Forecasting Patient Outcomes in Kidney Exchange

Naveen Durvasula[1], Aravind Srinivasan[2], John Dickerson[2]

[1]University of California, Berkeley [2]University of Maryland, College Park

Motivation

End stage renal disease (ESRD) affects over 750,000 annually in the United States alone with a societal burden comparable to diabetes. Growing demand for donor kidneys is met through the deceased-donor waiting list, direct donation, and **kidney paired donation (KPD)** programs. These programs allow patients with willing, but biologically incompatible donors to obtain a living-donor transplant by way of an organized market. The relative compatibility of exchange participants is encoded by a *compatibility graph*

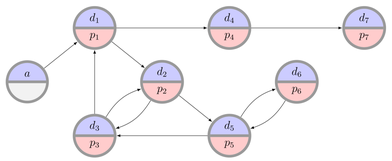

A kidney exchange finds the socially optimal set of cycles/chains that maximize a weighted edge sum. Complicated matching policies aim to balance match efficiency, hospital incentives, and ethical constraints. Not all transplants are equal, and waiting times can differ dramatically depending on patient features. All of this leads to an overall lack of transparency. **We aim to make KPD more transparent by developing and evaluating a system that forecasts patient outcomes.**

A Simple Approach

We model the kidney exchange as a dynamic graph $G(T) = G(V(T), E(T))$ where the vertex set evolves by the equation

$$V(T) = V(T-1) \cup A(T) \setminus D(T)$$

Here, $A(T)$ and $D(T)$ denote the arrivals and departures to the exchange. The edge set $E(T)$ is determined by the matching policy. At time T, exchanges have access to a *match record* $\mathcal{R}_T$ given by

$$\mathcal{R}_T := \left\{ (v, O(v), W(v), Q(v)) \mid v \in \bigcup_{t=1}^{T} D(t) \right\}$$

where the outcome $O(v)$ denotes whether the vertex was matched, $W(v)$ denotes the waiting time, and $Q(v)$ denotes the quality of the transplant. We use a simple random-forest approach to forecast outcomes (O, W, Q) from the match record $\mathcal{R}_T$.

Abstract

Kidney exchanges allow patients with end-stage renal disease to find a lifesaving living donor by way of an organized market. However, not all patients are equally easy to match, nor are all donor organs of equal quality—some patients are matched within weeks, while others may wait for years with no match offers at all. We propose the first decision-support tool for kidney exchange that takes as input the biological features of a patient-donor pair, and returns (i) the probability of being matched prior to expiry, and (conditioned on a match outcome), (ii) the waiting time for and (iii) the organ quality of the matched transplant. This information may be used to inform medical and insurance decisions. We predict all quantities (i, ii, iii) exclusively from match records that are readily available in any kidney exchange using a quantile random forest approach. To evaluate our approach, we developed two state-of-the-art realistic simulators based on data from the United Network for Organ Sharing that sample from the training and test distribution for these learning tasks—in our application these distributions are distinct. We analyze distributional shift through a theoretical lens, and show that the two distributions converge as the kidney exchange nears steady-state. We then show that our approach produces clinically-promising estimates using simulated data. Finally, we show how our approach, in conjunction with tools from the model explainability literature, can be used to calibrate and detect bias in matching policies.

SOTA Simulation Framework

We developed a simulation framework to sample match records and patient trajectories. Our simulators are state-of-the-art, using real data from the OPTN exchange.

Batch Simulation (up until time T) to obtain $\mathcal{R}_T$:

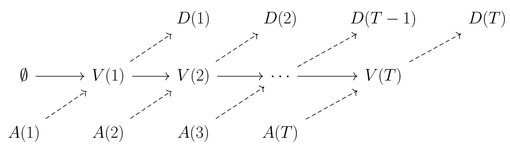

Trajectory Simulation (run τ times for S samples) to sample $(v, O(v), W(v), Q(v))$ for $v \sim f_P$:

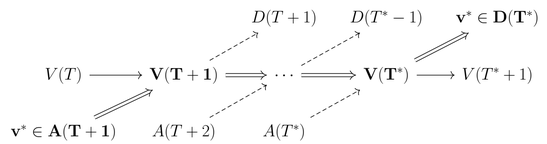

Distributional Shift and Steady-State Exchanges

D	REC_A	REC_B	DON_A	DON_B
1000	0.32	0.21	0.78	0.54
50000	0.22	0.15	0.78	0.48
Test	**0.24**	**0.17**	**0.79**	**0.50**

Distributional Shift. While we aim to forecast outcomes for patients *entering* the exchange, we only have data (in the form of match records) for patients who have *exited* the exchange. This creates a distributional shift. However, this shift appears to vanish as the exchange gets older! We analyze this effect through the lens of theory.

