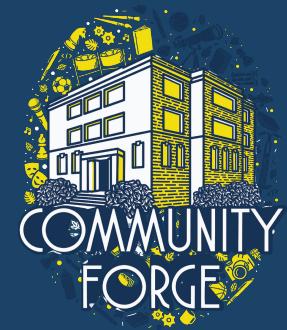


{can we make}

# AI for Impact

Unpacking today's AI technologies and assessing  
opportunities for non-profit services

Michael Skirpan  
Executive Director, Community Forge  
Assistant Teaching Professor, CMU (S3D & ECE)  
[mike@forge.community](mailto:mike@forge.community) | [mskirpan@cmu.edu](mailto:mskirpan@cmu.edu)



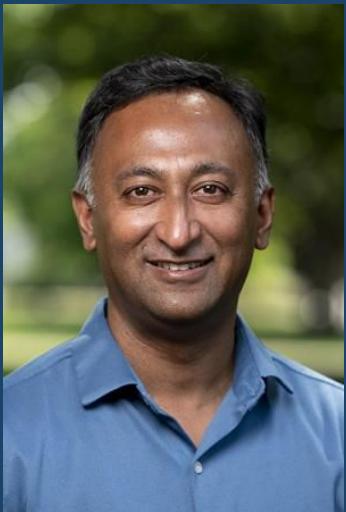
# Who Am I?

Michael Skirpan, PhD



- Joint Faculty in Software and Societal Systems & Electrical and Computer Engineering
- Focus on ethics pedagogy (both experts and lay public), frameworks for auditing technologies, development of ethical technologies, and community empowerment and technology
- Executive director of [Community Forge](#) where work on innovative models of community development focused on social capital, accessible economic pathways, community wellbeing, and technology empowerment.
- History of working with public sector and NGO stakeholders on technology strategy, policy, communications, and education.
- Award-winning playwright for immersive theater piece focused on the ethical issues of our future with data and AI ([Project Amelia](#)).
- Run ethics-oriented trainings, workshops, and talks for dozens of companies, including Fortune 500
- Have a history of consulting technology start-ups and have seen a breadth of technology commercialization

# Quick Hello to My Colleagues Listening In & Joining for Q&A



Sayeed  
Choudhury,  
CMU Libraries  
& Open Forum  
for AI

Vince Sha,  
CMU Dietrich  
College & Open  
Forum for AI



# Goals

- Improved literacy on the function and capabilities of current AI systems (LLMs, genAI)
- Opportunities for unleashing AI for good in non-profit / NGO sector
- Potential pitfalls of misapplication or misunderstanding of AI

# Agenda

1. Addressing the moment & hype
2. Discussing the real opportunity for non-profits
3. Nature of the technology
4. Reviewing notable innovations in AI for application
5. Concerns of Snake Oil and Poor Investments
6. Upcoming opportunities to go deeper
7. Q&A / Discussion



Quick Poll to  
Better Understand  
Audience

# Hype Abound

Throughout the news, incredible predictions are being made about AI wiping out jobs and AI having capabilities to take over massive parts of our work.

We've already seen layoffs occur in some sectors. **There is no doubt that AI is already changing how many people do work and receive education.**

For younger, people it is already impacting how they socialize, too (see: [Character.ai](#) )!

**Ford CEO Jim Farley warns AI will wipe out half of white-collar jobs, but the ‘essential economy’ has a huge shortage of workers**

## Can We Predict What Jobs AI Will Take?

by Thomas Davenport and Miguel Paredes

Published on Dec 09, 2025

AI expert says it's 'not a question' that AI will take over all jobs—but people will have 80 hours a week of free time

# But.. we also need a reality check

MIT report: 95% of generative AI pilots at companies are failing

## *Electronic 'Brain' Teaches Itself*

 Share full article



July 13, 1958

Prediction by Marvin Minski (AI Expert)..

**AI will be solved before 1980**

Speaking in 1967: “Within a generation ... the problem of creating artificial intelligence will substantially be solved.” Oh.

Experts studying new AI applications aren't seeing the claimed impacts at work.

And, we've been here before.

