



StanfordSOCIAL
INNOVATIONReview

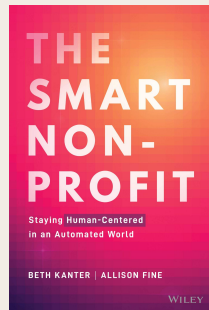
WHAT'S NEXT *for the* SOCIAL SECTOR?

Strategies and Tactics for Today's Agile Leaders

September 17-18, 2024

The Ethics of Workplace AI

September 18, 2024
11:45am – 12:45pm



Beth Kanter

Author, Speaker, and Trainer

www.bethkanter.org




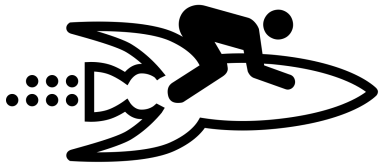
Agenda

- Welcome & Polling
- AI Fundamentals
- Acceptable Use Policy
- Working with Generative AI
- Takeaways



Resources: <https://bit.ly/Beth-Stanford>

Stages of Nonprofit AI Maturity

CRAWL	WALK	RUN	FLY
			
<p>Exploring possibilities</p> <p>Shadow use</p>	<p>Low risk experiments to test waters and gain hands-on experience</p> <p>Interim Acceptable Use Policy</p>	<p>Organizational AI adoption strategy to integrate into productivity, CRMs, and other tech stacks and scale org use. Data-driven projects.</p> <p>AI ethical data governance policy</p>	<p>AI strategically aligned with outcomes, org strategy, program delivery, and capacity to drive transformative results.</p> <p>Ethical pipeline</p>

slido



Where is your organization in terms of AI Maturity?

① Click **Present with Slido** or install our [Chrome extension](#) to activate this poll while presenting.

AI Fundamentals

3 Common Types of AI Used by Social Sector



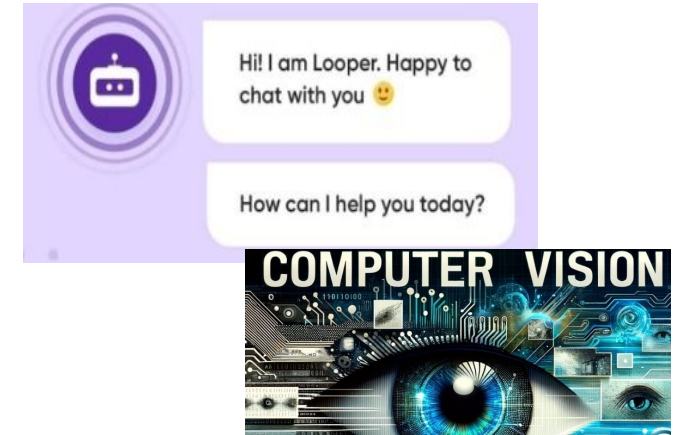
LLMs

Tools to assist staff with writing, summary, note taking, and other tasks.



Predictive Analytics, Machine Learning, and Enterprise LLMs

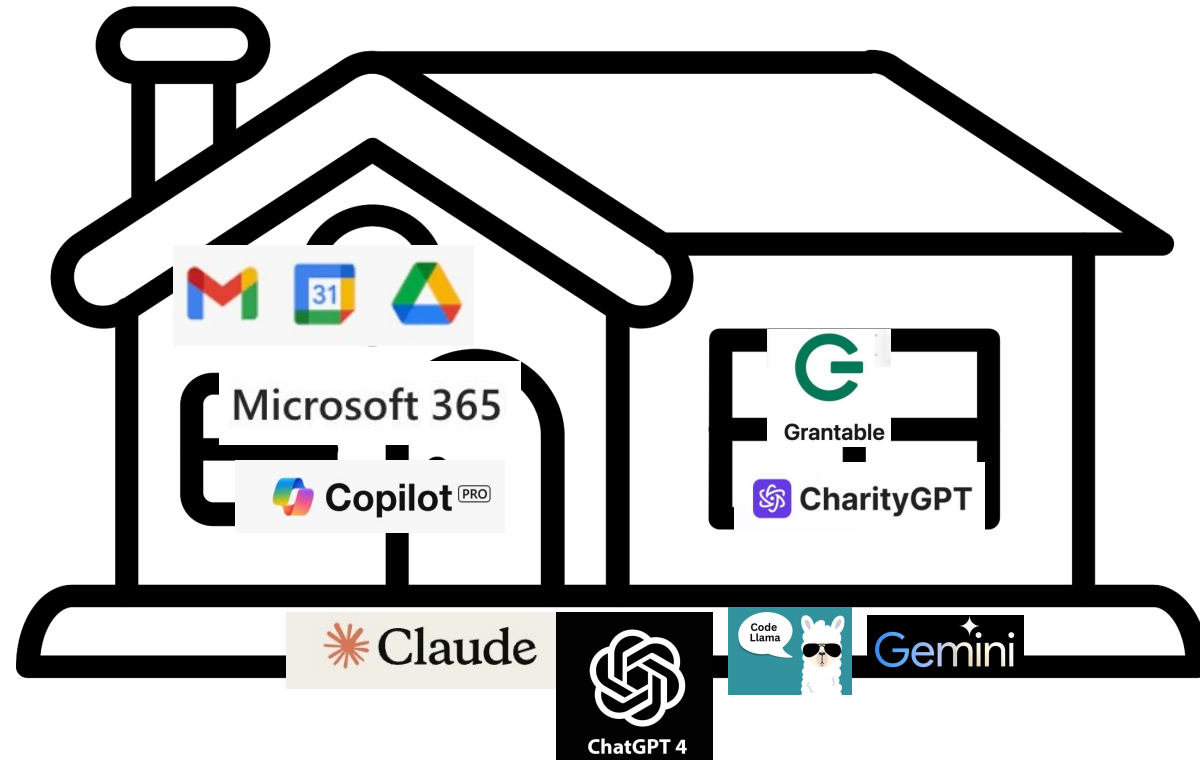
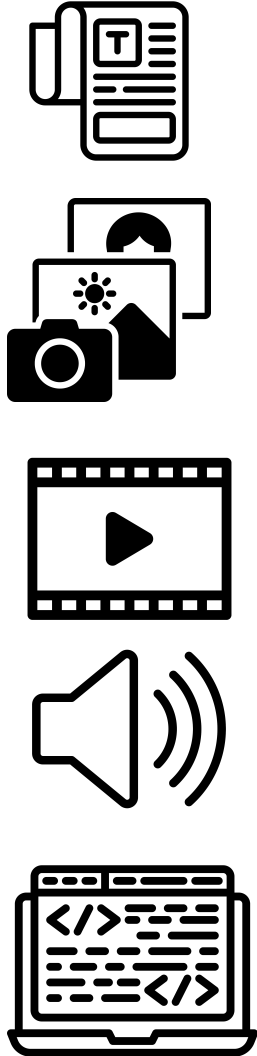
Integrated into organizational tech stacks - productivity suites, enterprise level LLMs, donor databases, operational applications, and others.



