Achieving Long-term Fairness in Sequential Decision Making

Yaowei Hu (yaoweihu@uark.edu)
Lu Zhang (lz006@uark.edu)

University of Arkansas

Fair Machine Learning

- Fair machine learning plays an important role in decision making tasks such as hiring, college admissions and bank loans.

- Various fairness notions: *demographic parity*, *equalized odds* and *counterfactual fairness*.

- However, the majority of studies on fair machine learning focus on the static or one-shot classification setting.

Decision Making Systems Are Dynamic

- In practice, decision making systems are usually operating in a dynamic manner such that the classifier makes sequential decisions over a period of time.

Example:

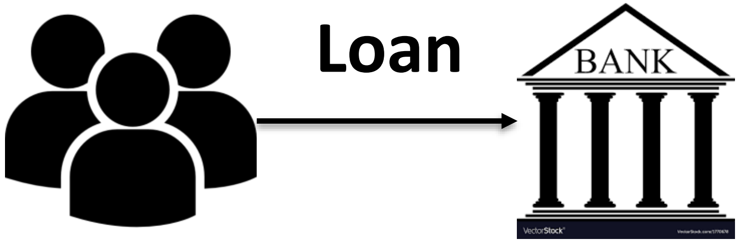

The bank uses the applicants' credit scores to make loan decisions

The bank's decisions will affect the applicants' credit scores

Feedback Loop

Long-term Fairness

- Fair decision making should concern not only the fairness of a single decision but more importantly, whether a decision model can impose fair long-term effects on different groups.

- This notion of fairness is referred to as long-term fairness in recent studies.

The Challenges

- The challenges of achieving long-term fairness:

  - **Feedback Loop**. Without knowing how the population would be reshaped by decisions, enforcing any fairness constraint may create negative feedback loops and eventually harm fairness in the long run.

  - **Distribution Shift**. Ignoring the distribution shift will critically affect the achievement of long-term fairness, as long-term fairness is affected by all decisions made by the model along the time.

Main Contributions

- We propose a causality-based long-fairness notion.
- We propose a framework for formulating long-term fair sequential decision making as performative risk optimization problem.
- Experiments on both synthetic and real-world datasets show that the proposed method can achieve long-term fairness in multiple time steps.

Preliminaries

- Structural Causal Model

A structural causal model M is represented by a quadruple $< \boldsymbol{U}, \boldsymbol{V}, \boldsymbol{F}, >$ where,

1. $\boldsymbol{U}$ *is a set of exogenous random variables that are determined by factors outside the model.*
2. $\boldsymbol{V}$ *is a set of endogenous variables that are determined by variables in* $\boldsymbol{U} \cup \boldsymbol{V}$.
3. $\boldsymbol{F}$ *is a set of structural equations from* $\boldsymbol{U} \cup \boldsymbol{V}$ *to* **V.** *Specifically, for each* $V \in \boldsymbol{V}$, *there is a function* $f_V \in \boldsymbol{F}$ *mapping from* $\boldsymbol{U} \cup (\boldsymbol{U} \backslash \boldsymbol{V})$ *to* **V**, *i.e,* $v = f_V(\mathrm{pa}_{v,} u_{\mathrm{v}})$, *where* pa_v *and* u_{v} *are realization of a set of endogenous variables* $\mathrm{PA}_{\mathrm{v}} \in \mathbf{V} \backslash \mathrm{V}$ *and a set of exogenous variables* U_V *respectively.*

- We assume the Markovian model in this work.

Preliminaries

- Hard Intervention

An intervention on endogenous variable X is defined as the substitution of structural equation $f_X(\mathrm{PA}_X,\ U_X)$ with a constant x, denoted as $do(X = x)$ or *$do(x)$ for short*.

- Soft Intervention

An intervention on endogenous variable X is defined as the substitution of structural equation $f_X(\mathrm{PA}_X,\ U_X)$ with a new function $x = g_\theta(\boldsymbol{z})$, which is denoted by θ*.*

Notation Table

Symbol	Description
$D = \{(S, \boldsymbol{X}^t, Y^t)\}_{t=1}^{l}$	A temporal dataset
$S = \{s^+, s^-\}$	A time-invariant protected attribute
$\boldsymbol{X}^t$	A set of time-dependent unprotected attributes
$Y^t = \{+1, -1\}$	A time-dependent class label
$\hat{Y}^t = h_\theta(S, \boldsymbol{X}^t)$	A predictive decision model parameterized by θ. $\hat{Y}^t$=1 if $h_\theta(S, \boldsymbol{X}^t) \geq 0$ and $\hat{Y}^t$=-1 if $h_\theta(S, \boldsymbol{X}^t) < 0$

Causality-based Long-term Fairness

- Based on SCM, we assume a time-lagged causal graph G for describing the causal relationship among variables over time.