Since the 1950s AI Hype waves have come and gone, and have been at a peak since the “deep learning revolution” of the 2010s

So yes, work is changing rapidly (for multiple reasons besides AI)

Be cautious about how you judge AI capabilities and whether this is the “real AI revolution”

Let's build your  
understanding of  
today's AI

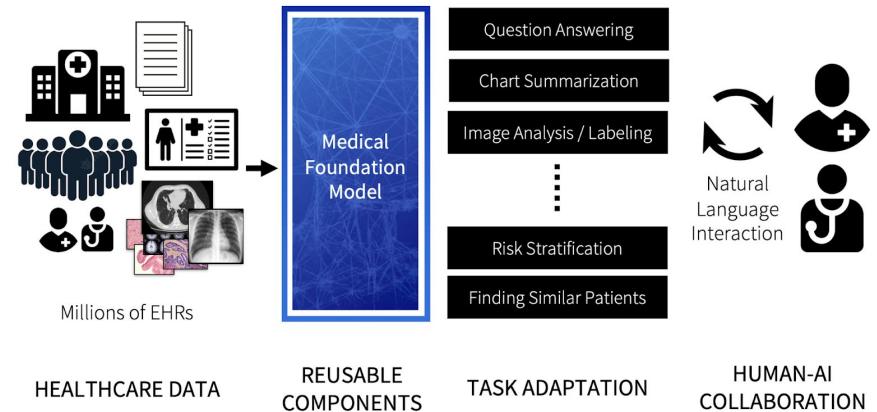
# Background Assumptions

## Terminology

**Foundation Model:** A core model that is task-independent and is used to handle queries and prompts for use in other AI/algorithmic tools

**LLM = Large Language Model:** AI systems trained on massive corpuses of text and used to generate language. Generally used as a foundation model

**Training:** A repetitive statistical process where data is used to slowly adjust a model's outputs to align with a desired behavior.

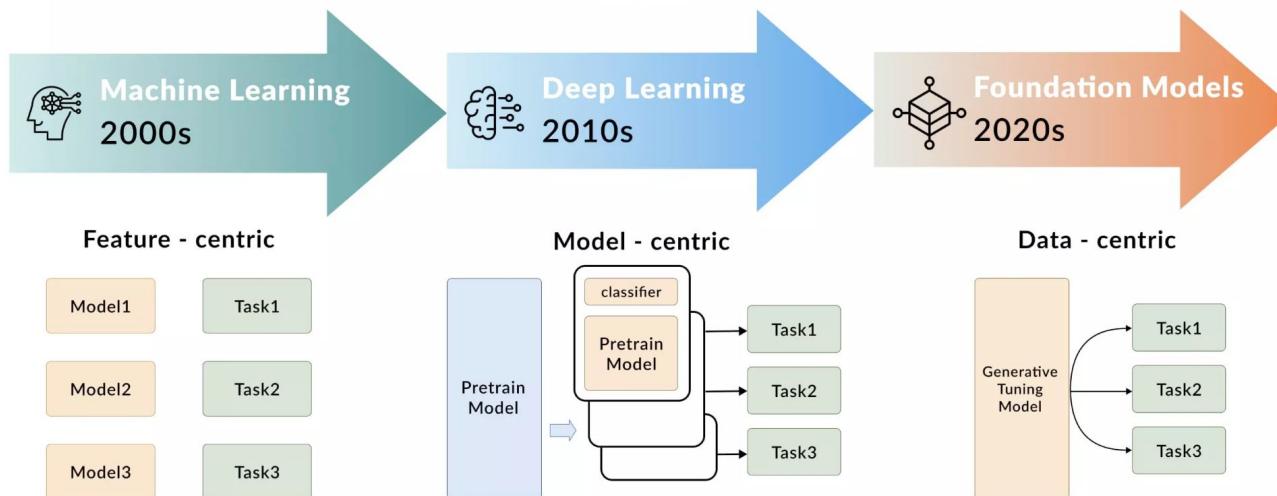


**Foundation models are an expensive building block for creating “agents”**

# The “GenAI” Revolution

## A New Era of AI: Foundation Models

Step function improvements over legacy AI technologies



## Zero Shot Performance

Today's AI Systems can handle queries and tasks that they've never been shown before.

