It is a framework for making your AI use visible, defensible, and genuinely useful to the project — and to you.

> **The three rules**
>
> **Think before you prompt.**
>
> AI amplifies your thinking. If your thinking is vague, the output will be too.
>
> **Document what you did.**
>
> If it is not recorded, it did not happen — and you cannot defend it.
>
> **You decide.**
>
> No decision in this project is delegated to AI. You are responsible for the work and its conclusions.

How this guide is structured: Plan — Do — Review

Good work follows a cycle. You plan before you act. You act with intention. You review what happened and use that to improve next time. This guide uses that structure throughout.

The sections are organised into three layers:

PLAN	Understand the development landscape. Locate your project. Define your milestone. Design your team before you build it.	2–5
DO	Build and run your virtual AI team. Execute tasks with method. Record as you go.	6–8
REVIEW	Report transparently on what you did and how. Reflect on what worked. Feed that back into the next cycle.	9–10

You will move through this cycle more than once on any meaningful project. Each pass makes the next one faster and better.

2 The digital team model

Throughout this guide, AI tools are treated as a virtual team of specialists — working under your direction. This is a practical organising principle, not a metaphor.

In industry, complex projects are delivered by teams: people with different expertise, different roles, and different responsibilities, coordinated towards a shared goal. The same logic applies here. You are the project lead. You define the goal, assign tasks, review output, make the decisions, and remain accountable for everything that gets delivered.

Your AI team — Claude, NotebookLM, Gemini, and others — each have different strengths and appropriate uses. Like any team, they do what you tell them. No more, no less.

> **What makes this model work**
>
> **You define the goal before you assemble the team —**
>
> not the other way around.
>
> **You assign roles based on what the stage requires —**
>
> not based on what tool is most familiar.
>
> **You know when you need a human —**
>
> some roles cannot be substituted. Plan the handover in advance.
>
> **You manage the team —**
>
> the output is always reviewed, validated, and decided upon by you.

The sections that follow explain how to build and run that team — starting with a question that applies to any project: when the project lead is not one person but a group, how does the model change?

3 Individual project or team project — what changes?

The digital team model in Section 2 applies in both cases. What changes is the layer of human coordination required before the virtual team can function effectively.

Individual project (including a defined sub-project within a larger team)

If you are working alone — or if you have been assigned a clearly scoped sub-project with sole ownership — the model is direct. You are the project lead. You do the planning work, build the virtual team, run the tasks, keep the records, and own the deliverable.

This is the simplest case, and the rest of this guide addresses it directly. The planning layer in Section 5, the task cycle in Section 6, and the record-keeping in Section 7 all apply without modification.

Team project — where sub-projects and responsibilities are not yet defined

When a group of people shares responsibility for a project and roles have not yet been assigned, a coordination problem exists before any AI problem does. The digital team model assumes one project lead; a group without defined responsibilities has no clear lead. Without deliberate agreement up front, the predictable failures follow:

- Two people independently prompt the same question and get different answers — neither knowing the other did it
- One person drives all AI-assisted work while others disengage — producing an uneven Methods section and uneven learning
- The group assembles individual AI outputs at submission time and discovers they do not fit together

None of these are AI problems. They are coordination problems that AI makes more visible. The solution is to resolve the coordination questions before the virtual teams are built.

Structuring a team project: one virtual sub-team per person

The most effective approach: divide the project into sub-projects based on the work packages identified during stage-gate planning (Section 5), and assign each team member as the lead of one sub-project and its associated virtual AI team.

Each sub-project lead is then responsible for:

- The planning work for their scope — the question, the approach, the virtual team design
- Running and managing their virtual team — prompting, reviewing, validating, deciding
- Record-keeping for their scope — the AI log that feeds into the shared Methods section
- Reporting back to the human team at each milestone — what was produced, what was decided, what is needed from others

The human team operates as a coordination layer above the virtual sub-teams: aligning on goals at milestones, reviewing each other's outputs, making collective decisions, and integrating the work into shared deliverables.

