

Recommendations as Treatments: Debiasing Learning and Evaluation

ICML 2016, NYC

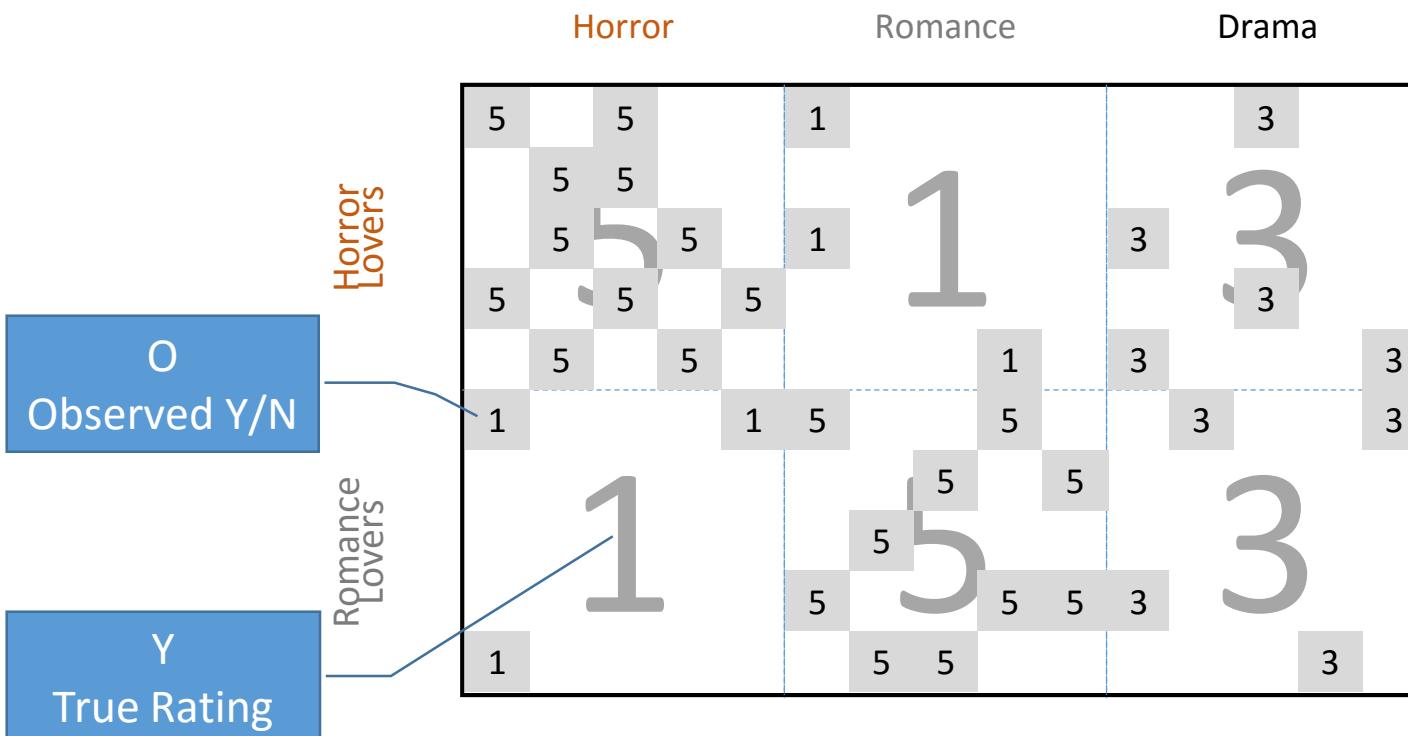
Tobias Schnabel[†], Adith Swaminathan[†],
Ashudeep Singh[†], Navin Chandak[§], Thorsten Joachims[†]



[†]Cornell University, [§]Google

Funded in part through NSF Awards IIS-1247637, IIS-1217686, IIS-1513692.

Movie recommendation



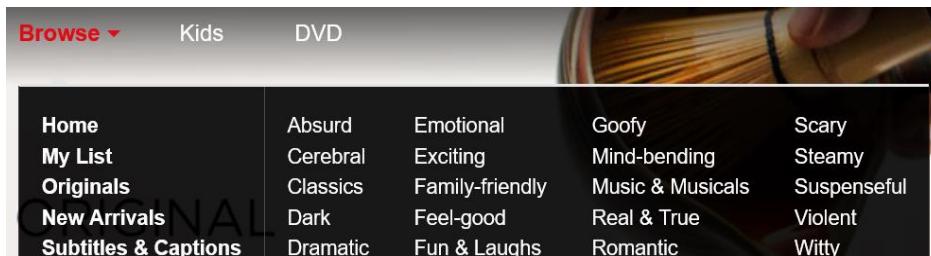
⇒ Data is Missing Not At Random (MNAR)