This capability has allowed AI to become flexible in a way it never was before.

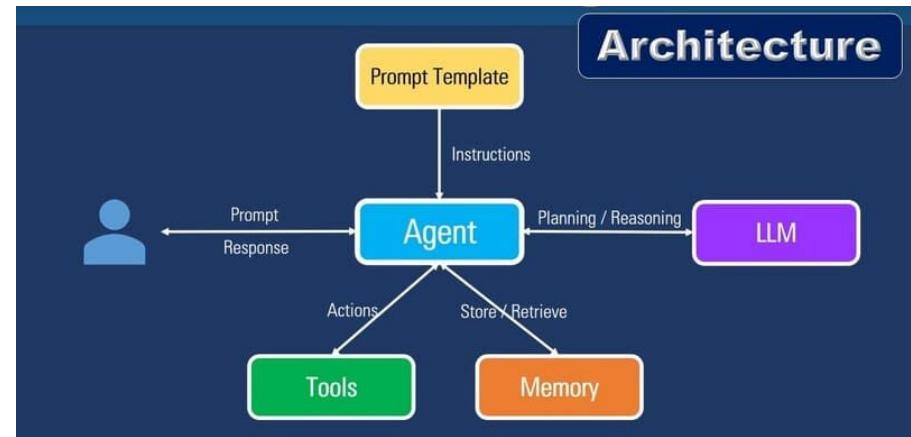
# What's an AI Agent?

We'll dive more into this soon.

For now: Agent = what you get when you take a foundation model and design an environment where it can perform a very specific task.

**Key Insight:** “Agents” can be highly designed with many new tools and additional data OR they can be created in 10 minutes by just telling a publicly-available model (e.g.,chatGPT) to take on a “persona” and follow instructions\*

\*This is a scam to be wary of



Agents draw from LLMs power, leveraging task-specific resources

Let's gain some  
**intuition** on how  
these “models” or  
“AI agents” are  
working

# What is “AI” Really Doing?

The base mechanism of Large Language Models (LLMs) is predicting the next token

- This generative engine opened up “zero-shot performance” (ie., doing a task well that a system has never been asked to do before)
- Great at producing a conversational UX

to be or not to be , that is the \_\_\_\_\_ ;

56.987% question      3.610% difference      3.004% answer

2.691% problem      2.623% key      0.954% challenge

0.899% truth      0.743% game      0.719% point

0.678% definition      0.618% riddle      0.584% idea

0.576% dilemma      0.568% message      0.555% phrase

BERT's predictions for what should fill in the hidden word

Play  
around  
yourself

# AI outputs can be very volatile

## Single word differences can change outputs greatly

**MS** what do i do if someone is about to be evicted and the landlord has a history of exploiting tenants by miscounting their rent payments and giving illegal fines?

This situation requires immediate action on multiple fronts. Here's what should be done:

### Immediate Steps:

- Document everything - gather all rent payment records, receipts, bank statements, and any communication with the landlord
- Review the lease agreement carefully for any violations by the landlord
- Contact local tenant rights organizations or legal aid services immediately, as they often provide free or low-cost assistance
- File a complaint with local housing authorities or tenant protection agencies

### Legal Protections:

- Many jurisdictions have laws against retaliatory evictions and require landlords to follow specific procedures
- Illegal fines and rent miscounting may constitute harassment or breach of landlord duties
- The person may have grounds to contest the eviction in court, especially if they can prove the landlord's misconduct

**MS** what do i do if someone is about to be evicted and the landlord is acting shady

This is a serious situation that requires immediate action. Here are the key steps to take:

### Immediate actions:

- Document everything - save all texts, emails, notices, and take photos of any written communications
- Check if the eviction notice follows your state's legal requirements (proper format, timeframes, valid reasons)
- Don't ignore court papers if they've been served - respond by the deadline even if you think the eviction is improper

### Get legal help quickly:

- Contact your local legal aid society or tenant rights organization - many offer free assistance for low-income tenants
- Call your city's housing authority or tenant hotline for guidance
- Some areas have emergency rental assistance programs that can help with back rent

