Machine learning for policy

Maximilian Kasy

Professor, Department of Economics, University of Oxford

November 10, 2021

What is machine learning?

Applications of active learning in policy

Pitfalls

Artificial intelligence as (automatic) decision-making

- The purpose of artificial intelligence (AI) is the construction of systems that autonomously make decisions.
- Such systems
  1. receive a sequence of inputs (percepts),
  2. process them, and
  3. make decisions interacting with their environment.
- The goal is to maximize a stream of rewards (or minimize a stream of losses).

Machine learning as one approach to AI

- There are different approaches to AI.
- Previous decades: Expert systems.
  Encode human knowledge in databases.
- Modern AI has had breakthroughs with an alternative approach:
  Learn from data using statistics.
  $\Rightarrow$ Machine Learning!
- This approach becomes ever more successful as
  1. more data and
  2. more computational power

  become available.

Different branches of machine learning

- Supervised learning:
  Predict outcomes from observed features.
- Unsupervised learning:
  Learn simplified representations of unstructured data.
- Active learning:
  Adaptive decision making, while learning which actions work better.
- Reinforcement learning:
  Current actions affect the evolution of the environment.

Supervised learning

- Objective: Minimize the error rate of predictions.
- Applications: Predict
  - Description of image from image itself.
  - Written text from recorded sound.
  - Translated sentence from original sentence.
  - Likelihood of repaying from loan applicant characteristics.
- Methods for supervised learning:
  - Deep learning (neural nets).
  - Lasso regressions.
  - Random forests.
  - ....

Active learning

- Objective: Achieve good average welfare over time.
- Repeated decision-making.
- Each decision has a dual purpose:
    1. Achieve good outcomes now (“exploitation”).
    2. Learn what works for future decisions (“exploration”).
- Good algorithms balance the two in just the right way.
- Most common version: “Multi-armed bandits.”
- Alternative: “Exploration sampling.”
  Learning quickly what policy is best.

What is machine learning?

Applications of active learning in policy

Pitfalls

Application I: Job search assistance for refugees in Jordan

- Jordan 2019, International Rescue Committee.
  - Participants: Syrian refugees and Jordanians.
  - Main locations: Amman and Irbid.
  - Sample size: 3770.
- **Context**: Jordan compact.
  Gave refugees the right to work in low-skilled formal jobs.
- 4 **Treatments**:
  1. Cash: 65 JOD (91.5 USD).
  2. Information: On (i) how to interview for a formal job, and (ii) labor law and worker rights.
  3. Nudge: A job-search planning session and SMS reminders.
  4. Control group.
- **Conditioning variables** for treatment assignment: 16 strata, based on
  1. nationality (Jordanian or Syrian),
  2. gender,
  3. education (completed high school or more), and
  4. work experience (having experience in wage employment).