Example adapted from (Steck et al., 2010)

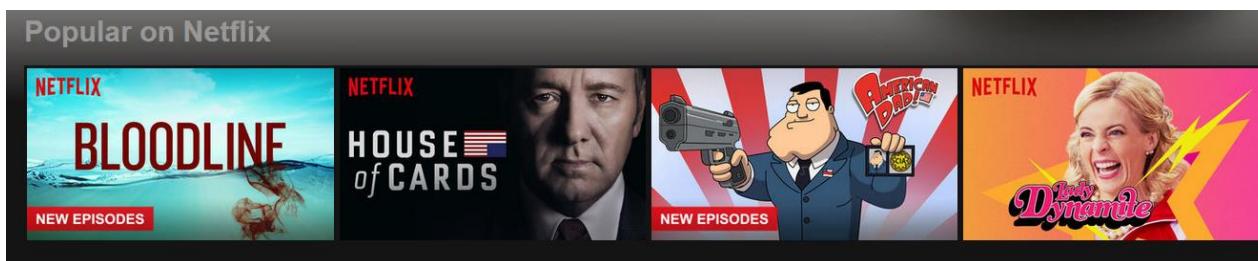
Selection Bias in Recommendation

- **Why is there selection bias?**

- User-induced bias (e.g., browsing)



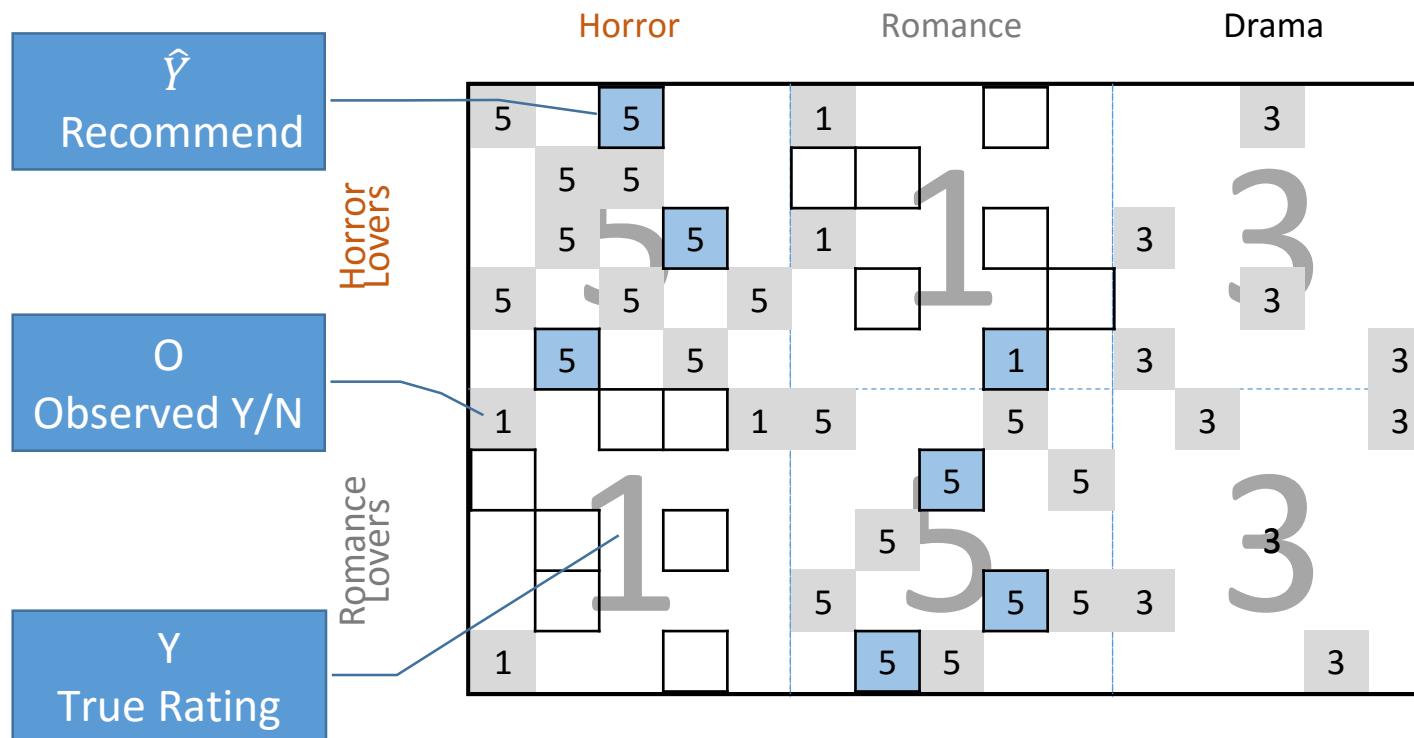
- System-induced bias (e.g., advertising)



- **Question:** What happens if we ignore selection bias?