What the team must agree before anyone starts

Before any virtual team is built or any AI tool is used on project work, the human team needs to reach agreement on five things:

1	Sub-project ownership	Who leads which scope of work? Overlapping or unassigned scope leads to duplicated effort, conflicting outputs, or gaps nobody notices until submission.
2	Which tools the team will use	A shared tool set makes it easier to share prompts, compare outputs, and write a coherent Methods section. It does not need to be identical across sub-projects, but the choice should be deliberate and documented.
3	How AI use will be recorded	Each sub-project lead keeps their own log using the minimum fields in Section 7. Agree the format in advance so the logs combine into a shared Methods section without gaps or inconsistencies.
4	Milestone check-in points	When does each sub-project lead report back to the group? What do they need to bring? Agreeing this in advance prevents surprises at integration time.
5	How handovers work	When one sub-project's output becomes another's input, what form does it need to arrive in? Defining this early prevents the human coordination layer from becoming a bottleneck.

The key point

In a team project, AI use is a team decision — not an individual one. The goal, the tool choices, the process, and the record-keeping all need to be agreed before work begins. A team that resolves these questions first will produce better work and a more coherent Methods section than one that assembles independent AI outputs at the end and hopes they fit together.

4 Four zones of AI use

Not all AI use is equal. Here is a plain-language map of where your work should sit.

Not using AI at all.	**Prompting AI and submitting the output.**	**Using AI to expand, check or improve your own thinking.**	**Designing a deliberate AI-assisted workflow. Thinking first. Using AI as a tool. Recording everything.**
Acceptable, but you are leaving value on the table.	**Not acceptable. This is not your work.**	**Acceptable. Document what you did and how.**	**This is what we aim for.**

> **Our aim**
>
> Zone 4 — Intentional — is what we are working towards. Zone 3 is acceptable for most tasks. Zone 2 is not acceptable under any circumstances. Zone 1 is your choice, but you are limiting yourself.

PLAN Understand the landscape. Locate yourself. Define the milestone. Design the team.

5 Before you build your team — map the landscape

Section 2 introduced the digital team model and Section 3 addressed how it adapts when the project lead is itself a team. Before you can build your virtual team, you need to know what stage you are actually in — and what that stage demands from the team.

Every industry has a recognised development process from early idea to finished product in the market. The stages differ in label but not in logic: each one has defined inputs, expected outputs, typical activities, and a stage gate — a decision point that determines whether development continues and on what basis.

This planning work is what makes everything else coherent. Most students skip it, jump straight to a tool, and produce output that does not fit the question. The planning layer is not extra work. It is what makes the subsequent work useful.

Before you build an AI team, you need to answer five questions. AI can help you work through all of them — but you need to lead the thinking.

1	What does the full development pathway look like in this sector?	From early idea to marketed product — what are the recognised stages and stage gates? If you do not know, this is your first research task. AI can sketch the pathway; you validate it against industry sources.
2	Where on that pathway is this project right now?	Be honest. Is this discovery-stage work? A feasibility study? A pre-commercial proof of concept? Mistaking your stage leads to the wrong team, the wrong outputs, and the wrong questions.
3	What does this stage gate require?	What data, evidence, or decisions are needed to pass this stage? What would a funder, regulator, or senior reviewer expect to see? What would stop the project moving forward?
4	What does a typical development team look like at this stage?	In industry, who is in the room? A scientist, regulatory specialist, commercial analyst, project manager? This is your team blueprint — the roles you need to staff, whether with humans or AI.
5	Where will you need a human, and what do they need to receive?	Some roles cannot be substituted. Legal review, ethics assessment, expert validation — these require humans. Map the handover points early: what does the AI team produce, and in what form, for the human to pick it up effectively?

Example: a generalised development pathway

The stages below are drawn from life sciences but recognisable across biotech, foodtech, medtech, and adjacent sectors. Your project may use different labels. What matters is the logic: each stage narrows uncertainty, generates evidence, and produces a decision.