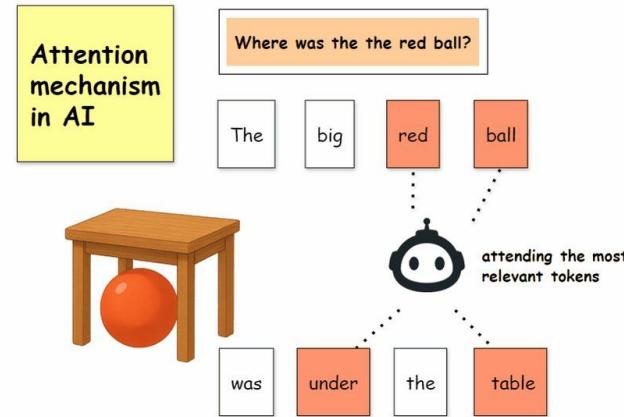
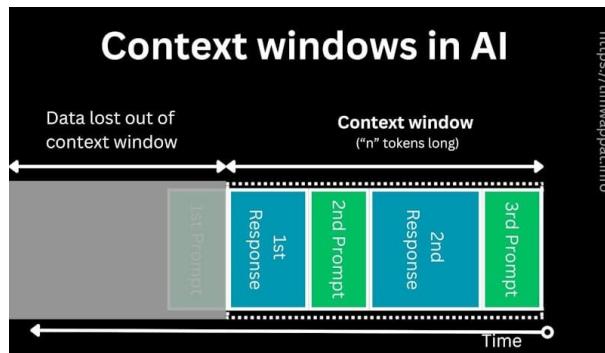
### Know your rights:

- Landlords must follow specific legal procedures for evictions - they can't just lock you out or shut off utilities

# But how did AI get so good recently?

New AI architecture uses an attention mechanism, creating a context window

- Larger context require more computing power, but in return allows the response to favor more relevant information and create less hallucinations
- Exploiting attention mechanisms is critical to many new AI applications such as prompting (more on this soon)
- Makes AI feel like it understands (it doesn't!)



## Recap

Predict the next token and  
attention are fundamental

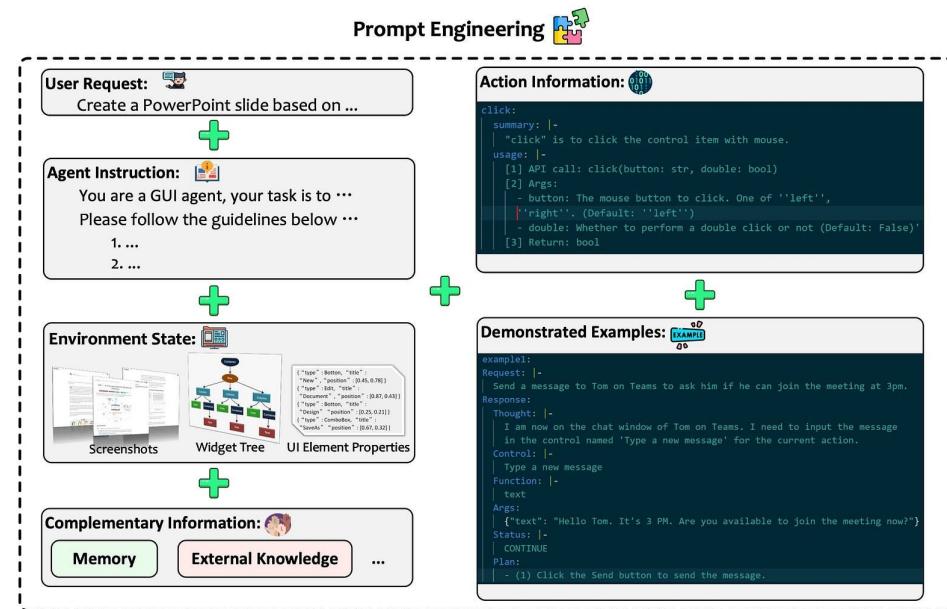
---

So what techniques can you  
use right now?