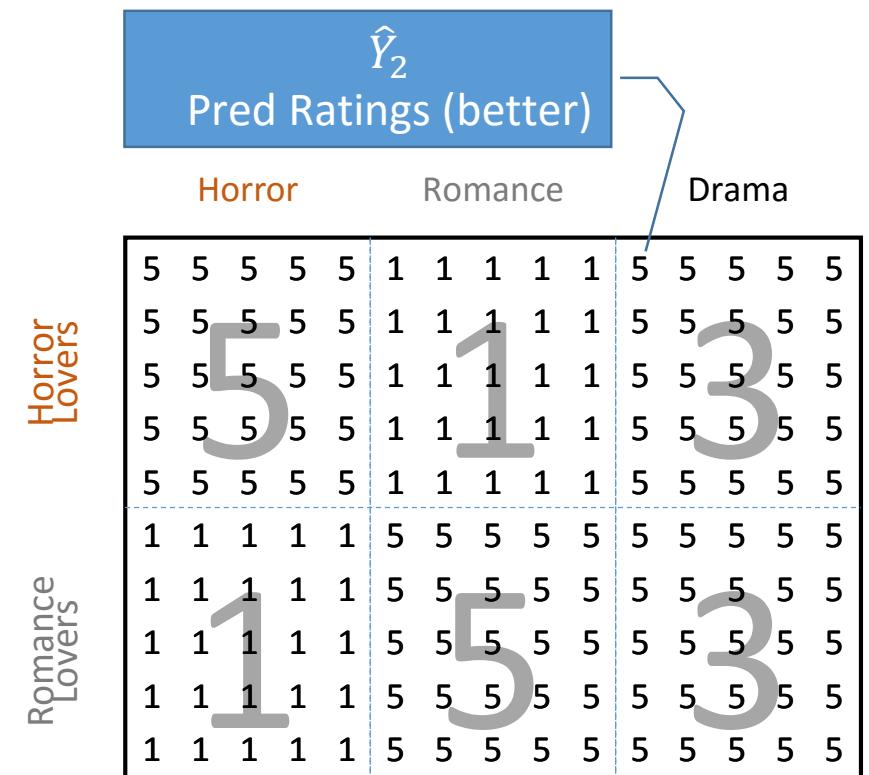
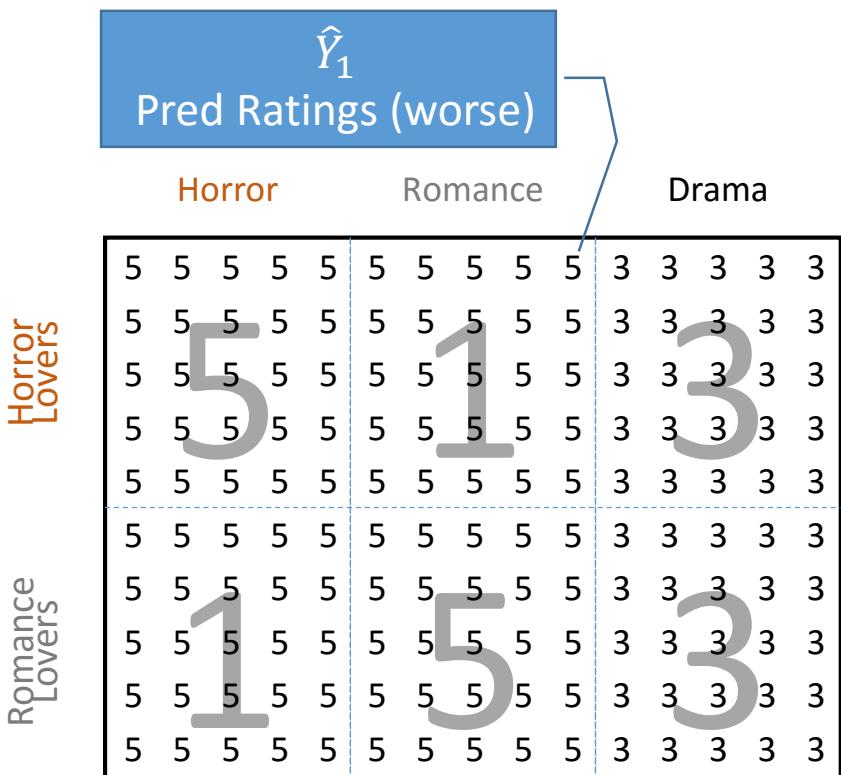
(Marlin et al., 2007; Steck, 2011;
Hernández-Lobato et al., 2014)