Discovery & Ideation	Concept definition, literature review, hypothesis, initial feasibility assessment	Scientists, domain experts, strategists	Is the concept technically feasible and worth pursuing?
Proof of Concept	Experimental data, prototype, technical validation, risk register	R&D, technical specialists, project manager	Does the core science or technology work at small scale?
Feasibility & Development	Scaled data, process definition, regulatory pathway, commercial model	R&D, regulatory, commercial, manufacturing	Can this be developed into a viable product at scale?
Validation & Approval	Compliance documentation, validated process, regulatory submission	Regulatory affairs, QA, legal, clinical or field teams	Does this meet the requirements of the relevant regulatory or approval body?
Launch & Commercial	Market entry plan, commercial agreements, post-launch monitoring	Commercial, marketing, operations, post-market surveillance	Is the product reaching the market and performing as expected?

Use this table as a template only. Research the actual stages and stage gates that apply to your sector and product type. That research is itself a meaningful early task — and AI can help you structure it.

Translating this into your project plan

Once you have answered the five questions, you have the inputs for everything that follows:

- Your stage gate defines the milestone you are working towards
- The expected outputs define what your deliverables need to contain
- The typical team roles become the blueprint for your virtual AI team
- The human handover points become explicit items in your project plan

DO Build your team. Execute tasks with method. Record as you go.

6 Working with your virtual team

You have mapped the development landscape, located your stage, defined your milestone, and identified the roles the work requires. Now you build the team and put it to work.

The practical method is straightforward: assign a role, provide context, set a specific task, collect the output, review it critically, and decide what to do with it. Then repeat for the next task or the next milestone.

> **What this means in practice**
>
> **They follow your instructions.**
>
> If your instructions are imprecise, the output will reflect that.
>
> **They do not know what is right.**
>
> They produce plausible-sounding responses. You determine correctness.
>
> **They will tell you what you asked for.**
>
> If you ask a leading question, you will get a confirming answer. Ask for critique explicitly.
>
> **They have no memory of consequences.**
>
> They cannot tell you if a decision will work. You carry that.

The six-step task cycle

Use this cycle for every substantive AI task. The thinking in steps 1 and 2 is the most valuable part — do not skip it.

1	Think first	Before opening any AI tool: what is the question? What do you already know? What is your working hypothesis or approach? Write this down, even in point form.
2	Define the task	What specifically do you want AI to help with? Searching literature? Structuring an argument? Checking logic? Drafting a section? Be precise.
3	Choose the right tool	Different tools are better for different tasks. A tool that summarises papers is not the same as one that generates prose. Know what you are using and why.
4	Prompt with intent	Give the AI context, a role, and a specific task. A vague prompt produces a vague answer. Record the prompt you used.
5	Review critically	AI output is a starting point. Check every factual claim. Verify references. Ask: does this actually answer the question I had? Does it contradict what I already know?
6	Decide and document	What did you use from the output? What did you discard and why? What decision did this inform? Record it. You made the decision — the AI did not.

On verification

Every factual claim AI makes needs a source you can cite independently. The more consequential the claim — the more it shapes your conclusions or recommendations — the more rigorously you must verify it. Three useful questions:

Can I find an independent source that confirms this?

Would my conclusion change if this turned out to be wrong?

Am I using this because it is correct, or because it is convenient?

Assigning roles

You can give an AI a specific role to work from. This is one of the most effective prompting techniques, and it is underused.

Example

Instead of: "Review my research proposal"

Try: "You are a senior industry consultant reviewing a student research proposal for a commercial aquaculture company. Your role is to identify gaps in the methodology, challenge assumptions, and flag any claims that are not supported by the evidence provided. Be specific and critical."

The second prompt produces a materially different — and more useful — response. The role gives the AI a perspective. The instruction to be specific and critical overrides its default tendency to be agreeable.

7 Record keeping

You do not need a special system. You need to record enough that someone else — or you, three weeks later — can understand exactly what you did and why.

