

Economics and Machine Learning: What can they teach each other?

Maximilian Kasy

Department of Economics, University of Oxford

September 2023

Economics and Machine Learning

- Economics shares with AI and machine learning (ML) the languages of
 - optimization, and
 - probability.
- But these fields also emphasize a number of distinct ideas.
- These distinct ideas matter, especially when we consider
 1. The use of AI for public good (as opposed to profit maximization).
 2. The ethics and social impact of AI.

Ideas from econ that matter for ML

1. Multiple agents
 - with unequal endowments,
 - conflicting interests, and
 - private information.
2. Welfare as utility
3. Aggregation via social welfare functions and welfare weights
4. Causal inference

Why these ideas from econ matter (1)

- ML tends to view everything as an **optimization problem**.
- Any potential issues are then understood as failures to optimize.
- Econ by contrast emphasizes
 1. **Conflicts of interest** and distributional impacts.
 2. Agency issues and asymmetric information.
 3. Externalities.

Examples from *AI ethics*:

1. Algorithmic bias and fairness.
 - Bias as a deviation from profit maximization?
 - Versus: The causal impact of automated decisions on the distribution of welfare.
2. Alignment and AI safety.
 - Value alignment as correctly specified reward function?
 - Versus: Conflict over the choice of objectives.

Why these ideas from econ matter (2)

- ML tends to consider **observable rewards** or losses.
- Normative economics emphasizes welfare as **utility**:
What people would choose.
- Utility is not directly observable.

Examples from *AI for public good*:

1. Labor market interventions.
 - Maximize employment probabilities?
Could be achieved via forced labor.
 - Versus: Maximize worker welfare by increasing their choice-sets.
2. Fertility and health in low income countries.
 - Minimize the number of births?
Could be achieved via forced sterilizations.
 - Versus: Maximizing women's autonomy in fertility and health decisions.

Papers that I will discuss

Cesa-Bianchi, N., Colomboni, R, and Kasy, M. (2023).

Adaptive maximization of social welfare

Kasy, M. (2023).

The political economy of AI:

Towards democratic control of the means of prediction

Kasy, M., and Abebe, R. (2021).

Fairness, equality, and power in algorithmic decision making

Kasy, M. (2023).

Algorithmic bias and racial inequality: A critical review

Introduction

Some “big ideas” from ML and Econ

Adaptive maximization of social welfare

The political economy of AI

Algorithmic bias and racial inequality

Conclusion

AI is automated decisionmaking

- AI systems maximize measurable **objectives**:

Russell and Norvig (2016), chapter 2:

For each possible percept sequence, a rational agent should select an action that is expected to maximize its performance measure, given the evidence provided by the percept sequence and whatever built-in knowledge the agent has.

- Leading approach: Machine learning (ML):
 1. Supervised learning.
 2. Targeted treatment assignment.
 3. Multi-armed bandits.
 4. Reinforcement learning.

Machine learning objectives

1. Supervised learning:

- Predict outcomes Y given features X .
- Prediction $g(X)$, prediction loss $l(g(X), Y)$.

2. Targeted treatment assignment:

- Assign a treatment W based on features X to maximize average outcomes Y among the treated.
- Assignment function $h(X)$, reward $h(X) \cdot Y$.

3. Multi-armed bandits:

- Maximize average outcomes over time. Cumulative reward $\sum_{t=1}^T Y_t$.
- Tradeoff between *exploration* and *exploitation*

4. Reinforcement learning:

- Expected cumulative reward $Q(X_t, W_t) = E[Y_t + Q(X_{t+1}, W_{t+1}) | X_t, W_t]$.
- Actions impact current reward and future state.

Adversarial bandits

- Canonical bandit problems:
 - Assign treatment sequentially.
 - Observe previous outcomes before the next assignment.
- Regret:
 - How much worse is an algorithm
 - than the best alternative in a given comparison set (e.g., fixed treatments).
- Two approaches for analyzing bandits:
 1. Stochastic: Potential outcomes are i.i.d. draws from some distribution.
 2. Adversarial: Potential outcomes are an arbitrary sequence.
- Adversarial regret guarantees:
 - Bound regret for arbitrary sequences.
 - We can do that because the stable comparison set substitutes for the stable data generating process.