Evaluating Recommendations under Selection Bias



⇒ Observed ratings are misleading due to selection bias

Evaluating Predicted Ratings under Selection Bias

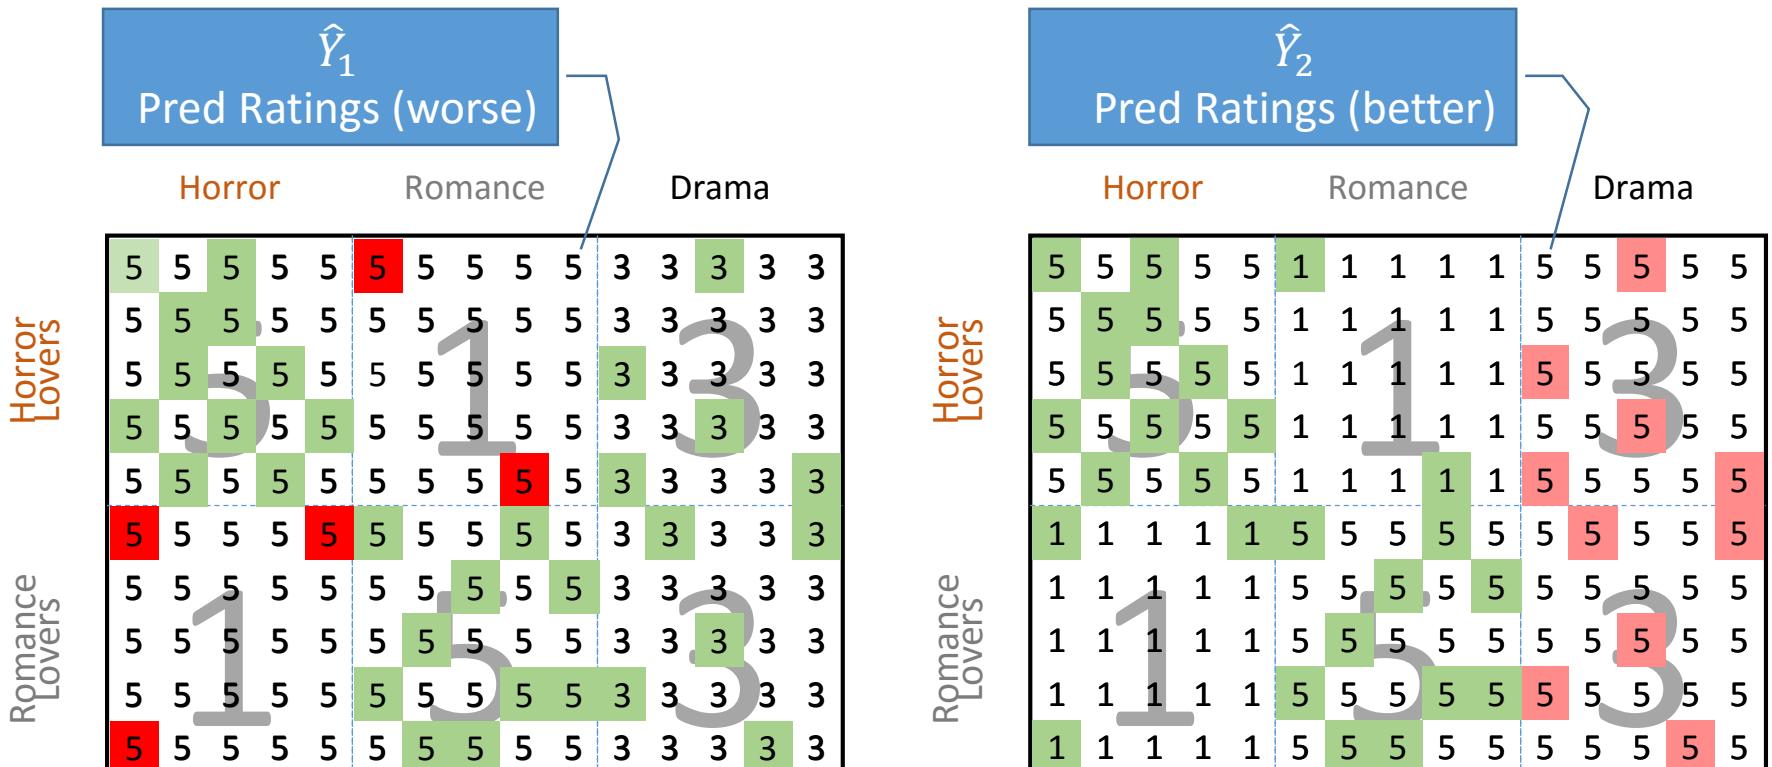


Evaluating Predicted Ratings under Selection Bias

		\hat{Y}_1 Pred Ratings (worse)														
		Horror					Romance					Drama				
Horror Lovers	Horror	5	5	5	5	5	5	5	5	5	5	3	3	3	3	3
	Romance	5	5	5	5	5	5	5	5	5	5	3	3	3	3	3
	Drama	5	5	5	5	5	5	5	5	5	5	3	3	3	3	3
	Horror	5	5	5	5	5	5	5	5	5	5	3	3	3	3	3
	Romance	5	5	5	5	5	5	5	5	5	5	3	3	3	3	3
	Drama	5	5	5	5	5	5	5	5	5	5	3	3	3	3	3
	Horror	5	5	5	5	5	5	5	5	5	5	3	3	3	3	3
	Romance	5	5	5	5	5	5	5	5	5	5	3	3	3	3	3
