PHILIP DENG

The AI *Gap*

Everyone is already behind.
Here's the bigger picture, and how to catch up.

Most of us are using AI by now. We found a few techniques we like, a prompt that works, a task it does well, and we stuck with them. Then life happened. Meanwhile the frontier kept moving, further than most of us realize. In 2023 these tools answered one question at a time. Today they read whole documents, research the web, use software, and work for hours on their own.

That's the gap: between what we use AI for and what it can do now. In one survey of grant pros, the top use is still trimming character counts. What we think it can't do is often a year or two out of date.

My hypothesis is that the gap is mostly overwhelm. There's so much, so fast, and we're hardworking, tired people. Five years ago I was a grant pro who had never written code. I got here by thinking clearly about the problem, talking it through with AI, and asking a lot of how-tos. This guide is the talk, made usable. Every hard part comes with something you can do.

PAGE 2
Where the gap is
What each task really takes, and a worksheet for your week.

PAGE 3
Build the system
Capture, a map, and loops that get better.

PAGE 4
Check it at scale
Review systems, evals, and one report start to finish.

PAGE 5
Reservations
Energy, privacy, ethics, detection, and where we aim.

PAGE 6
Start here
A four-week start, and the business side.

Assume it can do every task.
The real question is what it takes.

The opening poll asked which of these AI can do well today. Here's how I'd answer, sorted by what it takes to get it done well.

Summarize a 40-page RFP	TRIVIAL
Trim a draft to the character count	TRIVIAL
Read a funder's stated priorities	TRIVIAL
Write a narrative a funder prefers to a pro's	EASY
Read between a funder's lines	YOUR CONTEXT
Make a hard ethics call	YOUR CONTEXT
Make a tough go / no-go call	YOUR CONTEXT
Build the budget	A SYSTEM
Find funders that fit	A SYSTEM

TRIVIAL
Any AI chat, one clear request.

EASY
A good prompt, plus your own materials.

YOUR CONTEXT
It gets there once it knows what you know: your history, your funders, your community.

A SYSTEM
Data, tools, and checks built around the AI. Pages 3 and 4.

WORKSHEET · YOUR WEEK, TASK BY TASK

Do the same with your own week.

Look at your job the way the Anthropic Economic Index looks at every job: not by title, task by task. GPA's own survey found about 41% of a grant pro's week goes to writing, and writers rank 3rd of 785 jobs for overlap with what AI can do. List the seven tasks that take the most of your time, and mark what each one takes. Hand off the trivial ones this week. The ones that need a system are what the next two pages are for.

TASK	TRIVIAL	EASY	YOUR CONTEXT	A SYSTEM
	☐	☐	☐	☐
	☐	☐	☐	☐
	☐	☐	☐	☐
	☐	☐	☐	☐
	☐	☐	☐	☐
	☐	☐	☐	☐
	☐	☐	☐	☐

41% writing: GPA member survey, 2023 · Writers' rank: Microsoft Research, 2025 · Method: Anthropic Economic Index · Task levels: the speaker's judgment, October 2026

The people furthest ahead aren't better at prompting. *They built a system around the AI.*

You	DIRECT AND DECIDE
Agents	WITH SKILLS + TOOLS
The graph	HOW IT CONNECTS
The lake	EVERYTHING, RAW

Most of us work at the top: one person typing into one chat. The value is underneath. Everything you know lands in one place, the lake. The AI draws a map of how it all connects, the graph. Then agents work from that map, using skills you've written down and real tools. Almost nobody in our sector has built this yet, so starting now puts you ahead.

1 Capture everything CONTEXT ENGINEERING

Every call, email and report is signal. Some of it is already data. The rest evaporates unless you catch it.

ALREADY DATA · ROUTE IT	NOT YET DATA · CAPTURE IT
Funder emails	Program officer calls · *record and transcribe, with consent*
Drafts and submissions	Site visits · *a voice memo on the drive home*
Budgets and actuals	Why you passed on a grant · *one line, said out loud*
Awards, declines, reviewer notes	Hallway decisions · *a 30-second voice note*

START THIS WEEK
Pick one place where all of it lands, raw and unedited. Organizing it is the AI's job.