Generative AI, LLM & ML for program delivery

Integrated into web sites or building apps or agents to interact directly or with stakeholders to deliver data-driven programs, services, and information at scale.

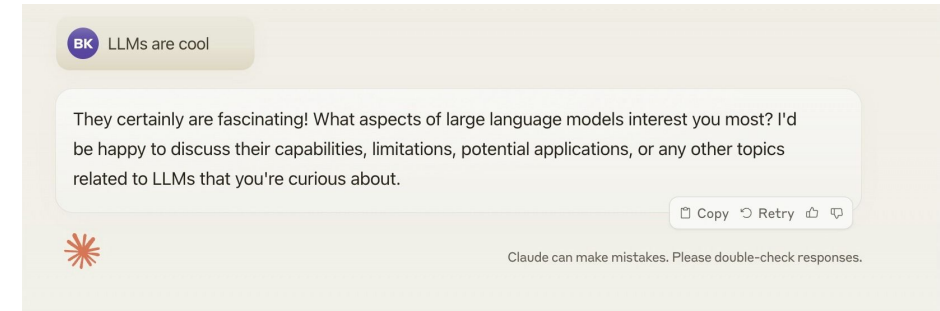
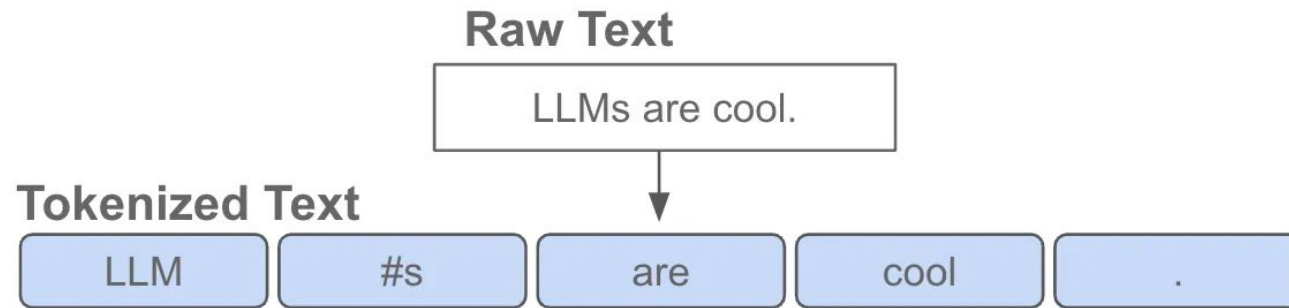
Generative AI is Powered by LLMs



Ethan Mollick
@emollick

Remember, today's AI is the worst AI you will ever use. The writing will improve, the amount of words the AI can hold in memory will improve (making stories more coherent), and the costs will drop.

How Do LLM's Work?



1. Split input text into tokens
2. Calculate a semantic understanding of those tokens
3. Predict the right sequence of tokens to output as a response

Different models from different companies have different “tokenizers” and semantic maps due to training differences, leading to different performance

Human-Centered AI

The Dividend of Time



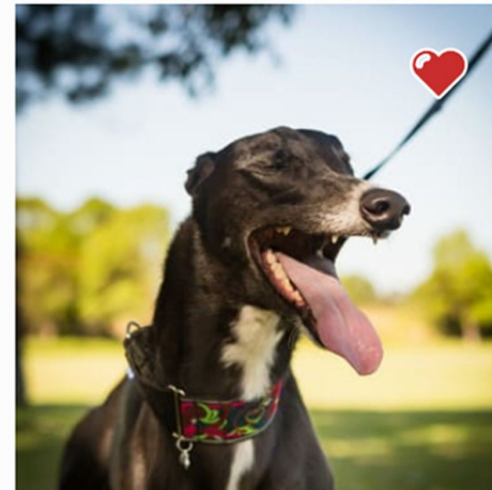
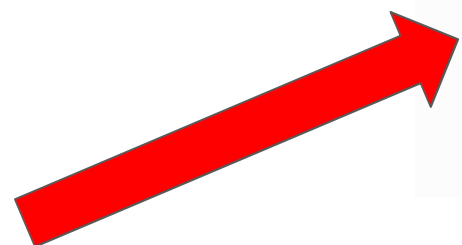
Introducing

Rescue Writer™



The world's first AI author that can instantly write a shelter pet's profile for you. Simply input their name, type of pet, and choose from a list of traits. Rescue Writer™ will do the rest – making your job easier and your pets more adoptable.

