

Increasing AI Literacy for Research Development Professionals

A Companion Guide

You Already Know Enough to Start —
And Starting Has Never Been More Urgent

Aimee Oke

Associate Director of Research, CVMBS

Colorado State University

aimeeoke.ai

Originally presented to the NORDP AI in RD Peer Mentoring Group | February 2026

How to Use This Guide

This companion guide is designed to stand alone. You don't need to have attended the live presentation to benefit from it. Each section maps to a concept from the original talk and includes the context you need to understand it, a visual or framework that makes it concrete, and one or two actionable tasks you can try this week.

Throughout this guide, look for  **Try This** boxes — these are specific, practical tasks designed to move you from reading about AI to actually using it. They're ordered from simple to more advanced, so start wherever feels comfortable and work forward.

The single most important thing I can tell you: the gap between reading about AI and using AI is where all the learning happens. Every expert I follow says the same thing. Start this week. Not next month. This week.

The Reality We All Know — Plus a New One

Before we talk about AI, let's be honest about where we are right now as research development professionals.

Death by 1,000 Cuts

Version control nightmares, cat-herding collaborators, reformatting data tables. Not hard tasks — just endless.

No Space to Innovate

Cognitive burden from repetitive work kills creative capacity. No room left for the strategic thinking that actually moves the needle.

More Hats, Fewer Hands

Reduced FTEs, expanded scopes, working nights and weekends. The portfolios got bigger; the teams didn't.

The Ground Is Moving

Tools and expectations changing faster than the workload lets you look up. This is the new challenge layered on top of the old ones.

The first three have been true for years. The fourth is new: the tools and expectations around us are shifting faster than our workload allows us to notice. That's what this guide is about — helping you look up.

Try This: The Frustration Inventory

Open any AI chatbot (ChatGPT, Claude, Gemini — even the free version). Type something like:

```
I'm a research development professional at a university. Here's what my typical week looks like: [describe your week – the tasks, the frustrations, the things that eat your time]. What patterns do you notice? Where do you see the biggest opportunities for me to reclaim time?
```

Don't worry about prompting "correctly." Just vent. Describe your week honestly. AI is surprisingly good at reflecting patterns back to you that you're too close to see. This isn't about getting AI to solve anything yet — it's about using it as a mirror.

The Pivot — And Then Another One

My AI journey started from burnout, not curiosity. After three training grants went in simultaneously, I was so exhausted I had to leave town. During the drive, I started listening to AI podcasts. That trip changed everything — not because of the technology, but because I realized the way we were working wasn't sustainable.

Then I read Ethan Mollick's *Co-Intelligence* and it clicked: this isn't a tech trend. This is going to shift everything. I paid \$20/month out of my own pocket for a ChatGPT subscription. The university wasn't going to cover it. But I thought: what's \$20/month to get rid of the drudgery?

The real shift came when I stopped using AI as an answer engine and started asking: where does this live in my actual workflow? Within a year, I went from pasting documents into ChatGPT to running a personal AI agent across multiple models. I'm telling you this not to overwhelm you — but because the distance between "I tried ChatGPT once" and where things are now is much larger than most people realize.

The takeaway: the on-ramp to AI doesn't require permission, a budget, or technical skills. It requires being frustrated enough with the status quo to try something different. Most of you are already there.

Try This: The \$20 Decision

Sign up for a paid AI subscription — Claude Pro or ChatGPT Plus (\$20/month). The free tiers are significantly behind. If \$20 feels like a stretch: if it saves you even 2 hours/month at your loaded hourly rate, it pays for itself in the first week. That's a case you can make to your supervisor.

Bonus: Start one AI podcast during your commute. Recommendations: AI Powered (Ethan Mollick), Lenny's Podcast (product-focused but transferable), or Latent Space (more technical). You don't need to understand everything. You need exposure.