2 Let AI make sense of it KNOWLEDGE GRAPH

It's two ideas. **Nodes** are things: a funder, a person, a grant, a program, a decision, a deadline. **Edges** are how two things connect: works at, awarded, pays for, knows, we passed. Don't build the web yourself. The AI reads the raw pile, draws the links, and keeps them current.

TRY IT
Give an AI a month of your notes and emails. Ask it to list every funder, person and deadline in them, and how each one connects to the others.

3 Build loops that get better SKILLS

It does the work → You review it → The fix goes into how it works → The next run starts better

A prompt is a one-off. A loop is a habit. A skill is just your lessons written down, in a file the AI reads every time. For example: *Funder reports: outcomes first, one real story, answer what they asked. Voice guide: how we sound, and what we never say. Numbers rule: every figure points back to its source.*

START THIS WEEK
The next time you correct the AI, don't just fix the draft. Write the fix into a short instructions file and give it that file every time.

AI made output cheap. *Checking is the new bottleneck.*

Anthropic says it about its own engineers: “human code review became the new bottleneck.” Grant work is next. Checking harder doesn't scale. A review system does.

a Never let the maker be the checker

One AI builds the draft. A separate AI, in a fresh session, attacks it with the opposite goal: find what's wrong. It never saw the drafting, so it has fresh eyes and no shared blind spots. Give it the funder's own scoring criteria.

b Put your eyes at the seams, not on every line

Split the work into runs that each start fresh and read only the last output. You check every handoff, and give the final draft a hard read before it ships.

Plan from your ask → YOUR CALL → **Research** from the plan → YOUR CALL → **Draft** from the research → YOUR CALL → **Review** attacks the draft → HARD READ

c Build evals: make checks pass or fail

A fuzzy check needs your judgment every time. A pass-or-fail check can run on its own, every time.

Is it compelling?	**Does every outcome have a number and a date?**	PASS | FAIL
Is the budget reasonable?	**Does the budget total match the narrative?**	PASS | FAIL
Did we follow the rules?	**Page limits and file names: a fixed rule. Budget math: a calculator, never a guess.**	PASS | FAIL

Make most of it pass or fail. Save your judgment for what's left: the story and the voice.

d Hand off more, one ring at a time

Widen what the AI can touch only as your checks earn trust.

1 · Proposes only
Suggests. Changes nothing on its own.

2 · Its own files
Drafts and edits only what it made.

3 · Your files
A stage, a folder, a status. Never your words, and it can be undone.

4 · Works outside
Sends, submits, spends. Always asks you first.

ONE REPORT, START TO FINISH · ILLUSTRATIVE, FICTIONAL FUNDER

You ask
“Alder's year-end report is due Friday. Draft it like last year's and flag anything off.” No files, no prompt.

It reads the map
Everything tied to this grant, by following the links out from the funder.

It gets things
Last year's report, the award letter, three attendance sheets, the budget. About two minutes.

It loads skills
Your funder-report rules, voice guide and numbers rule, before it writes a word.

It checks itself
Sources, budget, outcomes, word limits: pass. **Students served: 151 here, 142 in Q3. Fail.**

You answer one question
“151. Q3 was before the summer group.” Fixed, and noted for next year.

Still have reservations? Fair. *Here they are, one at a time.*

ENERGY

Not yes or no. How.

1. **Next to everything else we do, it's a rounding error.** One seat on a flight across the Atlantic is about the same energy as 400 prompts a day for 80 years. A year of heavy AI use is about a space heater running for 30 minutes.
2. **"Ration your prompts" is an old move.** The idea of a personal carbon footprint was popularized by an oil company's ad campaign. Both aim the light at you, and away from the system.
3. **The real fight is policy.** Data centers, the grid, solar and wind, at every level of government. Our organizations and funders have a voice in all of it.

PRIVACY

Not trust or don't. Know where it goes.

1. **It learned before you typed.** The model was trained on the public internet. Your one prompt didn't make it smart. The question is what happens to your words next.
2. **Your words take one of three roads.**

A free app	*may learn from your chats*
A paid business plan	*terms: no training*
Your own computer	*never leaves the building*

3. **Your work already lives in someone's cloud:** your email, donor lists, drafts. Read the terms of the plan you actually use, from a company you'd trust with your email. Terms change, so check again now and then.