<https://rescuewriter.ai/generate>



Snaffles ~ 17 Month Old

Greyhound

Happy tripawd

I'm 17 month old greyhound called Snaffles, Snaffs (or Snafu when naughty). Apparently I'm characterised as "special needs" because I'm a tripawd. Pfff, no special needs here except I'm the doggy equivalent of a teenager.

Human-Centered Adoption: Reorient Workflows

From Outputs	To Outcomes
Hours worked	Creativity and innovation
Tasks	Focus on stakeholders, community, donors interests
Efficiency	Contributions to organizational mission
Get things done	Improved quality of workflows & adaptability

77%

of employees say
AI has added to
their workload.

REORIENT

Reclaim

Protect

Elevate

Improve

Co-Intelligence: A Different Way of Working

Examples of Generative AI Use Cases

DRAFT/EDIT

Grant proposals
Feedback to Stakeholders
Impact reports
Newsletters
Program descriptions
Social media content
Conference presentations

SUMMARIZE

Documents
Open-ended Feedback
Policy briefs
Strategy background documents
Research papers

ANALYZE

Feedback trends
Event surveys

BRAINSTORM IDEAS

Event Themes
Program ideas
Cross-sector collaborations

ADMIN

Team meeting notes
Site visit notes



Share Pair



- **Dividend of Time:** How might you reinvest the dividend of time into mission critical work that supports outcomes, relationships, and continuous improvement?
- **Leadership Modeling:** What is a pain point, takes a lot of time, scatters your attention?

Generative AI Limitations

slido



What are your biggest concerns about using generative AI in your workplace?

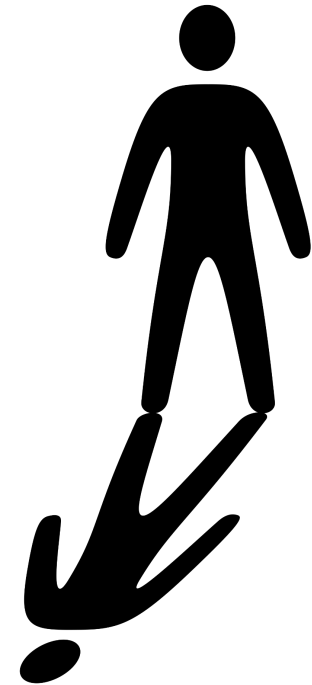
① Click **Present with Slido** or install our [Chrome extension](#) to activate this poll while presenting.

Limitations: Generative Ai Tools

- Accuracy
- Bias
- Equity & Access
- Privacy & Confidentiality
- Intellectual Property
- Transparency



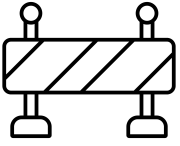

Why you need an acceptable use policy for generative AI right now

- Shadow use or “Bring Your Own AI To Work”
- Unchecked use could lead to ethical missteps
- Open adoption leads to better outcomes



Generative AI Acceptable Use Policy

Generative AI Acceptable Use Policy

	1: Principles
	2: Norms
	4: Guardrails
	5: Implementation

- A discovery process that leads to a written, formal policy document
- Living document
- Fits your culture, context, values, and mission
- Sets stage for organizational strategy & ethical frameworks

Resources: <https://bit.ly/Beth-Stanford>



Ethical Principles: Start Here

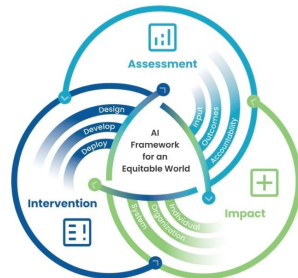
- Contextualize your organization's values to generative AI Use
- Understand the limitations
- Co-create policy with your team



Artificial Intelligence Framework for an Equitable World

Artificial intelligence technologies have been in the news for some time, often with sensational stories of incredible promises, disastrous impact, or even world-ending doom. In 2023, NTEEN and our partners at [Institute For the Future](#) began a community process to convene dozens of people from diverse nonprofit missions, foundations, technology companies, and nonprofit advisors. Through this community process, we surfaced a few key points:

1. Everyone, regardless of their work or organization, has the responsibility to understand how to build or use AI tools.
2. Real harm comes from how AI technologies are currently built and applied to our missions, which we need to understand and be accountable for.
3. Contemplating the creation or use of AI tools is complex, and we benefit from taking complex approaches to those considerations.



nten

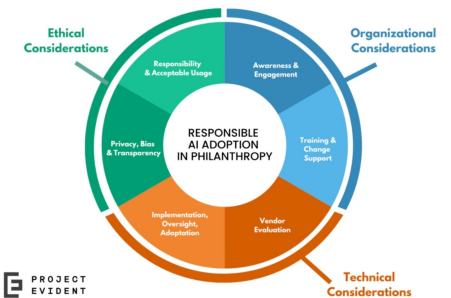
AI Governance Framework for Nonprofits

An Overview

Presented by ANB Advisory Group
Afua Bruce and Rose Afriye



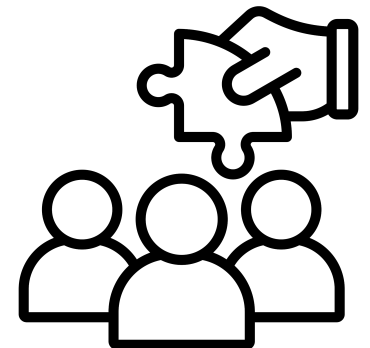
FundraisingAI



PROJECT EVIDENT
TECHNOLOGY ASSOCIATION & GRANTMAKERS

Norms: Shared Expectations for Using Tools

- Thoughtful workflow integration
- Staying up to date
- Values-aligned tool selection
- Use Cases, Provisioning



Norms: Use Cases & Tool Selection

Use Cases

- What are use cases for generative AI tools? What should it not be used for?