What Changed in the Last Month

Most people — even people who follow AI — are operating with a mental model that’s 3-6 months out of date. In this environment, 3-6 months is a generation. Here are four developments from late 2025 / early 2026 that illustrate the pace:

1. **Autonomous AI Work:** A team of 16 AI agents autonomously developed a functional C compiler from scratch in two weeks — 100,000 lines of code. A task that traditionally took human engineers decades. In January 2025, AI coded autonomously for about 30 minutes. By February 2026, it coded for two straight weeks.
2. **The \$285B Wake-Up Call:** A 200-line text file — AI workflow templates from Anthropic — triggered \$285 billion in stock market losses in 48 hours. Wall Street priced in what practitioners already knew: AI agents can now do work that entire SaaS companies were built to do.
3. **Non-Technical People Building Software:** Two CNBC reporters with no coding background built a working project management app in under an hour. Total cost: about \$10. This category of work — “vibe coding” — literally didn’t exist six months ago.
4. **AI In Your Tools Right Now:** AI is embedded inside Excel and PowerPoint today. Norway’s sovereign wealth fund reported saving 213,000 hours annually. Goldman Sachs deploys AI agents across compliance — in production, not pilots.

Why this matters for you: the specific tools will keep changing. Your ability to adapt to them is what matters. The people who start building that adaptation muscle now have a compounding advantage.

Try This: See It for Yourself

Pick one of the four developments above and search for a full article about it. Read it. Then open your AI tool and ask:

What are the implications of this for someone who works in university research administration? Be specific about how this might affect my daily work in the next 6-12 months.

This practices using AI as a thinking partner rather than a search engine.

You Don't Need to Be an AI Expert

You need to be a domain expert who can delegate.

AI is making execution cheap. What's becoming exponentially more valuable is judgment: knowing what good output looks like, knowing which question to ask, recognizing when something is technically correct but strategically wrong.

Ethan Mollick at Wharton challenged MBA students to build startups in four days using AI tools. The students who succeeded weren't the most technical — they were the ones who knew how to **scope problems, define deliverables, and recognize when output was off**. Those are management skills. Those are research development skills.

We scope proposals, we define deliverables, we evaluate output, we know our compliance landscape. We've been training for this moment without knowing it. The shift is learning to direct that expertise through AI rather than executing everything manually.

Try This: Delegate One Thing This Week

Think of a task you do regularly where you already know what "done well" looks like. Delegate it to AI. Give it the task, review the output like you'd review a new hire's work, and iterate. The goal isn't a perfect result on the first try. The goal is to experience the delegation loop: instruct → review → refine.

The New Mental Model: Delegation, Not Prompting

This framework (adapted from Ethan Mollick) gives you a way to evaluate when AI delegation is worth your time. Three variables matter:

- **Human Baseline Time:** How long the task takes you manually.
- **AI Process Time:** Time to request + wait + evaluate the output.
- **Probability of Success:** How likely AI produces acceptable output on the first or second try.

If the task takes hours and AI produces a draft you can evaluate in minutes — delegate.

Concrete example: Reviewing a 40-page NOFO. Human baseline: 2-3 hours. AI process time: paste it in, ask for a structured summary, review: about 15 minutes. Probability of success: high (structured extraction is what AI does well). The math works. Now compare: writing a specific aims page. Similar time investment, but probability of success is much lower — because it requires deep knowledge of the PI's research and the review culture. **The framework helps you decide where to delegate, not just whether.**

You improve the equation three ways: better instructions raise the probability of success, better evaluation skills reduce cycles, and faster evaluation lowers your time investment. All three are improved by subject matter expertise. That's your edge.

Try This: Run the Delegation Math

Pick three tasks from your week. For each, estimate: (1) how long it takes you, (2) how long it would take to get AI output and review it, and (3) your gut feeling on probability of success (high/medium/low). Start with the task that scores best — long human time, short AI time, high probability.

The CRIT Method: Your Power Prompting Framework

The CRIT method, developed by Geoff Woods (author of *The AI-Driven Leader*), transforms AI from a content machine into a strategic thought partner. CRIT stands for Context – Role – Interview Me – Task. It's the technique I use most frequently for problem-solving and getting to the root of an issue quickly.

C – Context	Give AI rich background about your situation. What would a new colleague need to know to help you?
R – Role	Assign a specific expert persona: “Act as a senior NIH program officer reviewing T32 applications.”
I – Interview Me	Tell AI: “Interview me. Ask up to 3 questions, one at a time, to deepen your understanding.” This is the game-changer most people skip.
T – Task	Give a clear, specific assignment after the interview. “Based on our conversation, draft a...”