The format is your choice: a running document, a section in your lab notebook, a structured log file, entries in a project notes tool. What matters is that the information is there.

These are the fields you need to capture for each substantive AI interaction:

Date	**When did you use AI for this task?**
Task	**What were you trying to do?**
Tool used	**Which AI tool or model?**
Prompt / input	**What did you ask? (Paste or summarise the prompt.)**
Output summary	**What did the AI produce? (Key points, not the full text.)**
What you used	**Which parts of the output did you use in your work?**
What you discarded	**What did you reject and why?**
How you verified	**How did you check correctness? (Cross-reference, expert check, manual review?)**
Decision made	**What decision or conclusion did this inform?**

Why this matters

In any serious research or industry context, methods are explained — the sources used, the analyses conducted, the tools applied. AI is a method. It needs to be described as one. Your record keeping is the raw material for that description.

8 Tools — a starting point

This list covers the tools most relevant to this project context. It is not exhaustive, and it will change — the capability gap between these tools has narrowed significantly and continues to shift. The skill is not knowing every tool; it is knowing what a given tool is actually good for, where it falls short, and being able to explain your choice.

For most research and analysis tasks, Claude, Gemini, and ChatGPT cover overlapping ground. The honest reason to prefer one over another is usually familiarity, workflow integration, or a specific task where one has a clear edge — not a categorical difference. Use what works, document what you used, and be prepared to justify it.

Claude (Anthropic)	Deep research support, structured planning, long-document analysis, concept development, critical review, drafting. Strongest for complex multi-step reasoning and precise instruction-following.	Leads on writing quality, nuanced reasoning, and large documents (200K context window). Check all factual claims independently — it can hallucinate confidently. Best for tasks that require sustained analytical rigour.
Gemini (Google)	Real-time web research, broad topic exploration, multimodal tasks (images, video). Native integration with Google Docs, Sheets, Slides, Gmail. Prompt generation for other tools.	Best choice if you live in the Google ecosystem. Strong for current information and research with live web access. Massive context window (up to 1M tokens). Less reliable than Claude for nuanced long-form reasoning.
ChatGPT (OpenAI)	Versatile general-purpose tasks: writing, summarising, coding, data analysis, image generation (DALL-E). Broadest ecosystem of third-party integrations.	The most widely used AI tool — large user community, mature plugin ecosystem, and strong structured reasoning. Covers similar ground to Claude and Gemini for most tasks. Verify factual claims as with any tool.
NotebookLM (Google)	Analysing a defined set of your own sources: papers, reports, lecture notes. Literature review, cross-checking conclusions, generating summaries, quizzes, audio overviews, and slide decks grounded in your sources.	Most reliable when restricted to sources you load — answers stay grounded and citable. Free tier: 50 sources per notebook, 50 chats/day. Now includes Deep Research (web), Data Tables, and PPTX export. Notebooks are silos — no cross-notebook search.
Perplexity	Fast web search with cited sources. Good for literature scans, finding recent papers or industry reports, and fact-checking claims where you need to see the source immediately.	Every answer includes direct citations to the web pages it drew from. Reduces the verification burden for factual lookups. Less suited to long-form reasoning or document analysis.
Gamma.app	Turning structured notes or outlines into draft presentations. Rapid prototyping of slide decks.	Output is a starting point — review content, accuracy, tone, and story flow carefully before presenting. Good for drafting; not a substitute for thinking through the narrative yourself.
Grammarly	Spelling, grammar, clarity, and readability of your own written work. Light rewriting and restructuring.	Does not generate content — it works on text you have written. Easy to attribute. Low risk of introducing factual errors. Good for final polish of reports and written deliverables.
Your brain	Everything that matters: the question, the plan, the judgement, the decision.	Non-negotiable. No AI tool substitutes for this. It is also the only tool in this list that will be held accountable for the outcome.