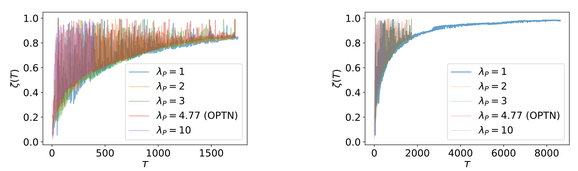

Steady-State Exchanges. Independently, practitioners have observed that kidney exchanges approach a steady-state over time. We confirm this empirically in our simulations. We measure distance to steady-state using the constant $\zeta(T) := \frac{\sum_{t=1}^{T} |D(t)|}{\sum_{t=1}^{T} |A(t)|} = \frac{|\mathcal{R}_T|}{|\mathcal{R}_T| + |V(T)|} \in [0, 1]$.

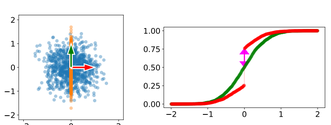

Shifted Directions. We say that a unit vector $\mathbf{z}$ is δ-shifted if the Kolmogorov distance between the one-dimensional projections of the data onto $\mathbf{z}$ is at least δ. We say that $\mathcal{R}_T$ is (γ, δ)-shifted if at least a γ fraction of all unit directions are δ-shifted.

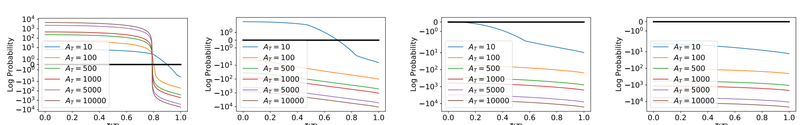

Theorem. *Suppose each vertex $\boldsymbol{v}_i \in \bigcup_{t=1}^{T} A(t)$ has features distributed as $\mathcal{N}(\boldsymbol{\mu}, \Sigma)$ where Σ is full rank. Then,*

$$\Pr\left[\mathcal{R}_T \text{ is } (\gamma, \delta)\text{-shifted}\right] \leq \underbrace{\left(\frac{e}{1-\zeta(T)}\right)^{A_T(1-\zeta(T))}}_{\text{Number of coalitions}} \overbrace{2^{\lceil \gamma d \rceil} \exp\left(-2A_T\zeta(T)\lceil \gamma d \rceil \delta^2\right)}^{\text{Probability that a fixed coalition is shifted}}$$

Experiments

Arrival Rate	D	Federated	$\lvert\mathcal{R}_T\rvert$	ζ	MAE $\left(\widehat{O}\right)$	IOU $\left(\widehat{W}_{95}\right)$	IOU $\left(\widehat{Q}_{95}\right)$
$\lambda_P = 1$	1500	No	56	0.397	0.258	0.451	**0.747**
$\lambda_P = 1$	4000	No	246	**0.953**	**0.191**	**0.644**	**0.761**
$\lambda_P = 1$	50000	No	4888	**0.984**	**0.130**	**0.653**	**0.632**
$\lambda_P = 2$	1500	No	157	0.477	0.221	0.336	**0.815**
$\lambda_P = 2$	4000	No	752	**0.882**	0.212	**0.620**	**0.809**
$\lambda_P = 3$	1500	No	285	0.523	**0.184**	0.386	**0.798**
$\lambda_P \approx 4.77$ (OPTN)	1500	No	593	0.509	**0.164**	**0.503**	**0.812**
$\lambda_P = 1$	1500	Yes	268	0.457	0.246*	0.232	**0.816***
$\lambda_P = 1$	4000	Yes	1224	**0.891**	**0.148***	**0.590**	**0.800***
$\lambda_P = 2$	1500	Yes	807	0.550	**0.145***	0.373*	**0.816***
$\lambda_P = 2$	4000	Yes	3773	**0.872**	**0.119***	**0.775***	**0.820***
$\lambda_P = 3$	1500	Yes	1434	0.488	**0.115***	0.421*	**0.815***
$\lambda_P \approx 4.77$ (OPTN)	1500	Yes	2652	0.537	**0.103***	0.449	**0.812**

Experimental Results. We bold steady-state parameters $\zeta > 0.8$, MAE scores < 0.2, and IOU scores > 0.5. We asterisk any federated learning experiments that improve relative performance.

Diagnosing Mechanism Behavior

Using SHAP and TSNE, we can use our model understand the underlying matching mechanism and detect potential miscalibrations.

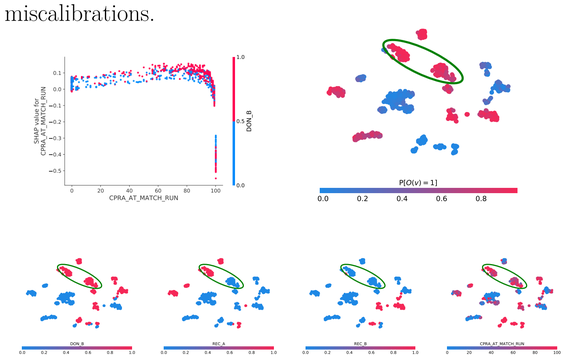

Visualizing Miscalibrations. Counterintuitively, we find that hard-to-match patients with easy-to-match donors are almost always matched

Summary of Contributions

1. We proposed a random-forest approach

	Old	Young
Large	O, W, Q	W, Q
Small	O, W, Q	Q

2. High values of ζ give a proxy for success
3. Our approach can be used to inform policy and make kidney exchanges more fair
4. We developed a state-of-the-art simulation framework