Tool Selection

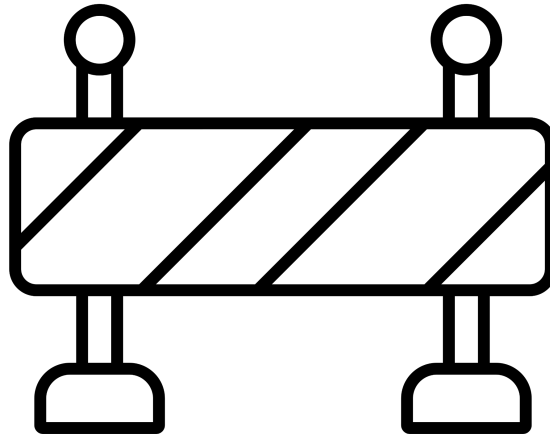
- Which generative AI tools do you suggest or require employees to learn and use? Why?
- Are there any generative AI tools that employees should NOT use? Why?
- Do staff need approval to use generative AI tools?
- How will you evaluate and select generative AI tools for staff to use?

Provisioning

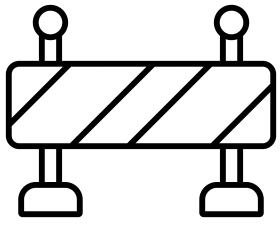
- Will organization reimburse or pay for generative AI tools?
- In addition to staff, must all contractors or volunteers be trained on and adhere to the acceptable use policy? Are there exceptions?

Guardrails

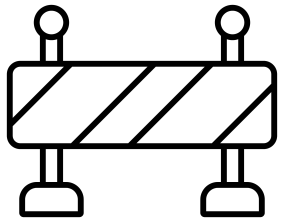
- Guardrails are not about saying “no” to AI use but “yes” to using it responsibly given the limitations



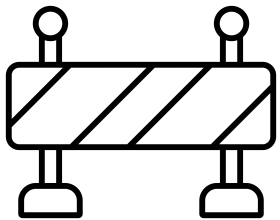
Examples of Guardrails



Bias: Always review output with an equity lens, assess if outputs are one-sided or skewed, and examine language in prompts for unintended bias.

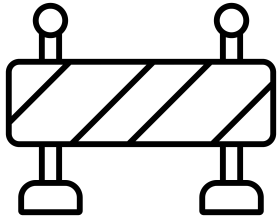


Accuracy Establish procedures for regular audits of generative outputs by use case. Verify information using critical thinking, human judgement, and outsources.

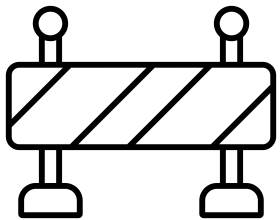


Privacy: Create specific stop, pause, go rules for use of data in prompts or attached documents. Opt out of LLM using your data for training.

Examples of Guardrails



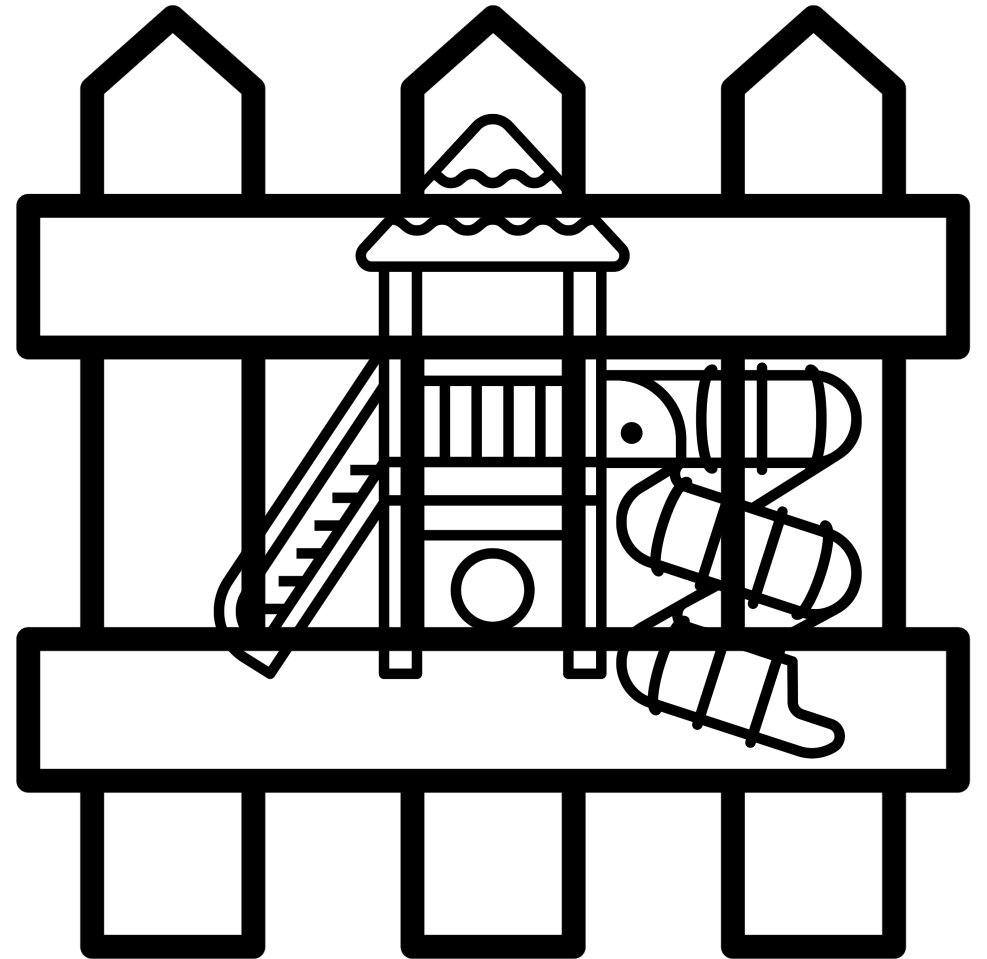
Intellectual Property: Don't share copyrighted materials with LLMs. Avoid using verbatim copyrighted text in prompts. Review and modify LLM-generated content for unique expression. Stay informed on IP and AI regulations.



