

Future of Decentralization, AI, and Computing Summit

Opening Remarks

Dawn Song
Professor, UC Berkeley
Co-Director, Berkeley RDI


Agenda



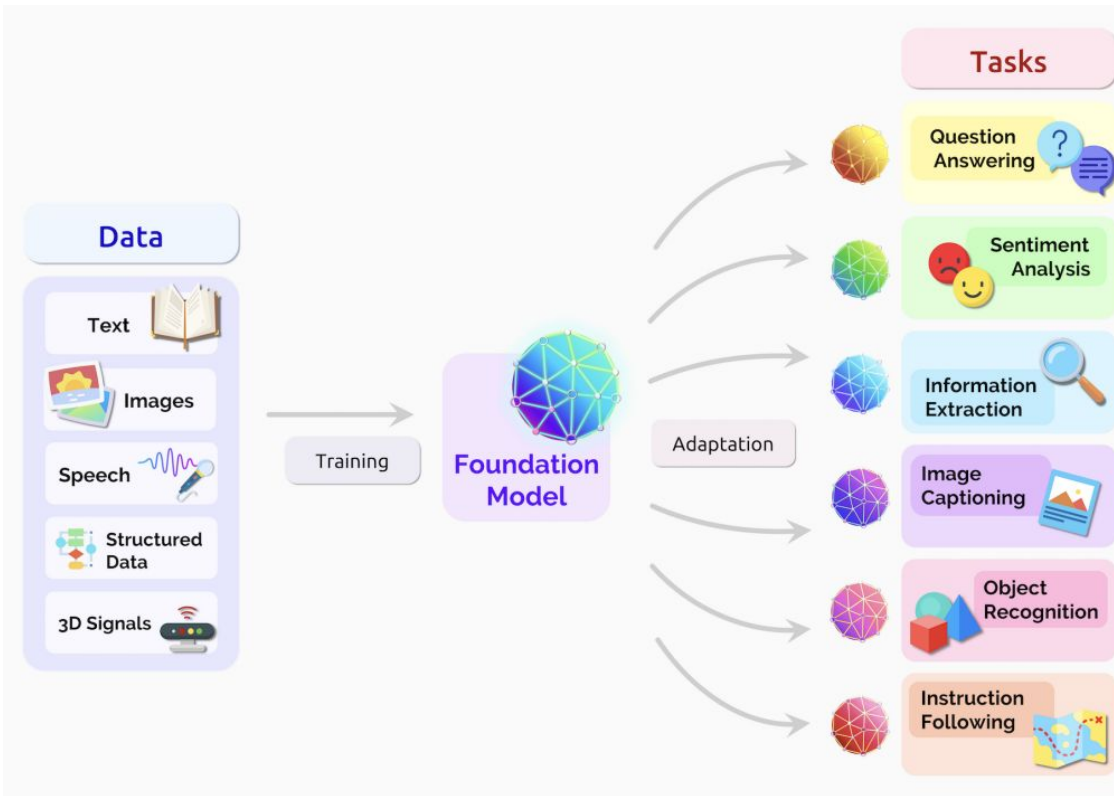
Exponential Growth in LLMs

Large Language Models (LLMs) & their associated bots like ChatGPT

● Amazon-owned
 ● Chinese
 ● Google
 ● Meta / Facebook
 ● Microsoft
 ● OpenAI
 ● Other



Powering Rich New Capabilities




<https://arxiv.org/pdf/2108.07258.pdf>

- | | | | |
|--|--|---|---|
| Q&A
Answer questions based on existing knowle... | Grammar correction
Corrects sentences into standard English. | Spreadsheet creator
Create spreadsheets of various kinds of dat... | JavaScript helper chatbot
Message-style bot that answers JavaScript ... |
| Summarize for a 2nd grader
Translates difficult text into simpler concep... | Natural language to OpenAI API
Create code to call to the OpenAI API usin... | ML/AI language model tutor
Bot that answers questions about language... | Science fiction book list maker
Create a list of items for a given topic. |
| Text to command
Translate text into programmatic commands. | English to other languages
Translates English text into French, Spanish... | Tweet classifier
Basic sentiment detection for a piece of text. | Airport code extractor
Extract airport codes from text. |
| Natural language to Stripe API
Create code to call the Stripe API using nat... | SQL translate
Translate natural language to SQL queries. | SQL request
Create simple SQL queries. | Extract contact information
Extract contact information from a block of ... |
| Parse unstructured data
Create tables from long form text | Classification
Classify items into categories via example. | JavaScript to Python
Convert simple JavaScript expressions into ... | Friend chat
Emulate a text message conversation. |
| Python to natural language
Explain a piece of Python code in human un... | Movie to Emoji
Convert movie titles into emoji. | Mood to color
Turn a text description into a color. | Write a Python docstring
An example of how to create a docstring for ... |
| Calculate Time Complexity
Find the time complexity of a function. | Translate programming languages
Translate from one programming language ... | Analogy maker
Create analogies. Modified from a communi... | JavaScript one line function
Turn a JavaScript function into a one liner. |
| Advanced tweet classifier
Advanced sentiment detection for a piece o... | Explain code
Explain a complicated piece of code. | Micro horror story creator
Creates two to three sentence short horror ... | Third-person converter
Converts first-person POV to the third-pers... |
| Keywords
Extract keywords from a block of text. | Factual answering
Guide the model towards factual answering ... | Notes to summary
Turn meeting notes into a summary. | VR fitness idea generator
Create ideas for fitness and virtual reality g... |
| Ad from product description
Turn a product description into ad copy. | Product name generator
Create product names from examples word... | ESRB rating
Categorize text based upon ESRB ratings. | Essay outline
Generate an outline for a research topic. |
| TL;DR summarization
Summarize text by adding a 'tl;dr:' to the en... | Python bug fixer
Find and fix bugs in source code. | Recipe creator (eat at your own risk)
Create a recipe from a list of ingredients. | Chat
Open ended conversation with an AI assist... |