Social welfare

Common presumption for many theories of justice:

- Normative statements about society are based on statements about individual welfare.
- Formally:
 - Individuals $i = 1, \dots, n$.
 - Individual i 's welfare v_i .
 - **Social welfare** is a function of individuals' welfare

$$F(v_1, \dots, v_n).$$

- This raises many questions:
 - Who is to be included among $i = 1, \dots, n$?
 - How to measure individual welfare v_i ?
 - How to aggregate to **social welfare**?

Individual welfare as utility

- Dominant in economics
- Formally:
 - Choice set C_i .
 - Utility function $u_i(x)$, for $x \in C_i$.
 - Realized welfare

$$v_i = \max_{x \in C_i} u_i(x).$$

- Double role of utility
 - Positive: Individuals choose utility-maximizing x .
 - Normative: Welfare is realized utility.

Optimal taxation

- Social welfare = weighted sum of individual utilities.
- Welfare weights:
 - Relative value of a marginal lump-sum \$ across individuals.
 - \approx Distributional preferences (rich vs. poor, healthy vs. sick,...).
- Envelope theorem:
 - Behavioral responses to marginal tax changes don't affect individual utilities.
 - They only impact public revenue (absent externalities).
 - \Rightarrow Impact on revenue is a sufficient statistic.
- Absent income effects:
 - Consumer surplus
 - = Equivalent variation
 - = integrated response function.

Causal inference

- Counterfactuals described by potential outcomes or structural functions:

$$Y^d = y(d, \epsilon).$$

- Automated decisionmaking requires to learn the causal effect *of algorithmic decisions*.
 - Conditional exogeneity is immediate.
 - Thus causal inference is trivial.
 - It is usually not even recognized as such in ML.
- But:
 - Discussions of fairness typically focus on inequality in treatment.
 - This is distinct from the impact on inequality in downstream welfare.
 - The distinction matters in the presence of pre-existing inequalities.

Introduction

Some “big ideas” from ML and Econ

Adaptive maximization of social welfare

The political economy of AI

Algorithmic bias and racial inequality

Conclusion

Adaptive maximization of social welfare

How should a policymaker act,

- who aims to maximize social welfare,

Weighted sum of utility.

⇒ Tradeoff redistribution vs. cost of behavioral responses.

- and needs to learn agent responses to policy choices?

Adaptively updated policy choices.

⇒ Tradeoff exploration vs. exploitation.

Cesa-Bianchi, N., Colomboni, R, and Kasy, M. (2023).

Adaptive maximization of social welfare

Adaptive maximization of social welfare

How should a policymaker act,

- who aims to maximize social welfare,
Weighted sum of utility.
⇒ Tradeoff redistribution vs. cost of behavioral responses.
- and needs to learn agent responses to policy choices?
Adaptively updated policy choices.
⇒ Tradeoff exploration vs. exploitation.

Cesa-Bianchi, N., Colomboni, R, and Kasy, M. (2023).

Adaptive maximization of social welfare

Adaptive maximization of social welfare

How should a policymaker act,

- who aims to maximize social welfare,
Weighted sum of utility.
⇒ Tradeoff redistribution vs. cost of behavioral responses.
- and needs to learn agent responses to policy choices?
Adaptively updated policy choices.
⇒ Tradeoff exploration vs. exploitation.

Cesa-Bianchi, N., Colomboni, R, and Kasy, M. (2023).

Adaptive maximization of social welfare

Adaptive maximization of social welfare

How should a policymaker act,

- who aims to maximize social welfare,
Weighted sum of utility.
⇒ Tradeoff redistribution vs. cost of behavioral responses.
- and needs to learn agent responses to policy choices?
Adaptively updated policy choices.
⇒ Tradeoff exploration vs. exploitation.

Cesa-Bianchi, N., Colomboni, R, and Kasy, M. (2023).

Adaptive maximization of social welfare

Setup: Tax on a binary choice

Each time period $i = 1, 2, \dots, T$:

- Policymaker (algorithm):
 - Chooses tax rate $x_i \in [0, 1]$.
- Agent i :
 - Willingness to pay: $v_i \in [0, 1]$.
 - Response function: $G_i(x) = \mathbf{1}(x \leq v_i)$.
 - Binary agent decision: $y_i = G_i(x_i)$.
- Observability:
 - After period i , we observe y_i .
 - We do *not* observe welfare $U_i(x_i)$.