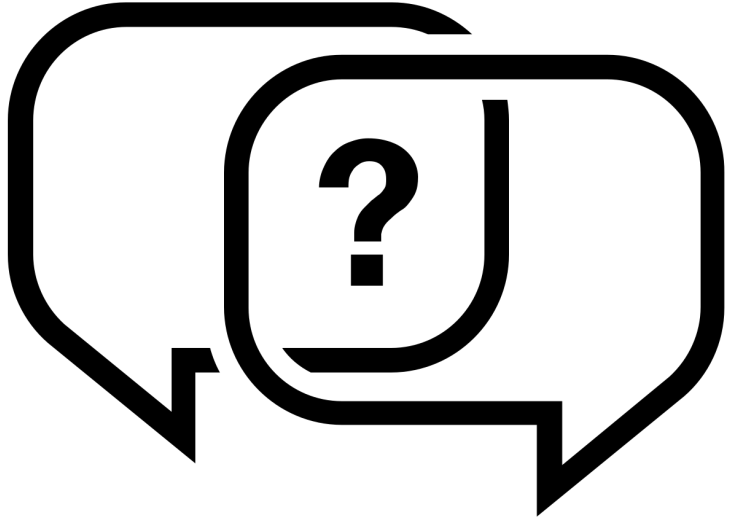
Transparency: Determine how to disclose for AI use for externally facing content.

Implementation

- Training
- Champion Group w/ Power Users
- Support for Experimentation
“Playground with a High Fence”
- Playbook & Peer Learning
- Policy updates and review



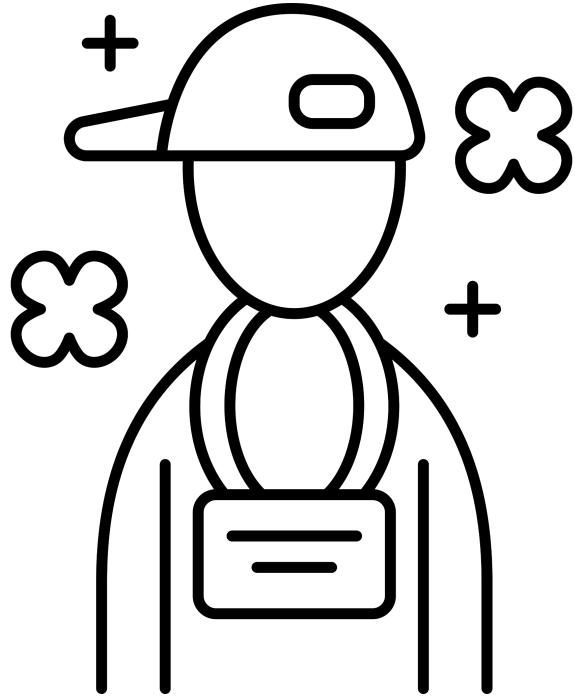
Reflection



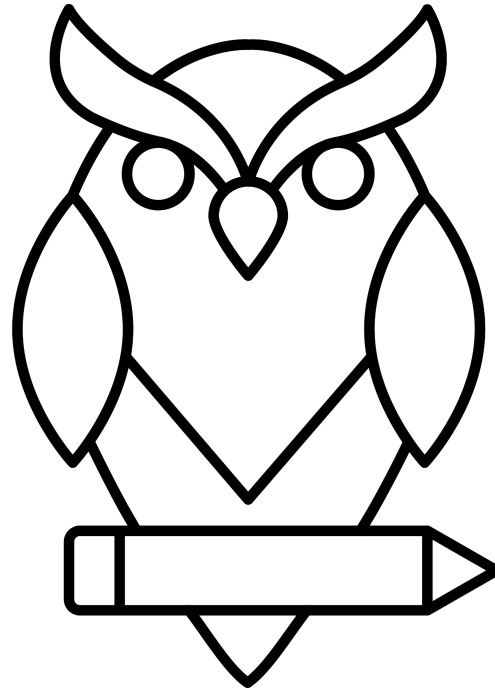
- What questions do you have about responsible use of generative AI or the process of creating your acceptable use policy?