The Interview step is the game-changer. Most people skip it. They try to give AI all the information upfront, then wonder why the output is generic. When you tell AI to interview you, it pulls context out of your head that you didn't think to include — and it keeps YOU as the thought leader driving the conversation. As Geoff Woods puts it: “I don't ask AI questions. I turn the tables and have AI ask ME questions.”

CRIT in Action: An RD Example

CONTEXT: I'm an Associate Director of Research at a veterinary college managing NIH T32 training grants. We're preparing for a competitive renewal and I need to demonstrate trainee outcomes in a way that resonates with reviewers. Our trainees span DVM, PhD, and postdoc levels.
ROLE: Act as a senior NIH study section reviewer who has evaluated hundreds of T32 training grant applications, with particular expertise in assessing trainee productivity metrics.
INTERVIEW: Before you begin, ask me up to 3 questions, one at a time, to understand our training program's strengths and the specific challenges I'm facing with the renewal.
TASK: After the interview, help me develop a framework for presenting our trainee outcomes data that will be most compelling to reviewers.

Try This: Your First CRIT Prompt

Take a real challenge from your work this week and write a CRIT prompt for it. Open Claude or ChatGPT and paste in your Context, Role, Interview, and Task. Let AI ask you questions before it produces anything. Notice how much better the output is when AI has had a chance to understand your specific situation.

Sticky note challenge: Write CRIT: Context, Role, Interview, Task on a sticky note and put it on your monitor. Use it every day for a week.

Bolting On vs. Rebuilding

This distinction is one of the most important concepts in this guide.

Bolting On	Rebuilding
<p>Using AI to speed up your existing process</p> <p>Paste a NOFO into ChatGPT, get a summary faster</p> <p>Use AI to draft an email you would have written anyway</p> <p>Ask AI to proofread a document</p> <p>Feels productive. Saves some time. But your workflow is fundamentally unchanged.</p>	<p>Rethinking the workflow from scratch with AI</p> <p>How would I design my NOFO review if AI handled 70% of extraction?</p> <p>What if I could run competitive landscape analysis in 20 minutes?</p> <p>What if I could build custom tools for each PI's specific needs?</p> <p>Changes what's possible. Changes your role.</p>

My example: I used to manually track T32 trainee publications in a spreadsheet. Bolt-on: I'd ask ChatGPT to help format citations. Useful, saved 30 minutes. Rebuild: I built a system that pulls publication data, cross-references it with our trainees, categorizes by grant period, and generates the tables NIH needs for the progress report. Same underlying task — completely different scale of impact.

Most people start by bolting on. That's fine — it's how you learn. But know that the bolt-on is the beginning, not the destination.

Try This: Easy Bolt-On Wins

Pick one of these and try it this week:

- Paste a NOFO and ask: "Extract all eligibility requirements, deadlines, budget caps, and page limits into a structured table."
- Take a rough budget justification and ask AI to "polish this for NIH formatting and tone."
- Paste a faculty draft email and ask: "Help me make this more concise and professional."

Then tape this question to your monitor: "What would the rebuild look like?" You don't have to answer it yet. Just start asking it.

The Four Skills That Thrive

These four skills determine who gets the most leverage from AI. Notice: three of the four are things you already have.

Taste The instinct for whether something is right — not by checklist, but by hard-won experience. You read a draft specific aims page and know in 30 seconds if it will land with reviewers.	Domain Judgment Specific, contextual understanding of how your domain works. You know which program officers align with your strengths and which faculty partnerships will survive a multi-year project.
Phenomenal Ramp The ability to learn fast when everything is evolving. The tool you learn this month may be replaced in three. The skill of rapidly adapting compounds. This is the new muscle.	Relentless Honesty Looking at your own work and asking: which parts require my judgment, and which could AI handle? If you spend 10 hrs/week on tasks AI handles in minutes, your time allocation doesn't match your value.

Try This: The Honest Audit

Block 20 minutes. Write down every task you did this week. Next to each, mark:

J = Requires my judgment (taste, domain knowledge, relationship context)

E = Execution (formatting, extracting, drafting, organizing)

B = Both (the execution portion is a delegation candidate)

Be ruthlessly honest. Everything marked E is a delegation candidate. This exercise often reveals that 30-50% of your week is spent on execution AI could handle.