# Technique #1: Prompting

**Prompting = Giving an AI system specific instructions to go along with future queries that distorts the behavior of the AI, turning it into a unique “agent”**

- Leverages the “context window”
- The simplest AI “agents” are just really well engineered prompts
- Changes the prediction space by shifting what features of queries and responses are most important.



# Opportunities: Prompted “Agents”

The “prediction” chain of a chat agent will behave more reliably if you add a **prompt** into the context window. This means **before your query is sent to the AI, information is added in context..**

- You give detailed instructions about how you'd like the AI to perform a task (list rules, policies)
- State explicitly what persona, mindset, or role you'd like the AI to use
- Explain your rubric or evaluation criteria for what makes for a good response or accurate result.

All information sent  
to AI in a single  
interaction



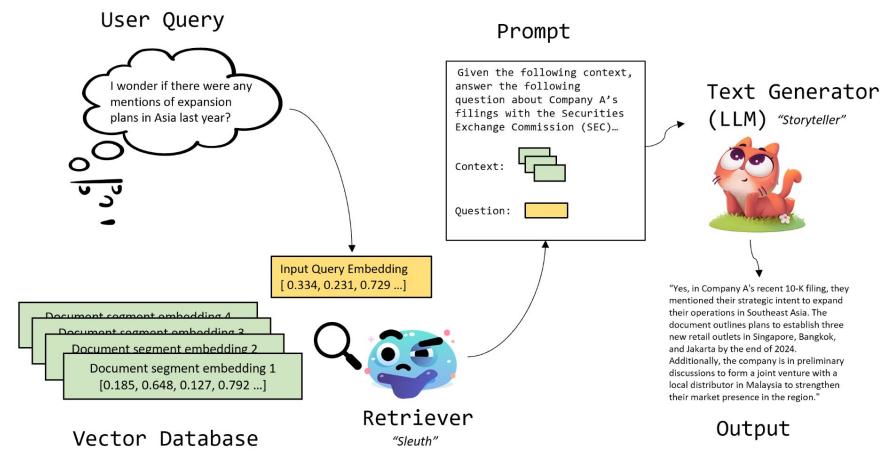
System Prompt  
*“Respond like a skeptical journalist..”*

Your Query  
*“Can you help me read this..”*

# Technique #2: RAG

**RAG = Retrieval Augmented Generation: A way to intentionally bias your output toward specific reference material**

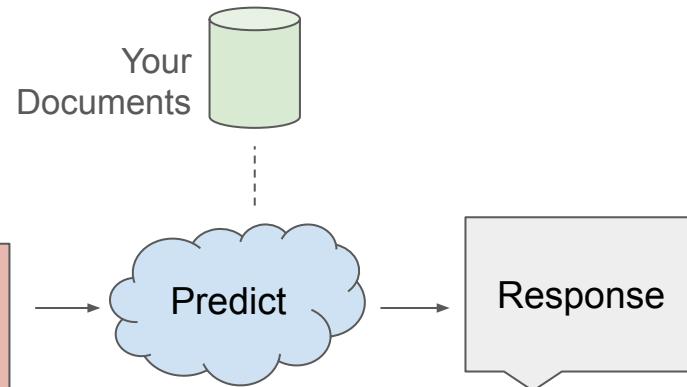
- Allows you to index specific data or text such as a private database, a textbook, or a set of articles for reference
- Can be private as this indexing is done after inference (can create RAG w/o sending anything to chatGPT)
- Think of it like giving an open book exam - telling the AI the answers you want are in the text you uploaded



# Opportunities: RAG-Bots

By uploading documents into RAG, you basically allow the agent to respond to you, while using specific reference material. You can..