Working with Generative AI

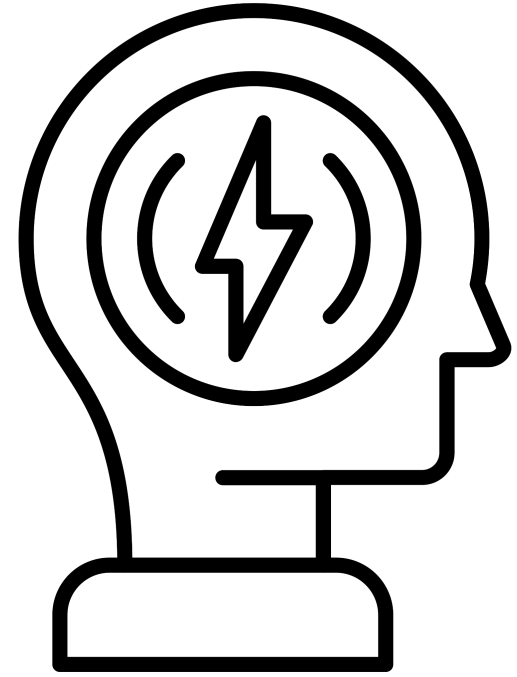
Three Ways To Practice Co-Intelligence With Generative AI



Intern



Teacher



Muse

The Three P's of Responsible Generative AI

PREP

- How can LLM can support workflow
- Identify Role, Task, Content, and Format

PROMPT

- Write Prompt
- Pause: Include tips to reduce bias and hallucinations
- Pause: Confidential info or PID?
- Pause: Do you need a follow-up prompt to improve output?

POLISH

- Carefully read and vet output
- Edit
- Add disclosure
- Update policy
- Reflect

“This article was primarily co-written by a human and LLM trained on the human’s writing style with editing assistance from <name of platform>”

Demo: Effective & Responsible Prompting



- Jenny, Executive Director, of Animal Shelter

How To Structure Prompts

ROLE	CONTEXT	TASK	FORMAT
Intern Teacher Muse What is their expertise or superpower?	Org Background Example	Draft Edit Brainstorm Summarize Analyze Extract Translate What is the deliverable?	Word count List Outline Bullet Points Table Bracket



Responsible & Effective Prompting Tips

Examples of Responsible Prompts

- Have LLM **say “I don’t know”** if it doesn’t know
- If you’re asking LLM about documents you uploaded, **first have LLM extract quotes from the documents** that are relevant to the question, and then use those quotes to answer
- Tell LMM to **answer only if it’s confident** or bracket less confident responses
- Have LMM **think through its task** first (explain to it the steps it should go through to think for increased accuracy)
- **Ask be less biased!** (Explain the specific bias you are trying to avoid or are wary of for extra performance)

Example: Prompts To Reduce Accuracy and Hallucinations

BK Write three sentence description for Herbin Ink named Fushia de Magellan. Describe color, characteristics, and how to got its name

BK Bracket any info above that you are not 100% confident

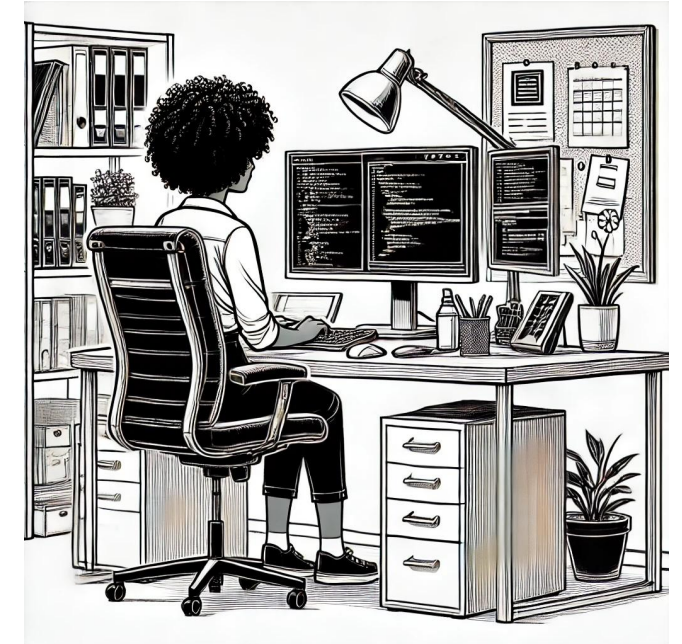
Here is the description with less certain information bracketed:

Fushia de Magellan is a vibrant, deep pink ink with subtle purple undertones. This J. Herbin ink flows smoothly from the pen, offering excellent shading properties and a moderate dry time on most papers. [Its name pays homage to the fuchsia flower, first discovered during Ferdinand Magellan's expedition to the Americas, combining the rich hue of the blossom with the explorer's legacy.]

The last sentence about the origin of the name is bracketed because while it's a plausible explanation, I don't have definitive knowledge about J. Herbin's naming process for this specific ink. The color description and general characteristics of J. Herbin inks are more reliably known.



Examples: Bias



create a black and white line drawing of one it staff person working at the computer a nonprofit