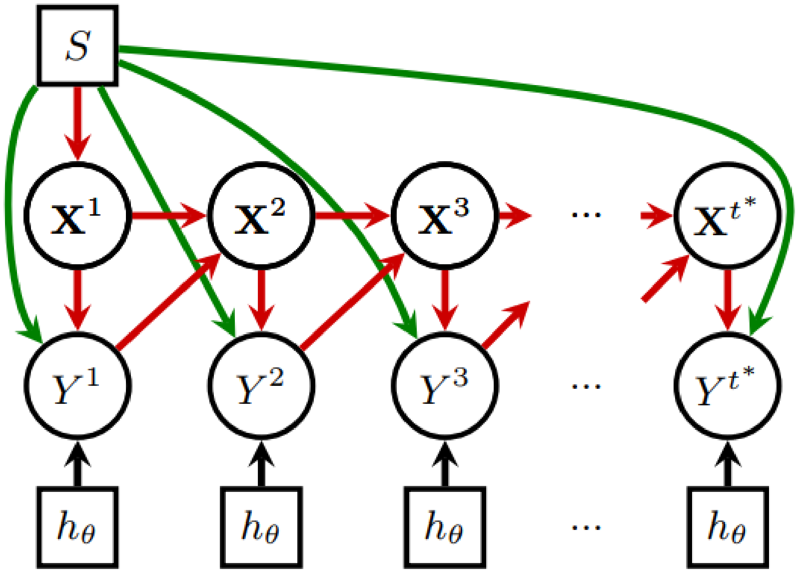

Causality-based Long-term Fairness

- **long-term fairness.** Formulated as path-specific effects that are transmitted in the time-lagged causal graph along certain paths.

$\boldsymbol{X}$ {
- $\boldsymbol{X}_i$ irrelevant attributes: justifiable in decision making, evolved autonomously or altered by external factors.
- $\boldsymbol{X}_r$ relevant attributes: the remaining

Definition 1 (Long-term Fairness). *The long-term fairness of a decision model* h_θ is measured by $P\left(\hat{Y}^{t*}(s_\pi^+, \theta)\right) - P\left(\hat{Y}^{t*}(s_\pi^-, \theta)\right)$ where π is a set of paths from S to $\hat{Y}^{t*}$ passing through $\boldsymbol{X}_r^1, \hat{Y}^1, \ldots, \boldsymbol{X}_r^{t*-1}, \hat{Y}^{t*-1}, \boldsymbol{X}_r^{t*}$, s_π represents the path-specific hard intervention and θ represents the soft intervention through all paths.

Loss Function and Short-term Fairness

- Two other requirements:

  - **Short-term Fairness.** The decision model should also satisfy certain short-term fairness requirement at each time step to enforce local equality, which may be stipulated by law or regulations.

Definition 2 (Short-term Fairness). *The short-term fairness of a decision model* h_θ at time t is measured by the causal effect transmitted through paths involved in time t , i.e., $P\left(\hat{Y}^t(s^+_{\pi^t},\theta)\right) - P\left(\hat{Y}^t(s^-_{\pi^t},\theta)\right)$, where $\pi^t = \{S \to \tilde{\boldsymbol{X}}_r \to \hat{Y}^t, S \to \hat{Y}^t\}$ with redlining attributes $\tilde{\boldsymbol{X}}_r$, s_π is the path-specific hard intervention and θ represents the soft intervention.