	Drama	5	5	5	5	5	5	5	5	5	5	3	3	3	3	3
	Horror	5	5	5	5	5	5	5	5	5	5	3	3	3	3	3

		\hat{Y}_2 Pred Ratings (better)														
		Horror					Romance					Drama				
Romance Lovers	Horror	5	5	5	5	5	5	5	5	5	5	1	1	1	1	1
	Romance	5	5	5	5	5	5	5	5	5	5	1	1	1	1	1
	Drama	5	5	5	5	5	5	5	5	5	5	1	1	1	1	1
	Horror	5	5	5	5	5	5	5	5	5	5	1	1	1	1	1
	Romance	5	5	5	5	5	5	5	5	5	5	1	1	1	1	1
	Drama	5	5	5	5	5	5	5	5	5	5	1	1	1	1	1
	Horror	5	5	5	5	5	5	5	5	5	5	1	1	1	1	1
	Romance	5	5	5	5	5	5	5	5	5	5	1	1	1	1	1
	Drama	5	5	5	5	5	5	5	5	5	5	1	1	1	1	1
	Horror	5	5	5	5	5	5	5	5	5	5	1	1	1	1	1

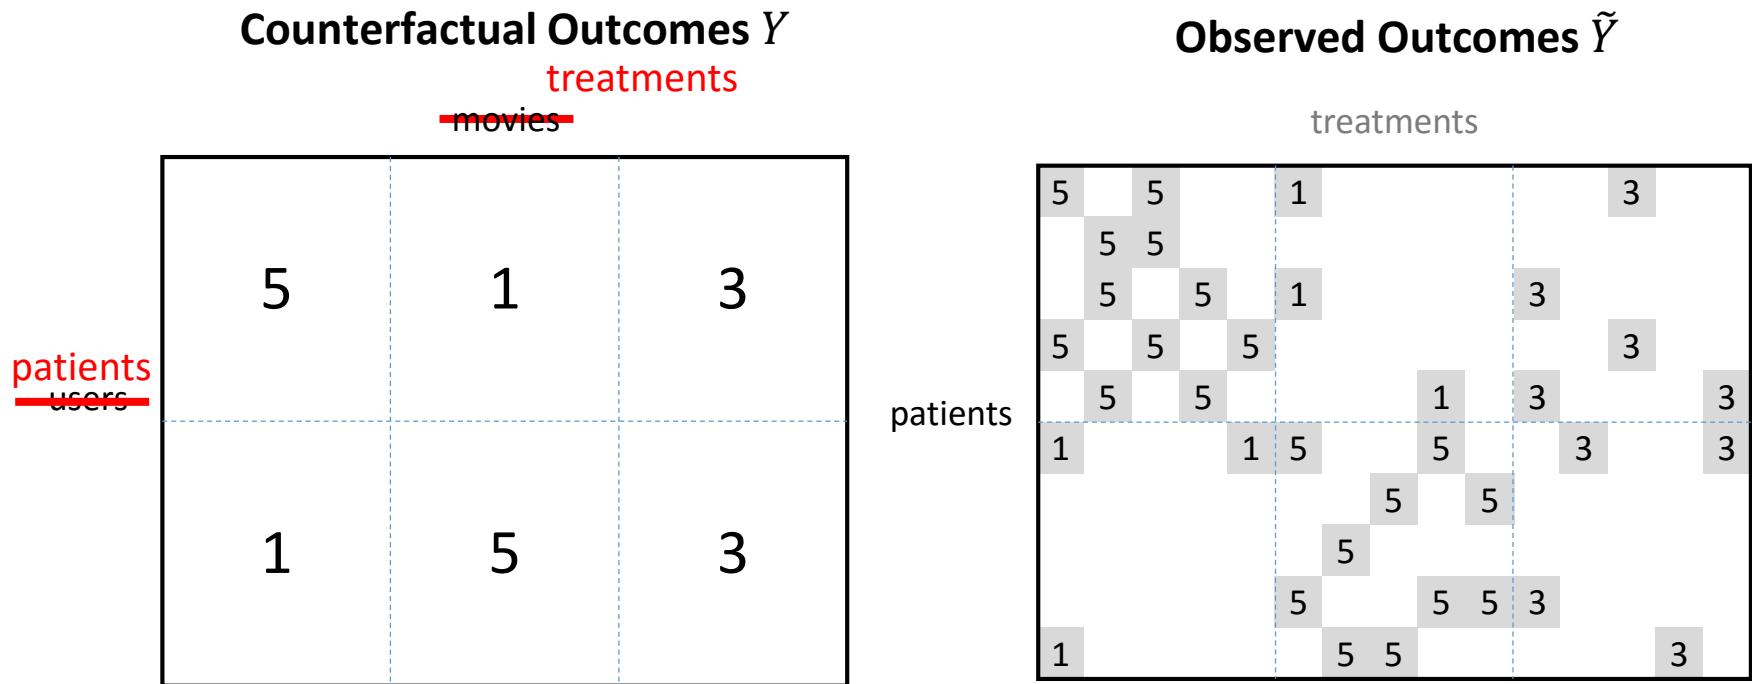
Evaluating Predicted Ratings under Selection Bias