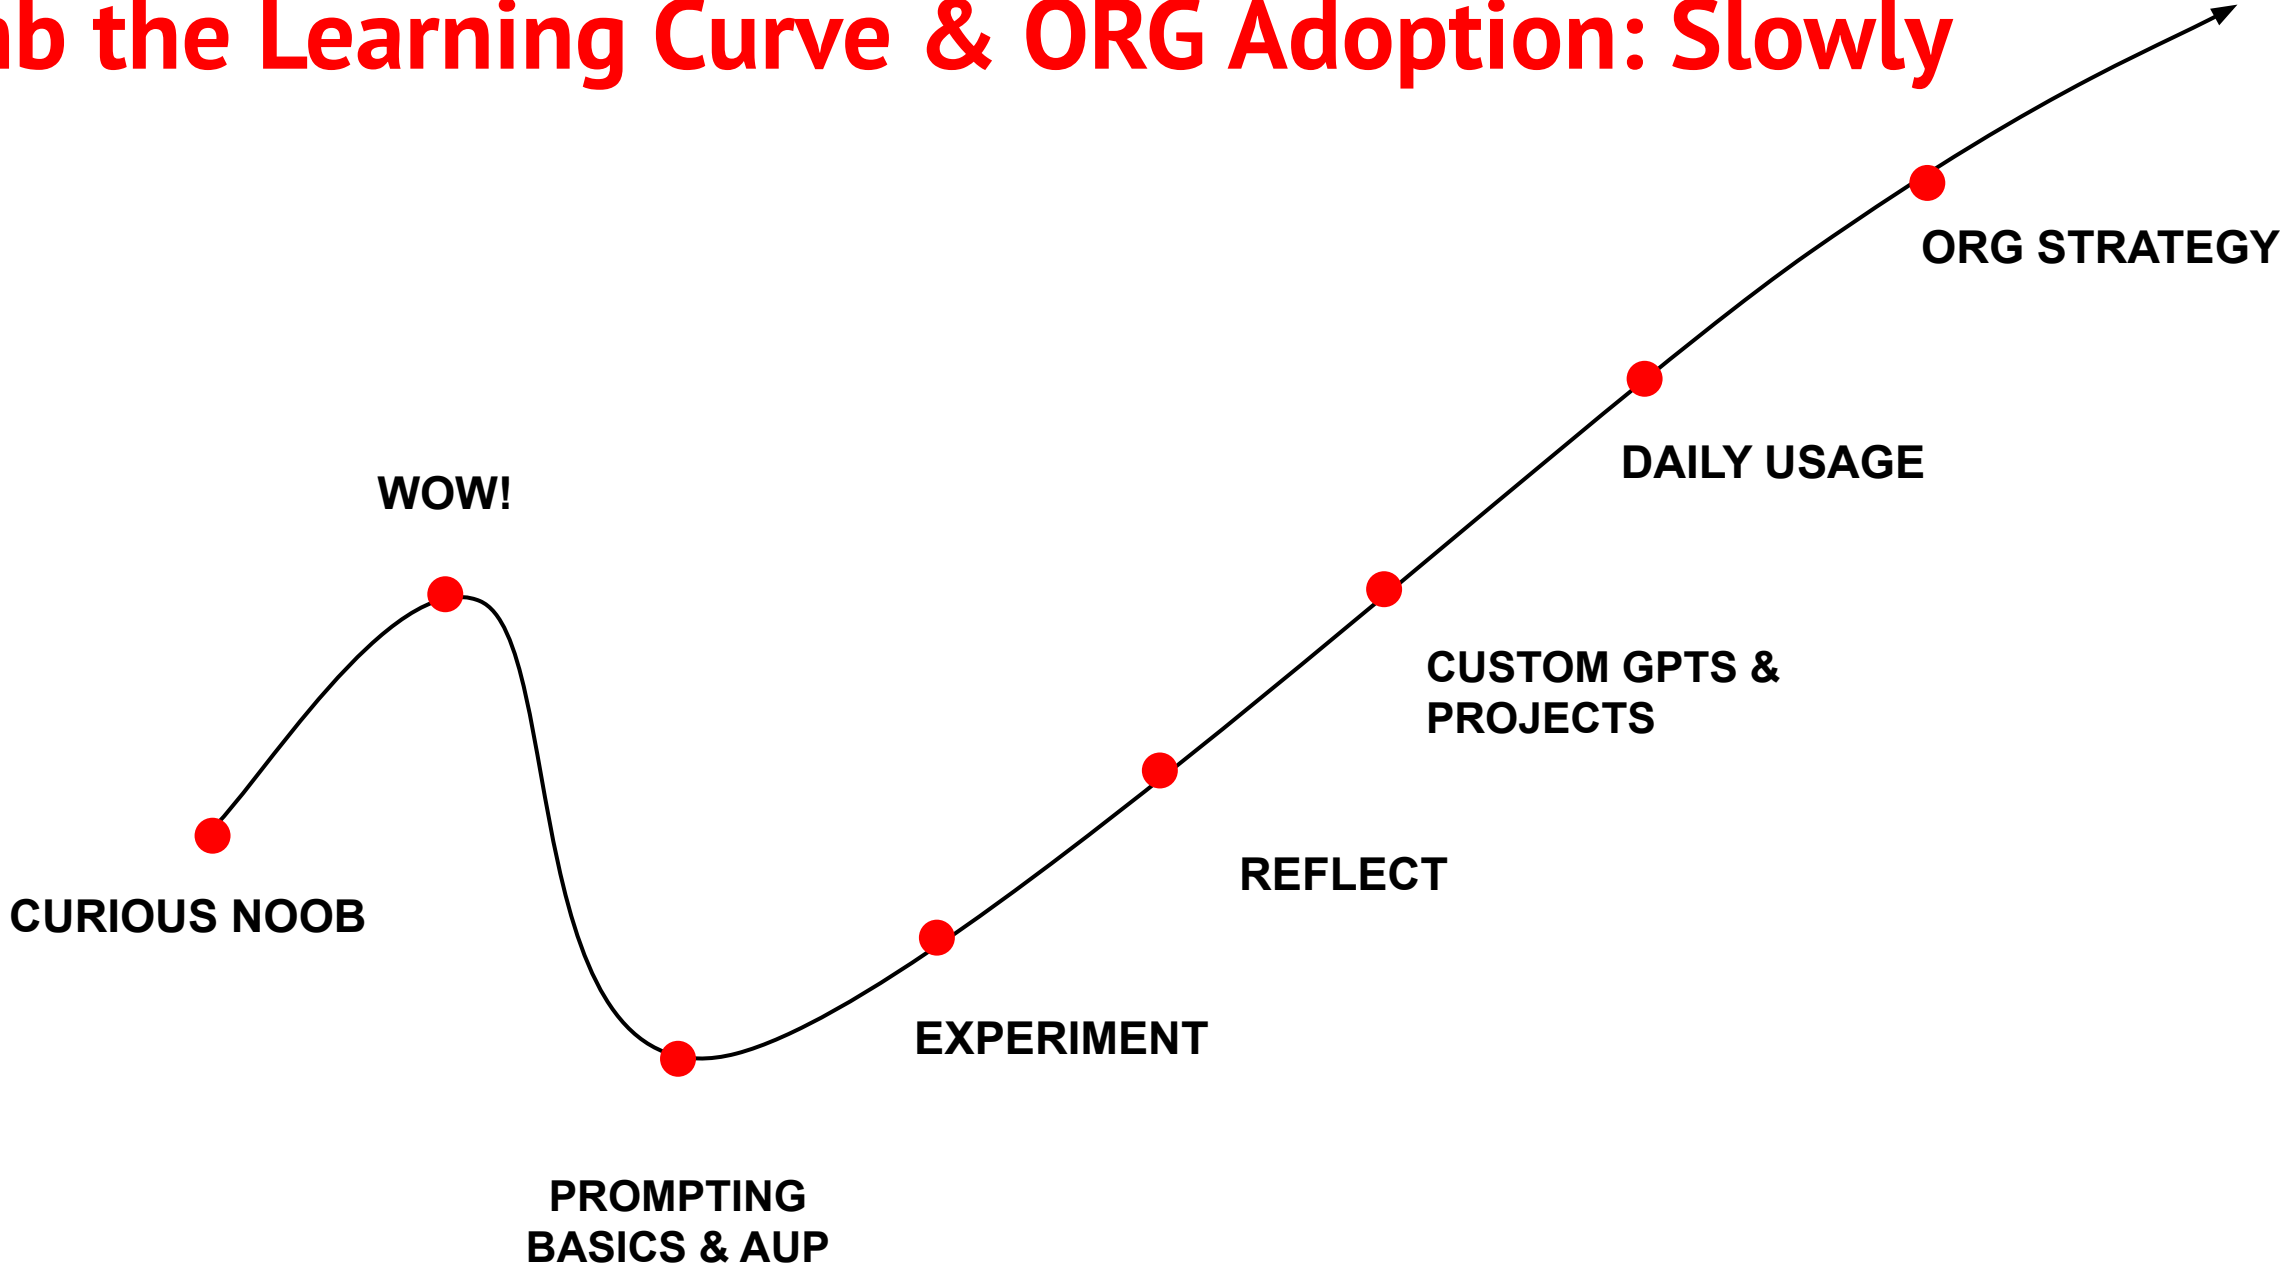
Examples: Privacy & Confidential Information