ETHICS

No new ethics. Old risks, made easier.

1. **My hot take,** and I'm open to having my mind changed. **Errors creep in:** an old problem, now copied at AI speed. **Spray-and-pray:** an old problem, now it costs nothing. Our code of ethics still works. Let's move forward with it.
2. **The bigger question is who gets our help.** About 1 in 10 organizations hand grant seeking to a contractor. 43% of nonprofits that file a 990 are run entirely by volunteers. Expertise and capacity keep flowing to the organizations that already have them.
3. **For a small org, every step is the hard part:** finding funding, building relationships, applying, reporting back. AI is lowering those walls. That lets stuck organizations through, and spray-and-pray too.

DETECTION

"Did you use AI?" Nobody can reliably tell.

Answer honestly and risk a penalty. A better question for both sides: is it accurate, true, and yours?

26%
of AI-written text was caught by OpenAI's own detector. They shut it down in 2023 for low accuracy.

61%
of essays by non-native English writers were flagged as AI by seven detectors, on average.

2 in 3
foundations use AI themselves. 98% have no AI policy for grantees.

WHERE WE AIM · THE RACE TO THE TOP

One direction is a race to the bottom: agents writing for agents, funders hiding, everyone flooded. Nobody wants that one. Here's the other direction.

1 · Match
Worthy organizations, new and mature, meet the funders who want their work, before anyone writes a word. Grant pros build the new paths.

2 · Better
AI welcome on both sides. Applicants make a true, compelling case. Funders give a fair read and deep, helpful feedback.

3 · Fewer
One application, one review, shared. Other funders can award on a reputable funder's review. No more writing the same application over and over.

The point is more access, to funding and to expertise, for the organizations that never had either. Some of it is already working: Lever for Change hands donors a vetted bench of 620 organizations.

Energy: Andy Masley and Hannah Ritchie, via "Our sector is getting AI and energy wrong" (Grantable, 2026) · Footprint campaign: BP and Ogilvy, 2004–06 (NPR/WBUR, 2023) · Contractors: GrantStation, State of Grantseeking 2026 · Volunteer-run: Candid, Aug 2026 · Detection: OpenAI, 2023; Liang et al., *Patterns*, 2023; CEP, Sep 2025 · Lever for Change, 2026

Build an AI practice. *Start this month.*

EARLY IN YOUR CAREER?

This matters most for you. You'll do this work alongside AI for decades. Build the practice now, while you're still learning the craft, and the two grow together.

A FOUR-WEEK START

WEEK 1

Hand off the trivial

Take the trivial tasks from your worksheet and give them to AI completely. Pick the one place everything lands.

WEEK 2

Write your first skill

A voice guide, or your rules for funder reports. One page the AI reads every time.

WEEK 3

Add a reviewer

A fresh session that attacks your next draft against the funder's own criteria.

WEEK 4

Make three checks pass or fail

Turn three fuzzy questions you always ask into checks that run on their own.

LOVE THE CRAFT? FIND YOUR HUMAN HALLMARK

Hand off
the parts that drain you.

Protect
the parts you love.

Make it your name
the writing, how you steward a client, the relationships, the professionalism.

My human hallmark: _____

RUN A GRANTS BUSINESS? THE MODEL MOVES TOO

PRICING

by the hour → by the result

When a draft takes an hour instead of a week, hours stop being the product.

DELIVERY

drafts → systems they keep

Leave each client with their own lake, skills and review checks.

SCALE

a few clients → more orgs

Including the small ones that could never afford you before.

COMPETITION

other consultants → their own AI

Still yours: judgment, relationships, your human hallmark.

ABOUT

Philip Deng is a grant pro first. He spent his early career in nonprofits as an executive director, development manager and grants consultant, then co-founded Grantable in 2020, where he is CEO.

philipdeng.com · grantable.co

SOURCES

Cover: tool timeline from OpenAI, Google DeepMind, Anthropic and METR; uses from Cyndi MacKenzie Consulting, *Grant Writing in the AI Era* 2024–2026 (n=191, self-selected). Page 4: Anthropic on code review, 2026. Page 2 and 5 sources are on those pages.