Social welfare and cumulative regret

- Social welfare: Weighted sum of public revenue and private welfare:

$$\begin{aligned} U_i(\mathbf{x}) &= \underbrace{\mathbf{x} \cdot \mathbf{1}(\mathbf{x} \leq \mathbf{v}_i)}_{\text{Public revenue}} & + & \lambda \cdot \underbrace{\max(\mathbf{v}_i - \mathbf{x}, 0)}_{\text{Private welfare}} \\ &= \mathbf{x} \cdot \mathbf{G}_i(\mathbf{x}) & + & \lambda \cdot \int_{\mathbf{x}}^1 \mathbf{G}_i(x') dx'. \end{aligned}$$

- Cumulative welfare for a constant policy \mathbf{x} / actual policy choices \mathbf{x}_j :

$$\mathbb{U}_T(\mathbf{x}) = \sum_{i \leq T} U_i(\mathbf{x}), \quad \mathbb{U}_T = \sum_{i \leq T} U_i(\mathbf{x}_i).$$

- Adversarial regret:

$$\mathcal{R}_T(\{\mathbf{v}_i\}_{i=1}^T) = \sup_{\mathbf{x}} E \left[\mathbb{U}_T(\mathbf{x}) - \mathbb{U}_T \left| \{\mathbf{v}_i\}_{i=1}^T \right. \right].$$

The structure of observability

Choice x_i reveals $G_i(x_i)$. But

$$U_i(x) - U_i(x') = [x \cdot G_i(x) - x' \cdot G_i(x')] + \lambda \int_x^{x'} G_i(x'') dx''$$

depends on values of $G_i(x'')$ for $x'' \in [x, x']$!

Different from standard adaptive decision-making problems:

- Multi-armed bandits:
Observe welfare for the choice made.
- Online learning:
Observe welfare for all possible choices.

Lower and upper bounds on regret

Theorem

- *There exists a constant $C > 0$ such that for any algorithm: there exists a sequence $(\mathbf{v}_1, \dots, \mathbf{v}_T)$ for which*

$$\mathcal{R}_T(\{\mathbf{v}_i\}_{i=1}^T) \geq C \cdot T^{2/3}.$$



- *Consider the algorithm “Tempered Exp3 for social welfare.” There exists a constant C' such that for any sequence $(\mathbf{v}_1, \dots, \mathbf{v}_T)$,*

$$\mathcal{R}_T(\{\mathbf{v}_i\}_{i=1}^T) \leq C' \cdot \log(T)^{1/3} \cdot T^{2/3}.$$

Compare to the lower bound for stochastic / adversarial bandits: $C \cdot T^{1/2}$.

Monopoly pricing, and reserve price setting for auctions, are bandit problems!

Construction for the proof of the lower bound



Parameters: $\lambda = 0.95$, $a = 0.116$, $b = 0.003$.

Introduction

Some “big ideas” from ML and Econ

Adaptive maximization of social welfare

The political economy of AI

Algorithmic bias and racial inequality

Conclusion

The ethics and social impact of AI

- Concerns about the impact of AI:
 1. Fairness, discrimination, and inequality.
 2. Privacy, data property rights, and data governance.
 3. Value alignment and the impending robot apocalypse.
 4. Explainability and accountability.
 5. Automation and wage inequality.
- Corresponding efforts to regulate AI.
- How can we think systematically about these questions?

Kasy, M. (2023).

*The political economy of AI:
Towards democratic control of the means of prediction.*

The ethics and social impact of AI

- Concerns about the impact of AI:
 1. Fairness, discrimination, and inequality.
 2. Privacy, data property rights, and data governance.
 3. Value alignment and the impending robot apocalypse.
 4. Explainability and accountability.
 5. Automation and wage inequality.
- Corresponding efforts to regulate AI.
- How can we think systematically about these questions?

Kasy, M. (2023).

The political economy of AI:

Towards democratic control of the means of prediction.

Key arguments

1. AI systems maximize a single, measurable **objective**.
2. In society, different individuals have **different objectives**.
AI systems generate winners and losers.
3. Society-level assessments of AI
require trading off individual gains and losses.
4. AI requires democratic control
of algorithms, data, and computational infrastructure,
to align **algorithm objectives** and **social welfare**.