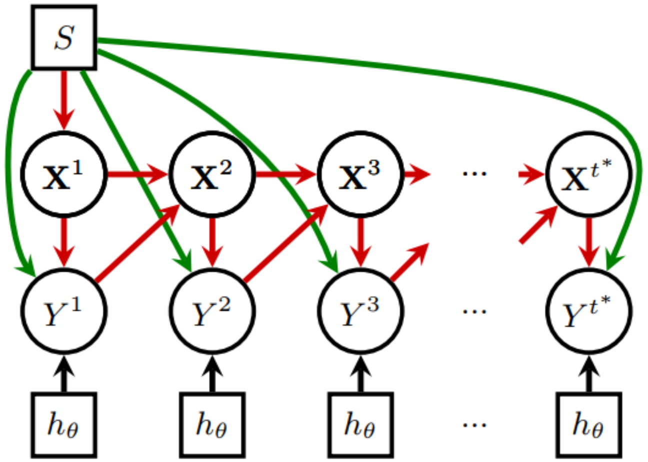

Loss Function and Short-term Fairness

- Two other requirements:
  - **Institution Utility.** It is a natural desire for a predictive decision model to maximize the institution utility.

Definition 3 (Institute Utility). *The institution utility of a decision model* h_θ is measured by the aggregate loss given by $\sum_{t=1}^{t*} E\left[L\left(Y^t, \hat{Y}^t\right)\right]$ where $L(\cdot)$ is the loss function.

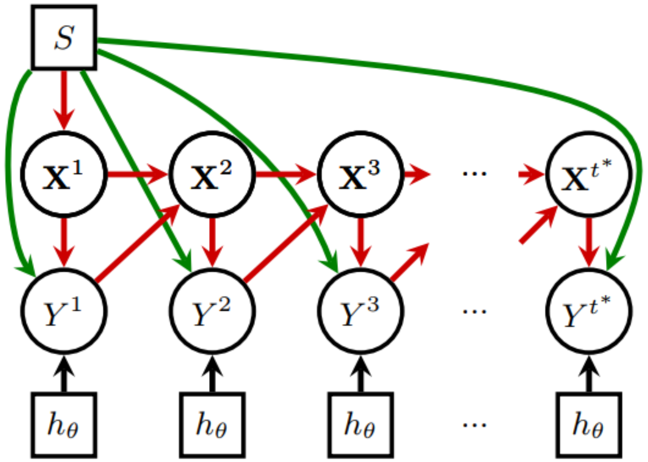

Soft Intervention for Model Deployment

- In all three definitions, we use soft intervention for modeling the decision model deployment.
  - We treat the deployment of the decision model at each time step as to perform a soft intervention on the decision variable.
  - The change to underlying population could be inferred as the post-intervention distribution after performing the soft intervention.

Learning Fair Decision Model

The goal is to learn a functional mapping $h_\theta:\ (S, \boldsymbol{X}^t)\ \rightarrow Y^t$ parameterized with θ. By meeting the two requirements of institution utility and short-term fairness, the functional mapping will achieve long-term fairness.

Problem Formulation 1. *The problem of fair sequential decision making is formulated as the constrained optimization:*

$$\arg\min_{\theta} \sum_{t=1}^{t*} E\left[L\left(Y^t, \hat{Y}^t\right)\right]$$

$$s.t.\ \ P\left(\hat{Y}^{t*}(s_\pi^+, \theta) = 1\right) - P\left(\hat{Y}^{t*}(s_\pi^-, \theta) = 1\right) \leq \tau_l$$

$$P\left(\hat{Y}^t\left(s_{\pi^t}^+, \theta\right) = 1\right) - P\left(\hat{Y}^t\left(s_{\pi^t}^-, \theta\right) = 1\right) \leq \tau_t,\ t = 1, \dots, t^*$$

where τ_l and τ_t are thresholds for long-term fairness and short-term fairness constraints, respectively.

Performative Risk Optimization

- Solving the optimization problem in Problem Formulation 1 is not trivial.

$$P\Big(\hat{Y}^t(s^+_\pi, \theta) = 1\Big) = \sum_{\mathbf{X}^1, Y^1, \cdots, \mathbf{X}^{t^*}} \Big\{ P(\mathbf{x}^1 | s^+) P_\theta(y^1 | \mathbf{x}^1, s^-) \cdots$$
$$\cdots P(\mathbf{x}^{t^*} | \mathbf{x}^{t^*-1}, y^{t^*-1}) P_\theta(Y^{t^*} = 1 | \mathbf{x}^{t^*}, s^-) \Big\},$$

- Convert Problem Formulation 1 to the form of performative risk optimization.
- The general formulation of the performative risk optimization is

$$\mathrm{PR}(\theta) = \mathop{\mathbb{E}}_{Z \sim \mathcal{D}(\theta)} \ell(Z; \theta) \,.$$

Performative Risk Optimization

- By adopting the surrogate function $\phi(\cdot)$, we can reformulate both loss function and constraints.