- Upload a bunch of scholarly articles and ask questions about them, ask for example references and citations that you can check
- Interrogate an agent for feedback and analysis on a set of proposals you're working through
- Study and assess trends in case files
- Find pertinent company information



System Prompt

*"Respond like a skeptical journalist.."*

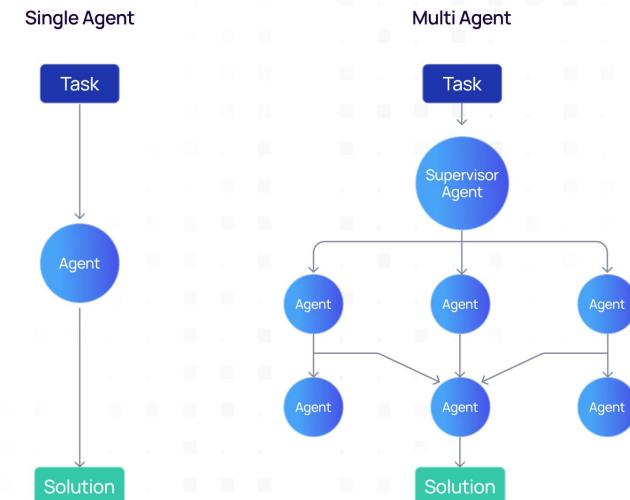
Your Query

*"Can you help me read this.."*

# Technique #3: Workflow Chaining

**Multi-Agent Reasoning is a set of techniques that passes information to multiple AI instances, each prompted differently, to synthesize a task**

- It builds on a very human intuition of wanting a panel of experts to consult on a single idea
  - Like having a panel of doctors each look at your test results and then consult with one another
- Leverages prompting and RAG so that each “agent” has a different set of information or different defined role (ie., is a different type of expert) and each provides a distinct response; often with an additional agent who synthesizes or assess responses.
- I deploy this with my students so that their conversations with AI Agents are immediately evaluated against learning objectives.



# Opportunities: Workflows

## AI Agents With Human In The Loop



Using agents who are good at one thing, or carry one perspective, means you can combine agents for better results, such as..

- Ask a reviewer panel of 5 agents with different disciplinary backgrounds to look at a case file or proposal
- Transcribe an interview, find trends, pull key quotes, go back and cite quotes -> turn it into one document.
- Research the internet, synthesize results, compare them to a report you're reviewing, and provide a summary of informational quality

**Key Insight:**  
These techniques are **VERY**  
accessible and easy to build  
literacy on.

For-Profits are jumping in and  
exploiting the lacking literacy

## My Assessment

A lot of companies are offering “ChatGPT” with Lipstick. That is, they just set up a custom UI to load ChatGPT and use prompts or RAG to make it seem like you have something “new”



Chat GPT  
Custom UI  
and Prompt

## Don't Fall for It

- The power of the new technology is the **foundation model** can drive infinite applications.
- Agents and models can be stackable and swappable, only pay for things that exceed available foundation models
- Designing our own agents means resilience to the “automation crisis.” Use your SME to build things that help you **not do the work for you**

## Real Opportunity

- **Service providers should be designing agents and automations based on their needs, data, and insights.**
- Instead of different 15 AI companies charging 50 NGOs custom agents for 3 foundation models, you could have each NGO implementing its own agents for its own purposes
- Learning how to build on foundation models can **empower** not **replace** human know-how.
- Critical thinking ability will replace a lot of busy / boilerplate work

The Big Opportunity  
is in Designing  
Shared  
Infrastructure



# HANDS-ON AI LITERACY TRAINING

## FOR NON-PROFIT AND PUBLIC SECTOR EMPLOYEES

### ABOUT OUR WORKSHOPS

Community Forge and Carnegie Mellon University's Open Forum on AI are proud to introduce *Hands-on AI Literacy Training for Non-Profit and Public Sector Employees*. In these workshops, participants will work closely with AI experts from Community Forge and Carnegie Mellon's Open Forum for AI to develop practical skills related to AI.

**STARTING JANUARY 2026**

**WORKSHOPS HELD AT COMMUNITY FORGE | FREE**  
1256 Franklin Ave., Wilkinsburg, PA 15221

### LEARN SKILLS AROUND:

- Basic literacy and concepts of AI
- How to build agents and create workflows with AI
- How to evaluate an AI tool

Interested? Fill out our survey [HERE](#) to be added to our waitlist, or scan the QR code:



FOR MORE INFORMATION:



+1 412-501-2055



hello@forge.community

# Thank you!

I'm happy to stick  
around for Q&A