Source: openai


Stable Diffusion: Fastest Repo to Reach 35K Stars on GitHub

Stable Diffusion Developer Adoption




Source: GitHub

Midjourney: The Largest Discord



ChatGPT: Fastest Platform to 100M Users



Many Risks & Open Challenges for Responsible AI

- Who controls AI?
 - centralized vs. decentralized control; open vs. closed source
- Trustworthiness
 - Robustness
 - Adversarial robustness
 - Out-of-distribution robustness
 - Test-time attacks vs. training-time attacks
 - Privacy
 - Fairness
 - Toxicity
 - Stereotype
 - Machine ethics
- AI Safety
 - Misuse/abuse of AI
 - Super intelligence

Importance of Mitigating Risk of Extinction from AI

Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war.



The New York Times

A.I. Poses 'Risk of Extinction,' Industry Leaders Warn

Leaders from OpenAI, Google DeepMind, Anthropic and other A.I. labs warn that future systems could be as deadly as pandemics and nuclear weapons.

Signatories:

AI Scientists Other Notable Figures

Geoffrey Hinton

Emeritus Professor of Computer Science, University of Toronto

Yoshua Bengio

Professor of Computer Science, U. Montreal / Mila

Demis Hassabis

CEO, Google DeepMind

Sam Altman

CEO, OpenAI

Dario Amodei

CEO, Anthropic

Dawn Song

Professor of Computer Science, UC Berkeley

Ted Lieu

Congressman, US House of Representatives

Bill Gates

Gates Ventures

Ya-Qin Zhang

Professor and Dean, AIR, Tsinghua University

Ilya Sutskever

Co-Founder and Chief Scientist, OpenAI

Shane Legg

Chief AGI Scientist and Co-Founder, Google DeepMind

Future of Responsible AI with Decentralization & Democratization

- Can we build a full decentralized, open-source stack for AI/ML:
 - Open-source decentralized AI/ML infrastructure for training and inference, with provenance, integrity, privacy guarantees
 - Open-source models and tooling
 - Personalized AI with privacy and trustworthiness
 - Decentralized, cooperative AI with incentives and social welfare
 - Democratic, decentralized process for AI governance & alignment

Open Challenges

- Is this technically feasible?
 - Is it possible to close the gap btw open source & closed source models?
 - Is it possible to scale decentralized training to large scale models?
 - Is it possible to build autonomous cooperative, decentralized agents?
- How to design proper incentive to maximize societal benefits?
- Can such an open source, decentralized system lead to more misuse/abuse? Is this at odds with AI safety guarantees?
- What are the different alternatives and possibilities that should be considered?

Spark the Initial Conversation & Cross Pollinate btw Diverse Domains

- AI & ML
- Security & Privacy
- Game theory & Economics
- Decentralization technology
- ...

Agenda

- **9:30 AM Session I: Open source LLMs**
- **10:30 AM Session II: Decentralized and Distributed Machine Learning Infrastructure**
- **11:40 AM Session III: Cryptography & ML; Personalization**
- **1:30 PM Session IV: Open-source LLMs tools & ecosystem**
- **2:15 PM Session V: Multi-agent Systems & Economics**
- **3:20 PM Session VI: Future Directions in AI + Decentralization**
- **4:30pm Lightning talks, Startup spotlights, Xcelerator Demo day, Reception**

Berkeley Center on Responsible Decentralized Intelligence (RDI)

Mission: Advance the science and technology of decentralization & decentralized intelligence, to make it universally accessible and empowering a responsible digital economy.

Campus-wide, multi-disciplinary: EECS/School of Engineering, Haas Business School, Law School



<https://rdi.berkeley.edu/>

Berkeley Engineering

BerkeleyHaas

Berkeley Law

Berkeley
Computing,
Data Science,
and Society



Diving Deeper

Berkeley
UNIVERSITY OF CALIFORNIA

Responsible GenAI and Decentralized Intelligence

CS294/194-196 Fall 2023

Instructors



Dawn Song Matei Zaharia

<https://rdi.berkeley.edu/responsible-genai/f23>

Agenda

- **9:30 AM Session I: Open source LLMs**
- **10:30 AM Session II: Decentralized and Distributed Machine Learning Infrastructure**
- **11:40 AM Session III: Cryptography & ML; Personalization**
- **1:30 PM Session IV: Open-source LLMs tools & ecosystem**
- **2:15 PM Session V: Multi-agent Systems & Economics**
- **3:20 PM Session VI: Future Directions in AI + Decentralization**
- **4:30pm Lightning talks, Startup spotlights, Xcelerator Demo day, Reception**