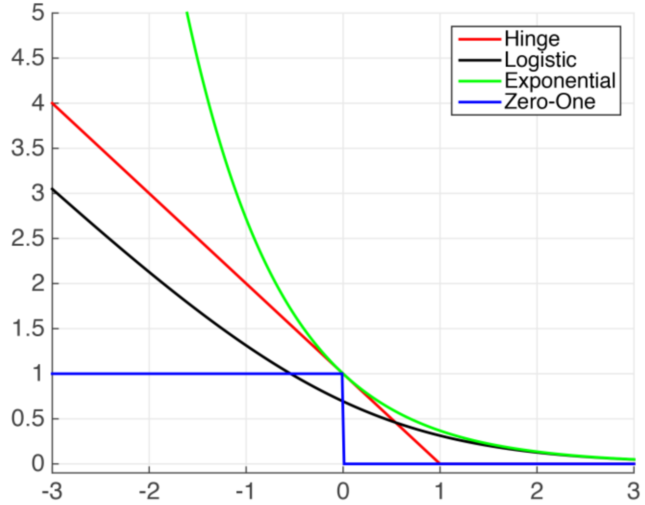

$$P\left(\hat{Y}^t(s_\pi^+,\theta)=1\right) = \sum_{\mathbf{X}^t} P(\mathbf{x}^t(s_\pi^+,\theta))\,\phi\left(-h_\theta\left(\mathbf{x}^t, s^-\right)\right)$$

$$= \underset{\mathbf{X}^t \sim P(\mathbf{X}^t(s_\pi^+,\theta))}{\mathbb{E}} \left[\phi\left(-h_\theta\left(\mathbf{X}^t, s^-\right)\right)\right].$$

Performative Risk Optimization

- Reformulate utility, short-term fairness and long-term fairness in the form of performative risk.

$$l_u(\theta) = \sum_{t=1}^{t^*} \mathop{\mathbb{E}}_{S,\mathbf{X}^t,Y^t \sim P(S,\mathbf{X}^t,Y^t)} \left[\phi\left(Y^t h_\theta(\mathbf{X}^t, S)\right)\right],$$

$$l_l(\theta) = \frac{1}{2}\Big\{ \mathop{\mathbb{E}}_{\mathbf{X}^{t^*} \sim P(\mathbf{X}^{t^*}(s_\pi^+, \theta))} \left[\phi\left(-h_\theta\left(\mathbf{X}^{t^*}, s^-\right)\right)\right]$$
$$+ \mathop{\mathbb{E}}_{\mathbf{X}^{t^*} \sim P(\mathbf{X}^{t^*}(s_\pi^-, \theta))} \left[\phi\left(h_\theta\left(\mathbf{X}^{t^*}, s^-\right)\right)\right] - 1 - \tau_l \Big\},$$

$$l_s(\theta) = \frac{1}{t^*} \sum_{t=1}^{t^*} \Big\{ \mathop{\mathbb{E}}_{\mathbf{X}^t \sim P(\mathbf{X}^t(s_{\pi^t}^+, \theta))} \left[\phi\left(-h_\theta\left(\mathbf{X}^t, s^-\right)\right)\right]$$
$$+ \mathop{\mathbb{E}}_{\mathbf{X}^t \sim P(\mathbf{X}^t(s_{\pi^t}^-, \theta))} \left[\phi\left(h_\theta\left(\mathbf{X}^t, s^-\right)\right)\right] - 1 - \tau_t \Big\}.$$

Performative Risk Optimization

Problem Formulation 2. *The problem of fair sequential decision making is reformulated as the performative risk optimization:*

$$\arg\min_{\theta} l(\theta) = \lambda_u l_u(\theta) + \lambda_l l_l(\theta) + \lambda_s l_s(\theta)$$

where λ_u, λ_l and λ_s are weight parameters and satisfy $\lambda_u + \lambda_l + \lambda_s = 1$.

Repeated Risk Minimization

- Repeated risk minimization (RRM) is an iterative algorithmic heuristic for solving the performative risk optimization problem.

