YOU LIVER AND YOU LEARN: MACHINE LEARNING FOR LIVER TRANSPLANT ALLOCATION POLICIES

6.867 Class Project

Josh Wilde
Kayla Cummings
Samuel Gilmour

MOTIVATION AND AIMS

- There are currently almost 17,000 patients on the waitlist for a liver transplant in the United States.
- Transplant offers are made according to a greedy policy based on patient scores calculated by the Model for End-Stage Liver Disease (MELD), which measures disease severity.
- Recent work (termed OPOM) by Bertsimas et al. (2019) uses optimal classification trees to compute patient scores.
- This project had two objectives:
  1. Use modern machine learning models to improve the measurement of disease severity and build **patient** policies.
  2. Develop alternative policies for which patient rankings depend on the **organ** being offered.

INPUTS AND WORKFLOW

DATA

- The Scientific Registry of Transplant Recipients (SRTR) dataset contains records from Jan 2002 to Sep 2016:

$\mathcal{P}$: Patient status updates, MELD scores, and waitlist removal dates due to transplant or death.

$\mathcal{O}$: Organ properties for each transplant.

LIVER SIMULATED ALLOCATION MODEL (LSAM)

- LSAM is the official simulation model used to test policies for allocation. It has two main components:

ACCEPTANCE MODEL	SURVIVAL MODEL
$A : \mathcal{P} \times \mathcal{O} \rightarrow \{0, 1\}$	$S : \mathcal{P} \times \mathcal{O} \rightarrow \mathbb{R}_+$

- Offers are made using a greedy algorithm according to the scores computed for each patient on the waitlist $\mathcal{W}$ by the test policy.

DATA

1 RANKING POLICY USING $\mathcal{P}$

- Static ranking policy.

2 RANKING POLICY USING $\mathcal{P}$ AND $\mathcal{O}$

- Static ranking policy.
- Dynamic ranking policy (dependent on current state of waitlist).

SIMULATE TRANSPLANT OUTCOMES

1 PATIENT POLICY

- Existing methods (MELD and OPOM) construct patient scores by predicting short-term mortality (a supervised learning problem):

$\mathbf{x}_i \in \mathbb{R}^K$ vector of patient status updates (from $\mathcal{P}$)

$y_i \in \{0, 1\}$ binary indicating patient death within three months (computed from $\mathcal{P}$)

- Neural network architectures and additional feature engineering achieve the same AUC as OPOM.

2 PATIENT-ORGAN POLICIES

- Patient policies based on disease severity are empirically fair across demographics, but transplant outcomes depend on both the patient and the organ they receive.
- We develop two new policies based on Bertsimas et al. (2013) which use patients and organs, and bias towards fairness:

STATIC POLICY	DYNAMIC POLICY
$s : \mathcal{P} \times \mathcal{O} \rightarrow \mathbb{R}_+$	$s : \mathcal{P} \times \mathcal{O} \times \mathcal{W} \rightarrow \mathbb{R}_+$

BUILDING THE STATIC POLICY

A Solve an offline bipartite weighted matching problem over $\mathcal{P}$ and $\mathcal{O}$ with linear fairness constraints.

B Obtain optimal $(\mathbf{x}^*, \mathbf{p}^*)$ and an equivalent pure matching problem.

C Regress the modified costs on patient and organ features to produce a scoring policy.

A

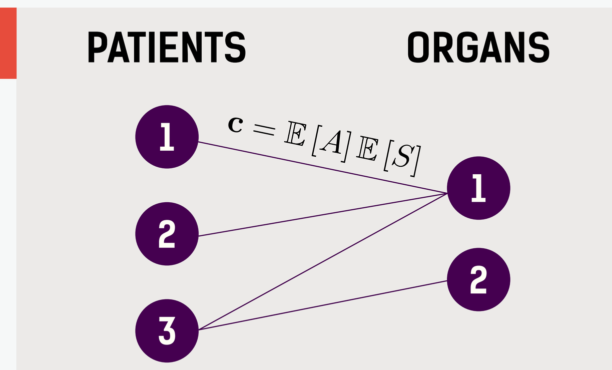

$$\begin{aligned} \max_{\mathbf{x}} \quad & \mathbf{c}^\intercal \mathbf{x} \\ \text{s.t.} \quad & x \in \{\text{match}\} \\ & A\mathbf{x} \leq \mathbf{b} \end{aligned}$$