Locations

Irbid

Amman

Assignment probabilities over time

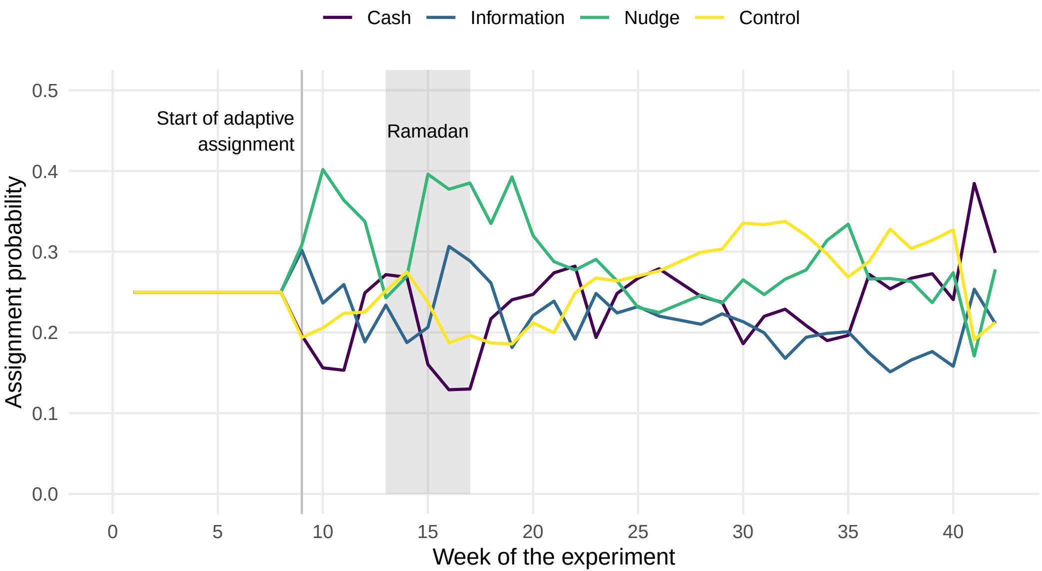

Application II: Agricultural extension service for farmers in India

- India, 2019.
  NGO Precision Agriculture for Development.
- **Context**: Enrolling rice farmers into customized advice service by mobile phone.

  > *[...] to build, scale, and improve mobile phone-based agricultural extension with the goal of increasing productivity and income of 100 million smallholder farmers and their families around the world.*

- Sample: 10,000 calls,
  divided into waves of 600.
- 6 **treatments**:
  - The call is pre-announced via SMS 24h before, 1h before, or not at all.
  - For each of these, the call time is either 10am or 6:30pm.
- **Outcome**: Did the respondent answer the enrollment questions?

Rice farming in India

Assignment shares over time

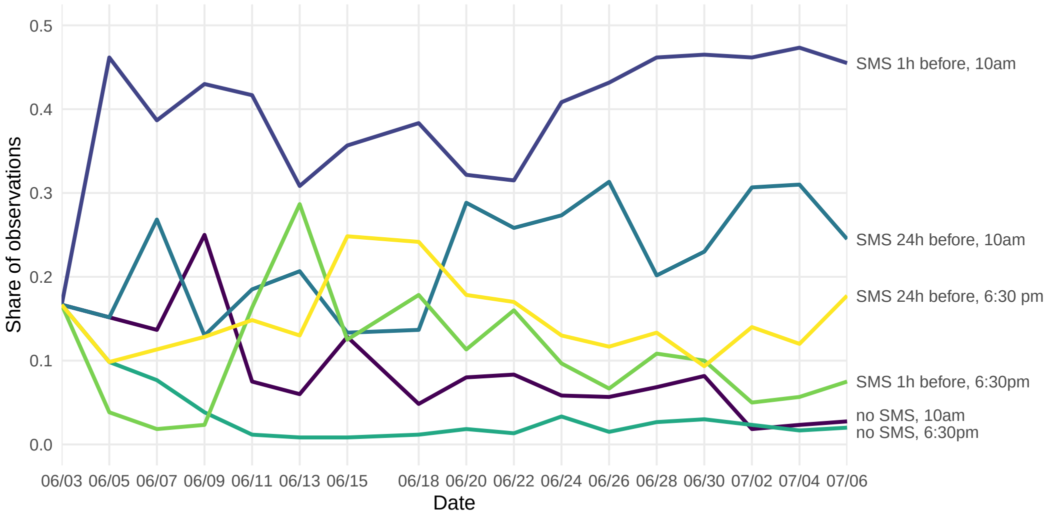

Application III: Matching refugees to host locations (simulations)

- Data for all refugees resettled by HIAS
  between January 2011 and December 2019.
- 8 demographic groups (types) based on
  - prime working age (25-54),
  - gender,
  - English-speaking.
- 17 affiliates (locations), with capacity constraints.
- Outcome Y_{jt}: Employed within 90 days of arrival.
- Simulations:
  - Calibrate success rates Θ_j for each type/affiliate combination.
  - Take actual capacity constraints.
  - Counterfactual matching using Thompson sampling.
  - Form posteriors using a hierarchical Bayesian model.