Algorithm 1: Repeated Risk Minimization

Input : Dataset $\mathcal{D} = \{(S, \mathbf{X}^t, Y^t)\}_{t=1}^{l}$, time -lagged causal graph $\mathcal{G}$, convergence threshold δ

Output: The stable model h_θ

Train a classifier on $\mathcal{D}$ according to Eq. (2) without the soft intervention to obtain the initial parameter θ_0;
$i \leftarrow 0$;
repeat
  Sampled the post-intervention distributions $P\left(\mathbf{X}^{t^*}(s_\pi^+, \theta_i)\right)$ and $P\left(\mathbf{X}^{t^*}(s_\pi^-, \theta_i)\right)$;
  Sampled the post-intervention distributions $P\left(\mathbf{X}^{t}(s_\pi^+, \theta_i)\right)$ and $P\left(\mathbf{X}^{t}(s_\pi^-, \theta_i)\right)$ for each t;
  Minimize $l(\theta)$ according to Eq. (2) to obtain θ_{i+1};
  $\triangle = \|\theta_{i+1} - \theta_i\|_2$;
  $i \rightarrow i + 1$;
until $\triangle < \delta$;
$\theta \leftarrow \theta_i$;
return h_θ;

Convergence Analysis of RRM

- The convergence of the RRM algorithm depends on the smoothness and convexity of the loss function, as well as the sensitivity of the distribution to the parameters.

Theorem 1. *Suppose that surrogated loss function* $(\varphi \circ h)(\cdot)$ *is* β *–jointly smooth and* γ*-strongly convex, and suppose that* $\boldsymbol{X}^{t+1}$ *are* c*-sensitive for any* t*, then the repeated risk minimization converges to a stable point at a linear rate, if* $2mc(t^{*}-1) < \frac{\beta}{\gamma}$.

Experiments

- Baselines:
  - **Logistic Regression (LR)**: An unconstrained logistic regression model which takes user features and labels from all time steps as inputs and outputs.
  - **Fair Model with Demographic Parity (FMDP)**: On the basis of the logistic regression model, fairness constraint is added to achieve demographic parity.
  - **Fair Model with Equal Opportunity (FMEO)**: On the basis of the logistic regression model, fairness constraint is added to achieve equal opportunity.

Experiments

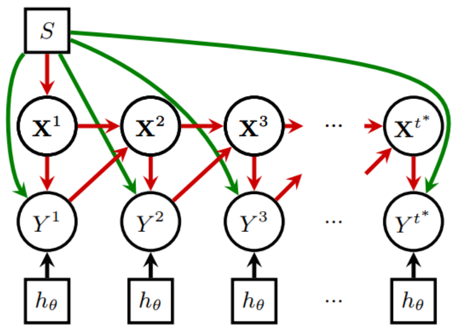

- Synthetic Data:

We simulate a process of bank loans following the above time-lagged causal graph, where S is the protected attribute like race, $\boldsymbol{X}^t$ represents the financial status of applicants, and Y^t represents the decisions about whether to grant loans.

We sample the predicted decisions from:

$$P(\hat{Y}^t) = \sigma(h_{\theta^*}(\mathbf{X}^t, S)),\ \hat{Y}^t \sim 2 \cdot \text{Bernoulli}(P(\hat{Y}^t)) - 1.$$

$\boldsymbol{X}^{t+1}$ is generated according to the update rule below:

$$\mathbf{X}^{t+1} = \begin{cases} \mathbf{X}^t - \epsilon \cdot \theta^t + b & \hat{Y}^t = 1, Y^t = -1 \\ \mathbf{X}^t + \epsilon \cdot \theta^t + b & \hat{Y}^t = 1, Y^t = 1 \\ \mathbf{X}^t + b & \hat{Y}^t = -1 \end{cases}$$

Experiments

- Semi-synthetic Data:

- Use the Taiwan credit card dataset (Yeh and Lien 2009) as the initial data at t = 1
- Extract 3000 samples and choose two features PAY AMT1 and PAY AMT2
- Generate a 4-step dataset using similar update rule

Yeh, I.-C.; and Lien, C.-h. 2009. The comparisons of data mining techniques for the predictive accuracy of probability of default of credit card clients. Expert Systems with Applications, 36(2): 2473–2480.