Confidentiality

Do not submit confidential company data, unpublished research, or personally identifiable information to any public AI tool. If in doubt, ask before you input. This applies to all tools above.

REVIEW Report transparently. Reflect on what worked. Feed that back into the next cycle.

9 What to include in your deliverables

Both written reports and presentations must make your AI use transparent. This is standard practice in any serious research or industry context — not a special requirement invented for this project.

Written reports

Include a Methods section — or equivalent — that describes how you conducted your research and analysis. This applies whether or not AI was involved.

The Methods section has two parts:

- The Plan. What was the research question or task? What was your approach before you used any tool? What development stage were you working at, and what did that stage require? This is the thinking you did before you opened any AI tool.
- The execution. Which tools did you use and for which tasks? What did you ask? How did you verify the output? What decisions did you make based on it?

This does not need to be long. A short paragraph or table for each tool is sufficient. What matters is that it is there — and that someone reading it could understand and reproduce your process.

Presentations

Presentations are more compressed, but the same principle applies. Include a Methods or Process slide — even a single slide — that covers:

- What question or task you were addressing and your initial approach
- Which AI tools you used, for what, and how you verified the output
- What decisions the AI-assisted work informed

Where a dedicated slide is not practical:

- Acknowledge AI-assisted work on the slide where it appears — a footnote such as "Analysis supported by Claude / NotebookLM" is sufficient
- Be prepared to explain your process in Q&A — what you asked, what you received, what you decided
- Do not present AI-generated text as if you wrote it. Paraphrase, adapt, and own the content.

> **The test**
>
> If someone asked you to walk through exactly how you arrived at any conclusion in your report or presentation — step by step, tool by tool, decision by decision — could you do it? If yes: you are working well. If no: you have a documentation gap.

The Review step closes the cycle

After each milestone, take stock. What worked in your AI workflow? What produced poor output? Which verification steps caught errors? Which tools were worth the time?

This reflection feeds directly into the next Plan phase. The cycle is not a one-time process. It is the method.

10 Three principles to work by

When working with AI tools — especially on a team project with real stakes — three principles are worth keeping in mind.

DNA

Inspector Barnaby's rule, originally for detective work. Applies equally well to AI-assisted research:

- Don't trust anyone
- Never believe anything
- Always check everything

This is the high-vigilance setting. Appropriate when the stakes are high, the source is unverified, or the claim is consequential. DNA deserves its own chapter — for now, treat it as the default posture for any AI output that will shape a decision.

Trust, but validate

The more practical day-to-day mode. You can work with AI output and trust its general direction — but you validate anything that matters before acting on it. As you build experience with a particular tool for a particular type of task, the validation burden may decrease. But it never disappears entirely.

You are the boss

Not a metaphor. You are responsible and accountable for every claim, conclusion, and recommendation in your deliverables. Your name is on it. Not the tool's.

This is not a burden — it is the point. The goal of this project is for you to develop real capability: to think, to research, to decide, and to communicate. AI accelerates parts of that process. It does not replace any of it.

This document is a starting point, not a finished policy.

Bring your ideas — on AI tools, workflows, prompts, or anything in this guide you think is wrong, incomplete, or could be better.

The best version of this guide will be written by the team that uses it.

Appendix — Materials and Methods

Creation of this document

This appendix documents how this guide was created, using the same structure — Plan and Execution — that the guide itself requires of students. It is included to model transparent AI-assisted practice, and to make the document's own provenance visible and reproducible.

Part 1 — The Plan

Research question and goal

To develop a practical, plain-language guide on intentional AI use for students working on research and development projects — covering framework, workflow, team coordination, record-keeping, and

deliverable requirements. The guide needed to be usable as a standalone student document and adaptable as a chapter for a broader publication on work and career development.

Starting point and context

The author brought an existing draft guide (SHRMP.bio Use of AI, V26-01, April 2026) and a University of Melbourne GenAI use framework (photographed poster, Biotechnology faculty). The draft established the core positions: AI as tool, use is expected, intentionality and record-keeping are non-negotiable, no decision delegation. The UoM framework was used as a reference point for zone positioning, not as a structural template.