Simulated employment by year

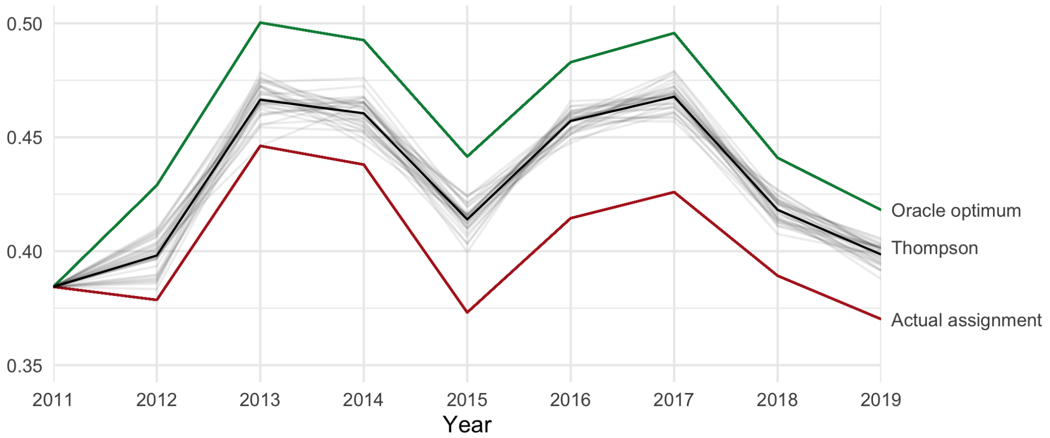

Simulated employment by type

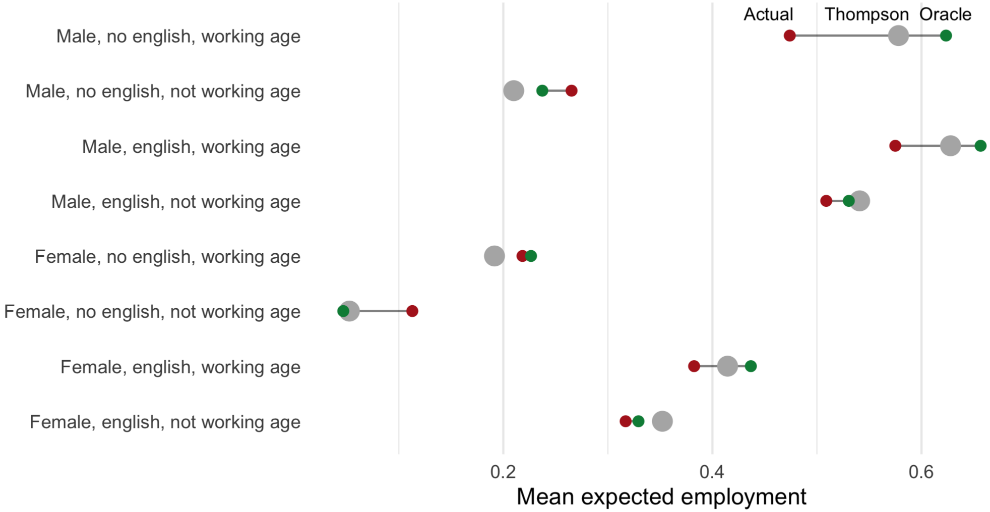

What is machine learning?

Applications of active learning in policy

Pitfalls

Pitfalls we encountered

1. Wrong outcome variable (Jordan experiment):
   - We targeted *formal* employment, *1 month* after the intervention.
     $\Rightarrow$ Little effect
   - It would have been better to target *all employment* at a *longer horizon*.
2. Wrong sample size / small effects:
   - If effects are too small, the algorithm can't adapt.
   - Benefits of adaptivity would have emerged later.
3. Wrong aggregation (refugee relocation):
   - Our simulations maximize *total employment*.
   - That led to a *decline* in employment for young non-English speakers.
   - The algorithm gave the best locations to those with the best prospects.

$\Rightarrow$ **CHOOSE THE OUTCOME THAT YOU ARE MAXIMIZING WISELY!**

Thank you!