Learning to See the Shape of Problems

I call this “parkour vision.” In parkour, practitioners develop the ability to look at urban architecture and see paths for movement that other people don’t notice — a wall becomes a launch point, a railing becomes a path. It’s a *paradigm shift in perspective* that changes how they view structures, obstacles, and space, enabling them to see potential for movement everywhere.

The same thing happens when you use AI regularly. You start to see the “shape” of problems differently. Not every problem is a chat problem. Some are tool-shaped, some are spreadsheet-shaped, some are software-shaped. Developing the instinct for what shape a problem is — that’s the new literacy.

I notice this...	It's shaped like...	So I reach for...
I keep doing this repetitive thing	Quick automation	A chat conversation
I need this analysis done the same way every time	Reusable tool	Custom GPT or saved prompt
I need to work with data in a structured way	Spreadsheet problem	Claude in Excel or AI-assisted analysis
I need something that doesn't exist	Vibe-code project	Builder platform with a clear spec
Complex, multi-step, reliable over time	Engineering problem	PRD + structured build process
This task should run itself	Agent problem	Defined agent with boundaries

You already do this. Every time you look at a researcher’s workflow and think “there has to be a better way to manage this data,” or “why are we entering the same information into three different systems” — that’s software vision. The difference now is that you can actually build the solution.

I didn’t start at the bottom of this table. I started with chat. The other shapes revealed themselves through practice. They will for you too. The friction you feel when a tool doesn’t quite fit — that’s the signal to look for a different shape.

Try This: Shape-Spotting

Keep a running list this week of moments where you think “this is tedious” or “there has to be a better way.” Don’t solve them. Just write them down. At the end of the week, try matching each one to a row in the table above. Most will be chat-shaped — and that’s perfect.

Level up: If you find yourself giving AI the same instructions repeatedly, turn it into a Custom GPT (ChatGPT) or a Project with custom instructions (Claude). Write once, reuse forever. That’s the step from row 1 to row 2.

Where Are You Right Now?

Use this quick self-placement to identify your starting point. (For the full self-assessment tool, visit <https://atomgrants.com/resources/the-ai-readiness-self-assessment-for-research-leaders>)

Ask yourself: How often do I use AI tools for actual work tasks (not just experimenting)? Can I describe what makes a good AI output in my domain? Have I modified any workflow based on AI capabilities? Rate yourself honestly:

Score	15-35	36-55	56-75
Level	Early Explorer	Emerging Practitioner	AI-Forward Operator
You are...	Curious, but AI isn't part of your daily rhythm yet.	Using AI in pockets, starting to see real value.	AI is amplifying your capacity. Thinking strategically.
Your Goal	Get one real delegation win this month.	Start rebuilding one workflow.	Build the adaptation habit.

No matter where you are, the goal is the same: move one level forward in the next 30 days. Not five levels. One.

Try This: Calendar Your Practice

Right now — not later, now — put a recurring weekly block on your calendar. Label it “AI Practice — Real Work.” Thirty minutes minimum. During this block, use AI on an actual task from your job, not tutorials or articles.

Start Here. But Know Where This Goes.

The Floor (Start Here) Read a 40-page NOFO manually: ~2-3 hours Paste NOFO into AI, review summary, send to faculty with your edits: ~15 minutes This is real. This works today.	The Ceiling (Where This Is Going) AI monitors new funding opportunities automatically Cross-references against faculty profiles Flags high-probability matches with briefs Ready before you get to your desk Monday This isn't science fiction. People are building this now.
--	---

The gap between the floor and the ceiling isn't about technical skill. It's about practice, curiosity, and a willingness to keep pushing the boundary of what you delegate.

Try This: The NOFO Speed Run

Find a NOFO relevant to your work. Paste the entire text into Claude or ChatGPT and use this prompt:

```
I'm a research development professional reviewing this funding opportunity. Please extract and organize: (1) Eligibility requirements, (2) Key dates and deadlines, (3) Budget limits and restrictions, (4) Required application components and page limits, (5) Review criteria and scoring weights, (6) Anything unusual I should flag for faculty.
```

Review the output. Edit it. Send it to a faculty member with your professional assessment added. Time yourself. That delta is your proof point.

When Execution Is Free, Judgment Becomes Priceless

AI cannot know that Dr. Chen's lab has unpublished preliminary data that makes a weak proposal strong. It cannot recognize that two co-PIs have a history that makes their "collaboration plan" fiction. That knowledge lives in you.