⇒ Observed losses are misleading due to selection bias

Recommendations as Treatments

- **Question:** How can we fix the effects of selection bias?
 - Connection to **potential outcomes framework**



⇒ Understand **assignment mechanism** (Imbens & Rubin, 2015)

Debiasing Evaluation

- **Assignment mechanism** for recommendation:
 - $P_{u,i} = P(O_{u,i} = 1)$
- Use **Inverse-Propensity-Scoring Estimator (IPS)** to obtain unbiased estimate:

$$\hat{R}_{IPS}(\hat{Y}|P) = \frac{1}{U \cdot I} \sum_{(u,i): O_{ui}=1} \frac{1}{P_{u,i}} (Y_{u,i} - \hat{Y}_{u,i})^2$$

(Little & Rubin, 2002; Cortes et al., 2008; Bickel et al., 2009; Sugiyama & Kawanabe, 2012).

Propensities P		
Horror	Romance	Drama
p	p/10	p/2
p/10	p	p/2

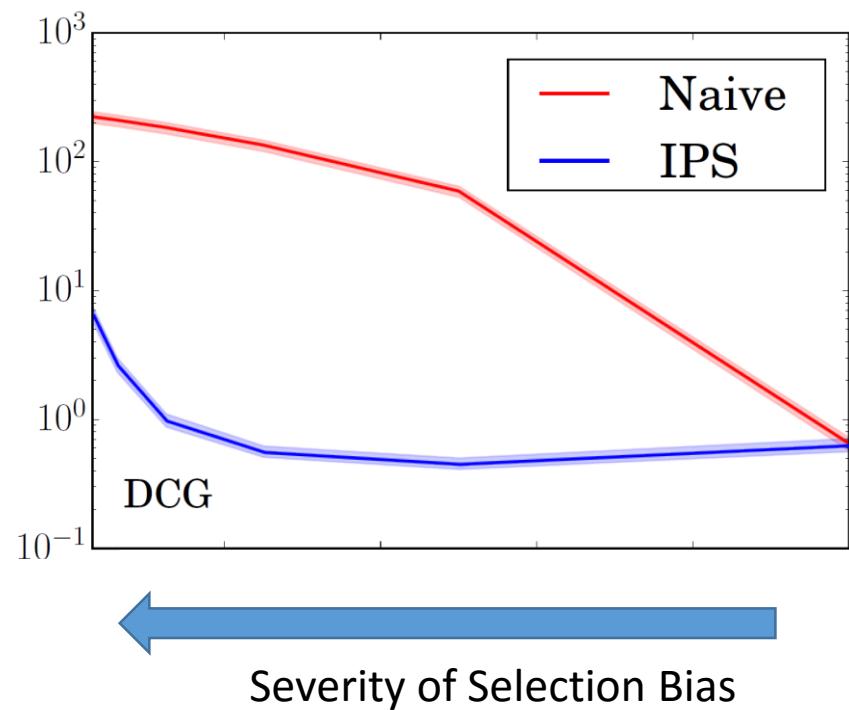