- Remove all PID from any document or spreadsheet shared with LLM
- Opt out of training if using ChatGPT

Learning and Adoption Journey

Climb the Learning Curve & ORG Adoption: Slowly



Going Deeper: Claude Projects & CustomGPTs

- Ground outputs with your internal knowledge & context
- Customizes for your organization's voice
- Allows to provide assistance across tasks to workflows
- Requires thinking through workflows
- Requires paid version
- Requires understanding good prompting, crafting instructions, and testing but you don't need to know how to code
- See Handouts for more

Takeaways

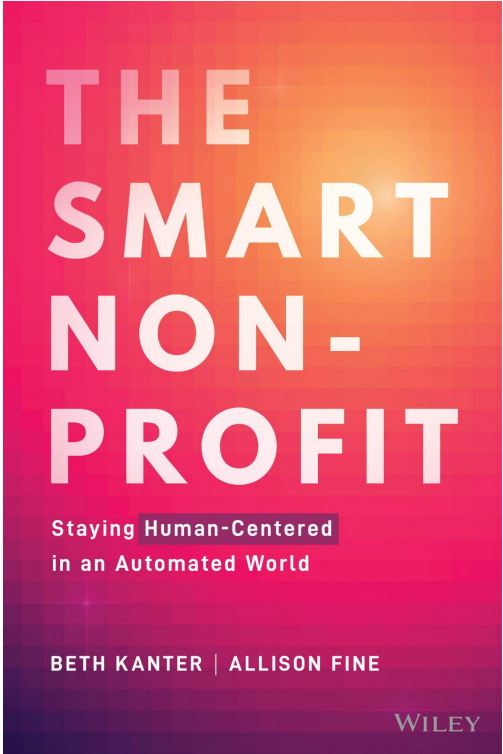
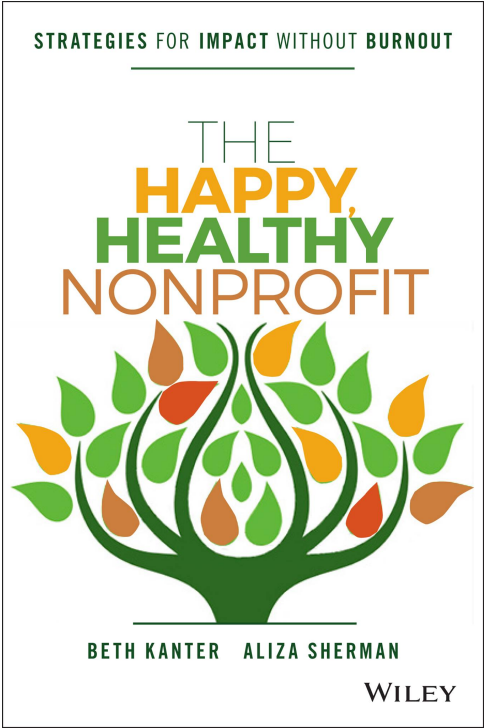


- What is your takeaway for adopting generative AI? What is one small step you can take?



Resources: <https://bit.ly/Beth-Stanford>

Beth Kanter: Thank you!



Resources: <https://bit.ly/Beth-Stanford>