If you're worried that AI is going to make your job irrelevant, I want to push back firmly. The opposite is happening. AI is making execution cheap and judgment expensive. Every source I've read converges on the same point: the scarce resource is no longer "who can produce the artifact." It's "who knows whether this is the right artifact."

Without you, AI produces polished noise. With you directing it, AI produces leverage.

Try This: The Judgment Test

Take something AI generated for you (from any exercise above). Before accepting or rejecting it, write down THREE things that are wrong, missing, or could be improved — things only someone with your specific domain knowledge would catch. This proves to yourself that your expertise makes AI useful, and trains the critical evaluation muscle.

What's Actually Stopping You?

Let's address the real barriers head-on:

- **"Budget constraints"** — Free tools work for experimenting. \$20/month unlocks the real capability. If it saves 2 hours/month at your loaded rate, it pays for itself immediately.
- **"Institutional resistance"** — Start with your own work. Build proof. The best argument for AI adoption is showing what you accomplished.
- **"Hallucination fears"** — Valid concern. This is exactly why domain expertise matters. AI doesn't need to be perfect. It needs to be good enough that review time is shorter than creation time.
- **"It's moving too fast"** — There is no settled state coming. Waiting to start is the most expensive decision you can make.
- **"Feeling like I'm already behind"** — You're not. The vast majority of RD professionals haven't started at all. Starting now puts you well ahead.

Try This: Build Your Proof Point and Find a Buddy

This week, complete one AI-assisted task and document: what the task was, how long it normally takes, how long it took with AI, and quality of output. Save this. You'll need it.

Then find one person and commit to a monthly check-in: "What did you try? What worked? What didn't?" Social accountability turns experiments into practice.

Your Monday Morning Move — And Every Monday After

This isn't a one-time exercise. This is a practice.

This Week

Pick one real task. Delegate it to AI. Evaluate the output like a new hire's work. Iterate.

Every Week

One hour per week with AI on real work. Not reading about AI. Not watching tutorials. Using it on actual tasks. The difference between week 1 and week 8 is dramatic.

Every Month

Set a calendar reminder: "Is there a better way to do this now?" The tools change every few months. Build in adaptation.

*The people who will be furthest ahead in six months aren't the most technical.
They're the ones who started now and didn't stop.*

A practitioner I follow uses a bike metaphor: with AI, going faster actually makes it easier to balance. The people who try to go slow — dipping in once every few months — have a harder time than people who commit to regular practice.

Try This: The Accountability Calendar

Set these three calendar items right now:

- 1. This week:** 30-minute block — "First AI Delegation"
- 2. Next week:** 1-hour block — "AI Practice — Real Work" (make recurring)
- 3. Next month:** 15-minute block — "AI Adaptation Check" (make recurring)

Keep Going

A year ago, I was sitting where many of you are. I was overwhelmed by my workload, skeptical about whether AI was really relevant to research administration, and honestly a little afraid of looking foolish trying something new.

That \$20 decision changed my career. Not because the tool was magic — but because I gave myself permission to start. And then I didn't stop.

None of the resources below matter if you don't start. The single most important thing you can do is try one task this week.

Resources

- aimeeoke.ai — Open-access frameworks, tools, and guides built for research professionals
- **Understanding LLMs: A Visual Guide** — An illustrated walkthrough of how large language models actually work (aimeeoke.ai)
- **AI-Readiness Self-Assessment** — Take it now, revisit in 30 days, and track your growth (co-created with Atom grants, linked at aimeeoke.ai)
- **Co-Intelligence** by Ethan Mollick — The best starting point for understanding why AI matters for knowledge workers
- **The AI-Driven Leader** by Geoff Woods — Where the CRIT framework comes from; essential reading for using AI strategically
- **“Management as AI Superpower”** by Ethan Mollick (January 2026) — The article that reframes AI as a delegation skill
- **The \$20 Starting Point** — Claude Pro or ChatGPT Plus. Use the best model available. Change when better ones arrive.

You already know enough to start.

The only thing between you and a fundamentally different relationship with your work is the decision to try — and the commitment to keep going.

Aimee Oke

aimeeoke.ai | aimeeoke@gmail.com | aimee.oke@colostate.edu

I answer every email. Reach out when you get stuck. Reach out when you have a win.