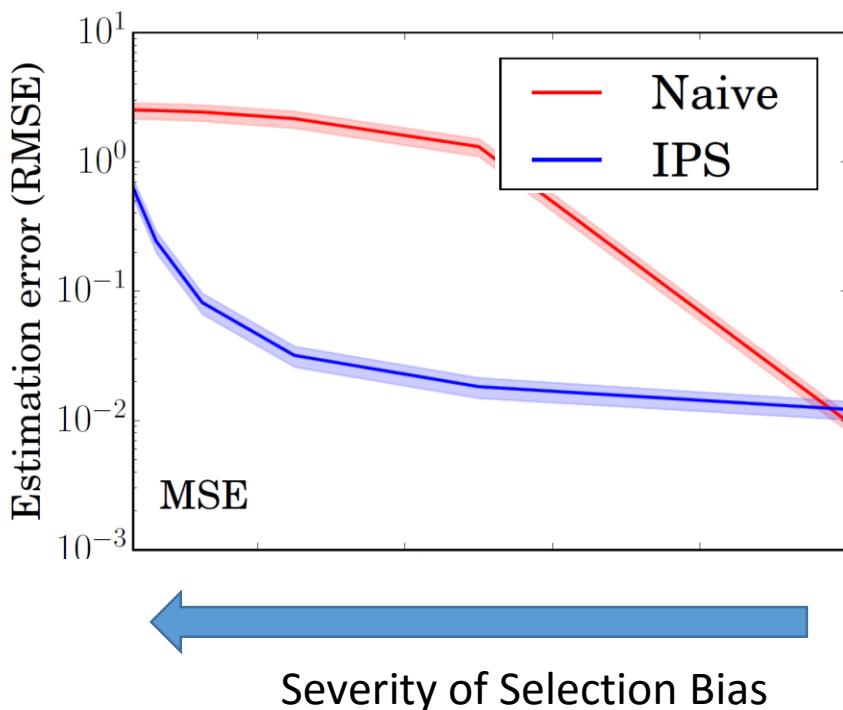
Propensity estimation

- Two settings:
 - **Experimental** – Propensities are under our control; known by design (e.g., ad placement)
 - **Observational** – Users self-select; need to estimate $P_{u,i}$
 - Estimate parameter of **binary** random variables:
$$P_{u,i} = P(O_{u,i} = 1 | X, \tilde{Y})$$
 - **Variety of models:** Logistic Regression, Naïve Bayes, etc.

Observations O			
Horror	Romance	Drama	
1 0 1 0 0	1 0 0 0 0	0 0 1 0 0	
0 1 1 0 0	0 0 0 0 0	0 0 0 0 0	
0 1 0 1 0	1 0 0 0 0	1 0 0 0 0	
1 0 0 0 0	0 0 0 0 0	0 0 1 0 0	
0 1 0 1 0	0 0 0 1 0	1 0 0 0 1	
1 0 0 0 1	1 0 0 1 0	0 1 0 0 1	
0 0 0 0 0	0 0 1 0 1	0 0 0 0 0	
0 0 0 0 0	0 1 0 0 0	0 0 1 0 0	
0 0 0 0 0	1 0 0 0 1	1 0 0 0 0	
1 0 0 0 0	0 1 1 0 0	0 0 0 1 0	

