The Means of Prediction
How AI Really Works (And Who Benefits)

Maximilian Kasy

Department of Economics, University of Oxford

Fall 2025

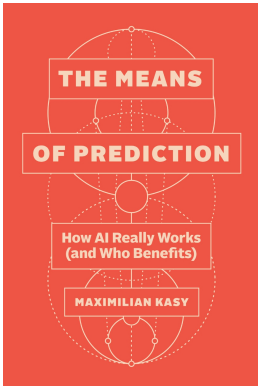

Are you scared of AI?

A popular dystopian story:

- AI will attain superhuman capabilities,
- will start to self-improve exponentially,
- and will threaten human existence in the name of self-preservation.

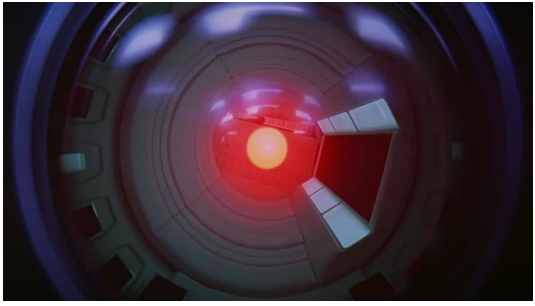

2001: A Space Odyssey

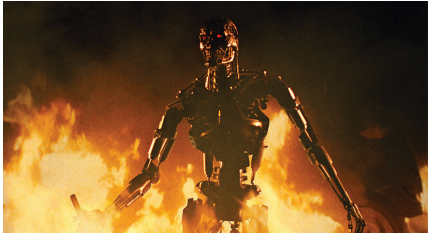

Terminator

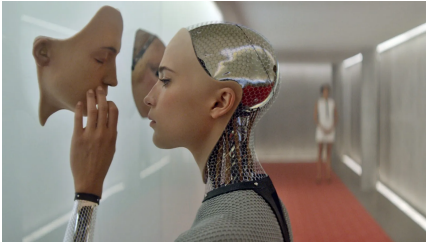

Ex Machina

- Such stories touch on our deepest fears:
  - Losing our livelihoods, autonomy, lives, and loved ones,
  - to inscrutable and inevitable forces.
- But they don't enable good decisions:
  - Make it seem like AI and its use are fate.
  - Obscure conflicts over who controls AI.
- Intentional obfuscation by tech players?

A more accurate story

1. AI is automated decision-making using *optimization*.
2. Key issue: Who gets to pick the *objectives* that AI optimizes?
   (Not: Did the AI fail to optimize?)
3. Power flows from control of AI *inputs*:
   data, compute, expertise, energy.
4. We need *democratic control* of AI objectives
   by those affected by AI decisions.

Some examples

Beyond the headline-grabbing large language models:

- Algorithmic management of gig-workers.
- Automatic screening of job candidates.
- Filtering and selection of social media feeds, search engine results.
- Ad targeting.
- Predictive policing and incarceration.
- Automated choice of bombing/assassination targets (e.g. Gaza).

How AI works

The political economy of AI

Regulating algorithms

How AI works

The book explains the foundations of machine learning and AI *without math*:

- AI is automated decision-making, maximizing some reward.
- Machine learning is AI using statistics.
  - Supervised learning: Prediction
    - Overfitting versus underfitting, tuning.
    - Deep learning, transformers.
  - Online learning: Choosing actions over time.
    - Exploration versus exploitation.
    - Planning.