Development stage

Discovery and ideation stage. The goal was a Version 1 document suitable for immediate use with student cohorts, with the expectation of iteration based on student and team feedback.

Intended audience

University students working on industry research projects, individually or in teams, at any stage from early ideation to pre-commercial development. The document was also designed to be sector-neutral and adaptable for use beyond the original student context.

Approach before AI tools were used

The author defined the document's core philosophical positions and non-negotiables independently, drawing on experience running student industry projects. Key positions established prior to AI-assisted drafting:

- AI use is expected and encouraged — but only with intentionality, method, record-keeping, and critical review
- No decision-making is delegated to AI
- The Materials and Methods concept applies to desk research and AI-assisted work, not only laboratory work
- The Plan–Do–Review cycle provides the structural spine
- The digital team model (AI tools as virtual specialists under a human project lead) is the organising mental model
- Stage-gate thinking should precede team assembly — you cannot design a team without knowing what stage you are in

Claude Sonnet 4.6 (Anthropic)	Primary drafting tool throughout. Used for: initial content architecture, section drafting, callout box writing, table content, iterative revision across 9+ versions in response to author feedback and corrections.	All output reviewed, corrected, and approved by the author at each iteration. Sections rewritten, reorganised, and extended based on author direction. No section accepted without author review and decision.	Author reviewed each version for accuracy, tone, completeness, and structural coherence. Factual corrections applied directly. Author retained all decision-making on content and structure.
Claude Sonnet 4.6 (web search)	Research on current state of AI tools: ChatGPT vs Claude vs Gemini capability comparison (2026); NotebookLM features, limits, and recent updates (June 2026).	Used to update and critically review the tools table (Section 8). ChatGPT added; NotebookLM description updated to reflect June 2026 feature set; Perplexity added as a distinct tool category.	Cross-referenced across multiple independent sources (April–June 2026). Key facts — model names, pricing, tier limits, feature availability — verified against at least two sources before inclusion.
Node.js / docx npm (programmatic build)	Document production. The guide was built programmatically using a Node.js script and the docx library, not designed manually. All formatting, tables, colour, and layout defined in code.	The build script is the canonical source of the document. Each version produced by running the script, allowing reproducible output and version-controlled iteration.	Each build validated using the docx skill validation script. Document opened and reviewed in Microsoft Word after each material change.
Author (human)	All of the following: initial brief and philosophical positions; provision of source documents; all iterative direction, correction, and approval; structural decisions at each version; final content sign-off.	The author's input shaped every version. Claude produced drafts; the author decided what was kept, changed, removed, or added.	Not applicable — the author is the decision-maker, not a tool whose output requires external verification.

Record-keeping approach

Development was conducted in a single Claude project session, providing a continuous conversation log of all prompts, outputs, author corrections, and decisions. Versions were named sequentially (V1.0 through V1.7) and produced as dated .docx files. The Node.js build script constitutes the primary record of document structure and content decisions; it is version-controlled and reproducible.

Decisions made

All structural, content, and formatting decisions were made by the author. Key decisions included: adoption of the Plan–Do–Review framework as the document spine; addition of the stage-gate development pathway section; introduction of the individual vs team project distinction; removal of project-specific branding for generic publication; inclusion of this appendix as a modelling exercise.

What was discarded

A worked example referencing a specific student project context was drafted and subsequently removed by the author on the grounds that it was too context-specific for a generic document. An earlier version of the tools table contained only six tools and was replaced after the author requested a critical review and web search to verify current tool capabilities. Several draft phrasings were revised or rejected for tone — the author flagged and corrected instances of AI-sounding language throughout.

> **Note on this appendix**
>
> This appendix was itself drafted by Claude Sonnet 4.6 based on the conversation record of the document's development, then reviewed and approved by the author. It follows the same two-part Plan / Execution structure required of students in Section 9 of this guide.