B

$$\begin{aligned} \max_{\mathbf{x}} \quad & (\mathbf{c} - A^\intercal \mathbf{p}^*)^\intercal \mathbf{x} \\ \text{s.t.} \quad & x \in \{\text{match}\} \end{aligned}$$

C

$$\begin{aligned} \mathbf{x}_{ij} &= [\mathbf{p}_i \quad \mathbf{o}_j] \\ y_{ij} &= (\mathbf{c} - A^\intercal \mathbf{p}^*)_{ij} \end{aligned}$$

BUILDING THE DYNAMIC POLICY

- For the dynamic policy, we construct static policies on smaller batches of data and use each as a datapoint in a new regression task.
- The independent variables are now the mean of patient features in each batch, and dependent variables are the static policy weights.

RESULTS

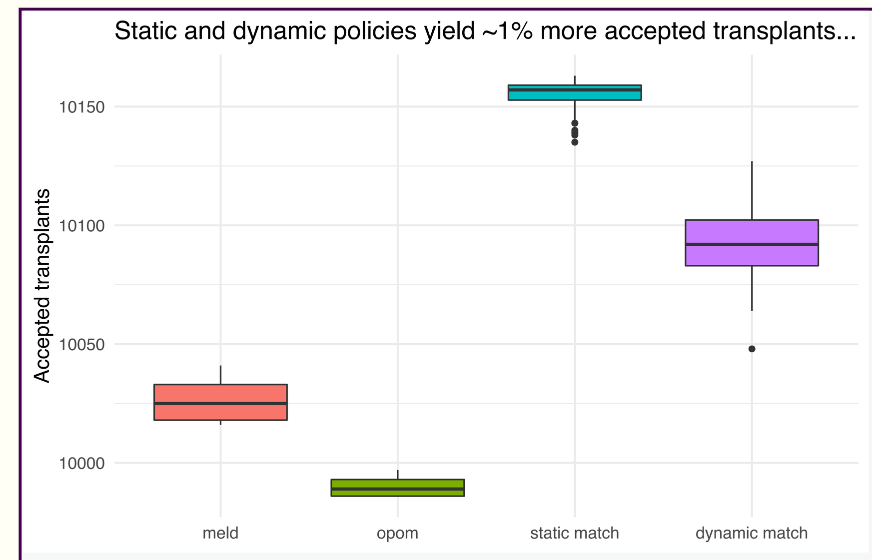

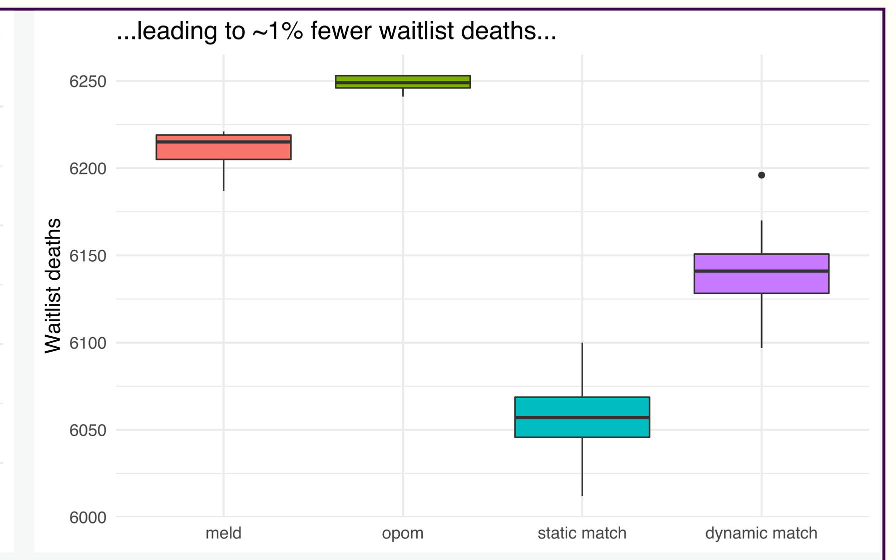

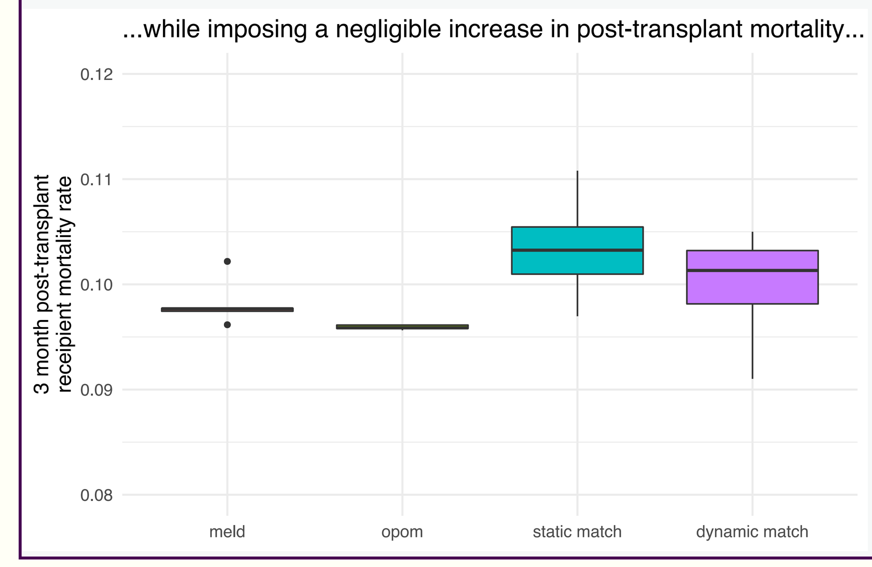

...and continuing to allocate organs fairly

	MELD	OPOM	Static Match	Dynamic Match
Race Non-White	0.30	0.30	0.30	0.30
Female	0.37	0.37	0.37	0.37
Elderly (65+)	0.12	0.12	0.12	0.12
Waitlisted 1+ years	0.12	0.12	0.10	0.13
Acute Hepatic Necrosis	0.05	0.04	0.05	0.05
Cirrhosis	0.08	0.08	0.08	0.08
Malignant Neoplasms	0.09	0.09	0.09	0.09
Liver disease	0.34	0.34	0.34	0.34
Obese (BMI > 30)	0.34	0.34	0.34	0.34
Blood Type O	0.47	0.47	0.47	0.47

CONCLUSIONS

- Price of interpretability in transplant candidate mortality prediction is near zero. OPOM comes close to approximating the Bayes Error when considering only patient features.
- Total accepted transplants increase and waitlist deaths decrease by ~1% in our simulations because we train policies based on projected outcomes for acceptance and post-transplant survival.
- Fairness constraints in our matching policies effectively maintain fair organ allocations within targeted demographics.
- Mortality outcomes do not improve by using other waitlisted patients' features to compute a given patient's ranking. We show that simpler static policies produce fair and efficient outcomes.

REFERENCES

Bertsimas, D., Kung, J., Trichakis, N., Wang, Y., Hirose, R., & Vagefi, P. A. (2019). Development and validation of an optimized prediction of mortality for candidates awaiting liver transplantation. American Journal of Transplantation, 19(4), 1109-1118.

Bertsimas, D., Farias, V. F., & Trichakis, N. (2013). Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. Operations Research, 61(1), 73-87.