2. Privacy, data property rights, and data governance

Standard view:

(Dwork and Roth, 2014)

- Differential privacy.
 - It should make (almost) no observable difference whether your data are in a dataset.
 - No matter what other information is available to a decisionmaker.
- Machine learning performance is unaffected by differential privacy.
- Related:
Individual property rights over data.

Alternate view:

(Viljoen, 2021)

- Primary use of data in ML is to learn *relationships*, not individual data.
⇒ Informational externalities.
(Acemoglu et al., 2022)
 - Privacy / property rights cannot prevent harms from AI.
- ⇒ Only democratic governance can address harms, not individual property rights.

3. Value alignment and conflicts of interest

Standard view: (Russell, 2019):

- Value alignment is a gap between human and **machine objectives**.
- Possible solutions:
 1. More careful engineering of objective functions.
 2. Infer objectives from observed human behavior (“inverse reinforcement learning”).

Alternate view:

- Value alignment is a gap between the **objectives of those controlling the algorithm** and the **rest of society**.
- Additionally:
Not everything is observable, imposing fundamental limits on optimization.
- Possible solutions:
 1. Democratic control to align **algorithm objectives** with **society**.
 2. Refrain from deploying AI in some consequential settings.

Introduction

Some “big ideas” from ML and Econ

Adaptive maximization of social welfare

The political economy of AI

Algorithmic bias and racial inequality

Conclusion

1. Algorithmic bias and racial inequality

Standard view:

(Pessach and Shmueli, 2020)

- Fairness \approx treating people of the same “merit” independently of their group membership.
- If an algorithm is maximizing **firm profits** then its decisions are fair by assumption.
- No matter how unequal the resulting outcomes within and across groups.
- Only deviations from profit-maximization are “unfair.”

Alternate view:

(Kasy and Abebe, 2021; Kasy, 2023)

- **Welfare** / equality \approx (counterfactual / causal) consequences of an algorithm for the distribution of welfare of different people.
- Fairness vs. equality:
 1. Improved prediction \Rightarrow Treatments more aligned with “merit.”
Good for fairness, bad for equality.
 2. Affirmative action / redistribution:
Bad for fairness, good for equality.

“Algorithmic bias” as deviation from profit maximization

- Job candidates get wage w (known), their marginal contribution to profits would be M (unknown).
- Employer / algorithm makes hiring decisions D based on covariates X (known).

$$d(X) = P(D = 1|X).$$

- X can be used to predict M , $m(X) = E[M|X]$.
- A test for deviation from profit maximization: Suppose

$$m(x) > m(x'), \quad d(x) < 1, \quad \text{and} \quad d(x') > 0.$$

Then profits could be increased by hiring more candidates with features x and fewer candidates with features x' .

- Most fairness definitions are based on variants of this condition.

The causal impact of an algorithm on the distribution of welfare

- Outcomes are determined by the potential outcome equation

$$Y = W \cdot Y^1 + (1 - W) \cdot Y^0.$$

- The realized outcome distribution is given by

$$p_{Y,X}(y, x) = \left[p_{Y^0|X}(y, x) + w(x) \cdot \left(p_{Y^1|X}(y, x) - p_{Y^0|X}(y, x) \right) \right] \cdot p_X(x).$$

- What is the impact of $w(\cdot)$ on a statistic ν ?

$$\nu = \nu(p_{Y,X}).$$

Examples: Variance, quantiles, between group inequality, social welfare.

Introduction

Some “big ideas” from ML and Econ

Adaptive maximization of social welfare

The political economy of AI

Algorithmic bias and racial inequality

Conclusion

Conclusion

- Ideas from Econ that matter for ML:
 1. Multiple agents with conflicting interests and private information.
 2. Welfare as utility.
 3. Aggregation via social welfare functions and welfare weights.
- Especially relevant for:
AI for public good, Ethics and social impact of AI.
- Versus the big commercial applications of AI:
Maximizing ad clicks, monopoly price setting.
- Ideas from ML that matter for econ:
 1. Variance/bias tradeoffs, data-dependent tuning.
 2. Sequential decisionmaking and exploration/exploitation tradeoffs.
 3. High-dimensional, non-traditional data formats.

Thank you!