Experiments

- Results of Synthetic Data:

Table 1: Accuracy, short-term and long-term fairness of different algorithms on the synthetic dataset.

Alg.	Metric	Time steps				
		$t=1$	$t=2$	$t=3$	$t=4$	$t=5$
RL	Acc	0.912	0.894	0.917	0.921	0.917
	Short	0.152	0.160	0.166	0.164	0.174
	Long	0.058	0.117	0.173	0.246	0.340
FMDP	Acc	0.735	0.706	0.704	0.708	0.725
	Short	0.212	0.216	0.224	0.220	0.232
	Long	0.180	0.306	0.376	0.431	0.481
FMEO	Acc	0.829	0.790	0.795	0.800	0.814
	Short	0.010	0.010	0.010	0.014	0.020
	Long	0.080	0.122	0.190	0.276	0.352
Ours	Acc	0.801	0.754	0.729	0.707	0.692
	Short	0.012	0.008	0.012	0.008	0.002
	Long	0.040	0.024	0.020	0.012	0.002

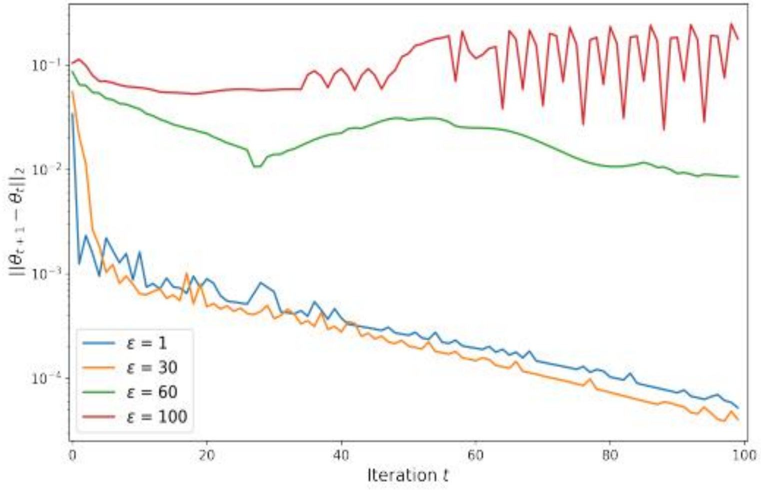

Figure 2: The convergence results for different values of ϵ on the synthetic dataset.

Experiments

- Results of Semi-synthetic Data:

Table 2: Accuracy, short-term and long-term fairness of different algorithms on the semi-synthetic dataset.

Alg.	Metric	Time steps			
		$t=1$	$t=2$	$t=3$	$t=4$
RL	Acc	0.828	0.826	0.841	0.816
	Short	0.015	0.018	0.021	0.012
	Long	0.038	0.088	0.243	0.433
FMDP	Acc	0.830	0.843	0.846	0.841
	Short	0.063	0.066	0.075	0.069
	Long	0.038	0.076	0.223	0.397
FMEO	Acc	0.824	0.830	0.830	0.813
	Short	0.072	0.075	0.087	0.078
	Long	0.006	0.045	0.156	0.295
Ours	Acc	0.648	0.648	0.680	0.687
	Short	0.006	0.006	0.003	0.006
	Long	0.064	0.043	0.016	0.003

Conclusions

- Proposed a framework to achieve long-term fairness in sequential decision making.
- Measured both long-term and short-term fairness as path-specific effects in a time-lagged causal graph.
- Formulated as a performative risk optimization problem, and repeated risk minimization is adopted to train the model on the datasets sampled from post-intervention distributions.
- Theoretically analyzed the convergence of the proposed algorithm.
- Conducted experiments on both synthetic and semi-synthetic datasets to show the effectiveness of the proposed framework and algorithm.

UNIVERSITY OF ARKANSAS

Achieving Long-term Fairness in Sequential Decision Making

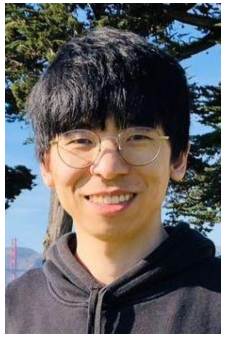

Yaowei Hu **Lu Zhang**

This work was supported in part by NSF 1910284, 1920920, and 1946391.