Tuning of supervised learning algorithms

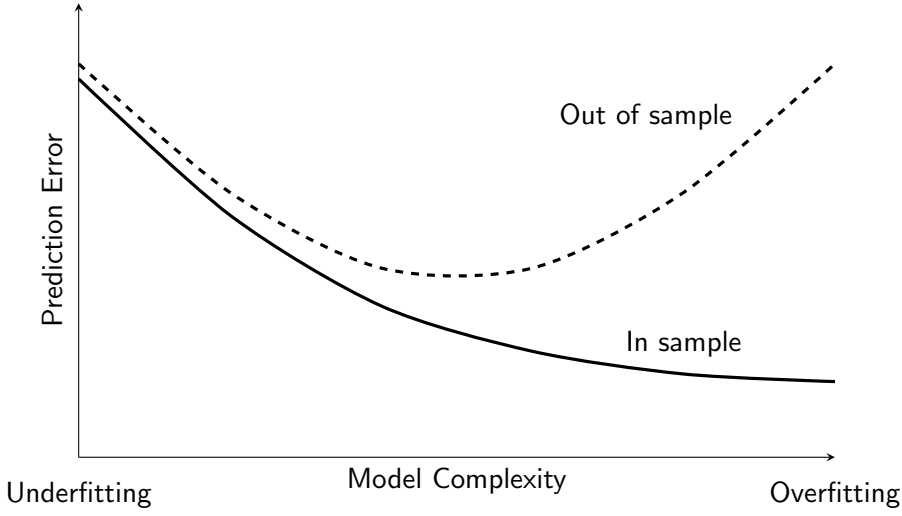

Scaling of large language models

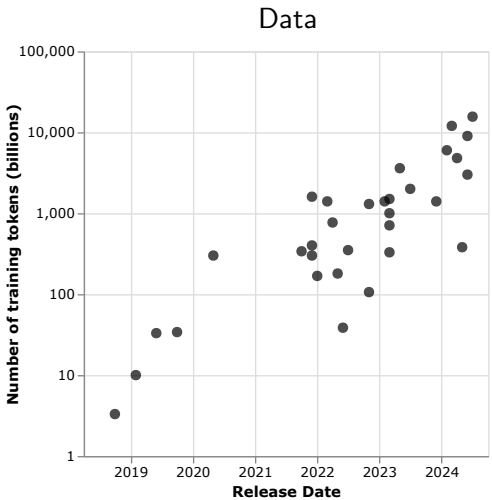

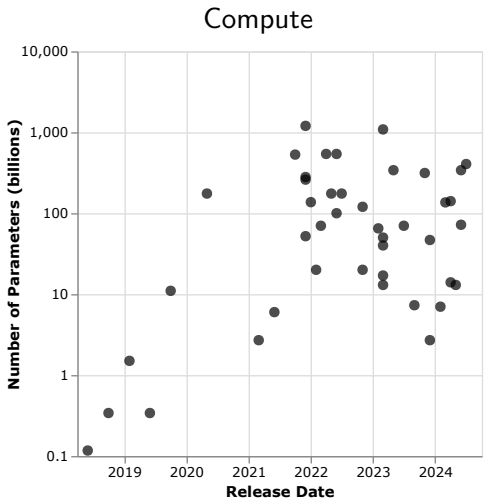

How AI works

The political economy of AI

Regulating algorithms

The means of prediction

- These foundations clarify what resources are needed for AI:
  - Data
  - Compute
  - Expertise
  - Energy
- Implications:
  - Potential for future improvements (domain-dependent).
  - Control of AI by controlling its inputs.
  - Contests over property rights, externalities.

Agents of change

Who can align AI objectives with social welfare?

- Interests, values, and strategic leverage.
- AI companies? Constrained by profit maximization.
- AI discourse should address others:
  - Workers (click-, gig-, tech-), consumers,
  - media and public opinion, state and law.
- Ultimate goal:
  - Democratic control of AI objectives
  - by those impacted by AI decisions.

Ideological obfuscation

- Ideology: Represents
  - Interests of a particular group as those of society at large.
  - Contingent choices as objective necessity.
  - Social relationships as technical ones.
- Popular AI stories that prevent change:
  1. *Man versus machine*: Obfuscates conflicts within society.
  2. *Intelligence explosion*: Not human choices but autonomous process.
  3. *Only experts understand AI*: Prevents democratic control of tech companies.
  4. *If we don't do it, China will*: Political inevitability.

How AI works

The political economy of AI

Regulating algorithms

Regulating algorithms

Ramifications of this perspective for various policy domains:

- Value alignment and the limits of AI
- Privacy and data ownership
- Workplace automation and the labor market
- Fairness and algorithmic discrimination
- Explainability of algorithms and algorithmic decisions

Thank you!

Book available for ordering here:

https://press.uchicago.edu/ucp/books/book/chicago/M/bo255887145.html