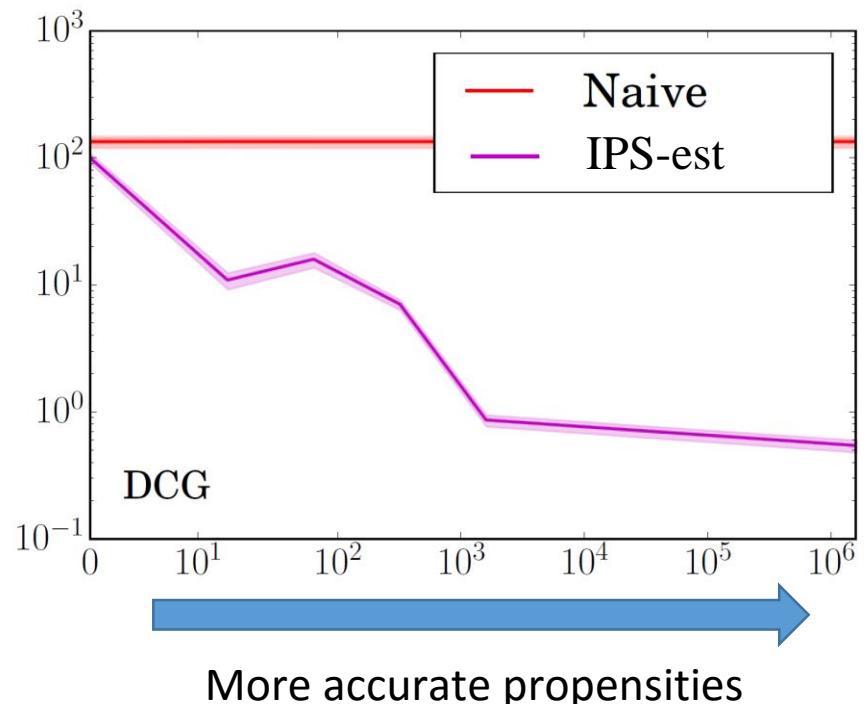
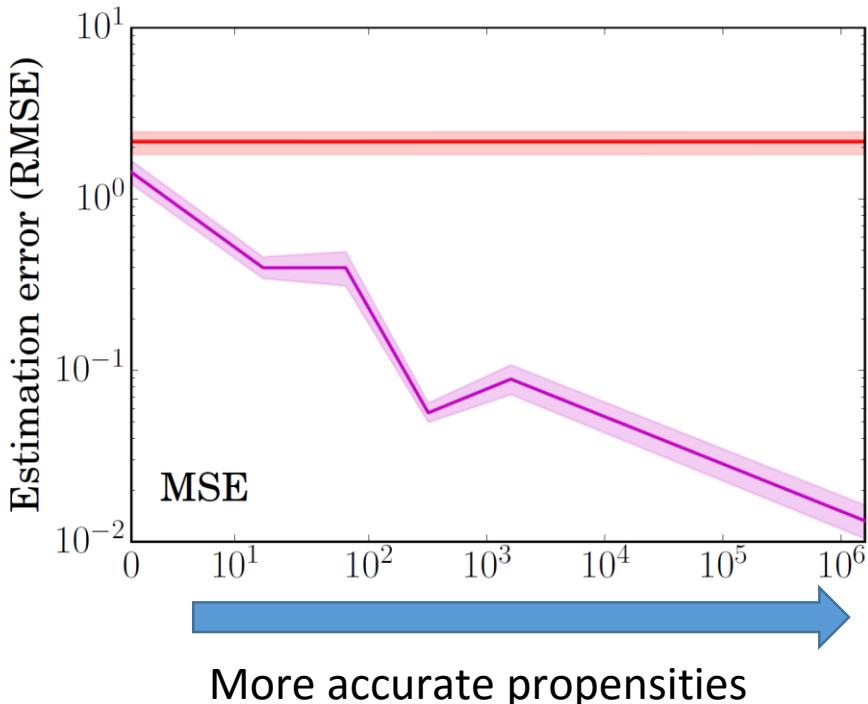
Debiasing Evaluation

- Robustness to selection bias:



Debiasing Evaluation

- Robustness to inaccurate propensities:



Debiasing Learning

- **Empirical Risk Minimization (ERM)** successful in many settings (Cortes & Vapnik, 1995)
- Use **ERM** together with **Inverse-Propensity-Scoring Estimator (IPS)**

$$\hat{Y}^{ERM} = \operatorname{argmin}_{\hat{Y} \in \mathcal{H}} \{ \hat{R}_{IPS}(\hat{Y} | P) \}$$

- For **matrix factorization** with MSE loss:

$$\hat{Y}^{ERM} = \operatorname{argmin}_{V,W} \left\{ \sum_{o_{u,i}=1} \frac{1}{P_{u,i}} (Y_{u,i} - V_u W_i)^2 + \lambda (\|V\|_F^2 + \|W\|_F^2) \right\}$$


propensity weight

Generalization Error

- **Theoretical insights:**
 - **Additional trade-off** between bias and variance
- With probability $1 - \eta$, capacity $|\mathcal{H}|$, maximum loss Δ :

$$R(\hat{Y}^{ERM}) \leq \hat{R}_{IPS}(\hat{Y}^{ERM} | \hat{P}) + \frac{\Delta}{U \cdot I} \sum_{u,i} \left| 1 - \frac{P_{u,i}}{\hat{P}_{u,i}} \right|$$

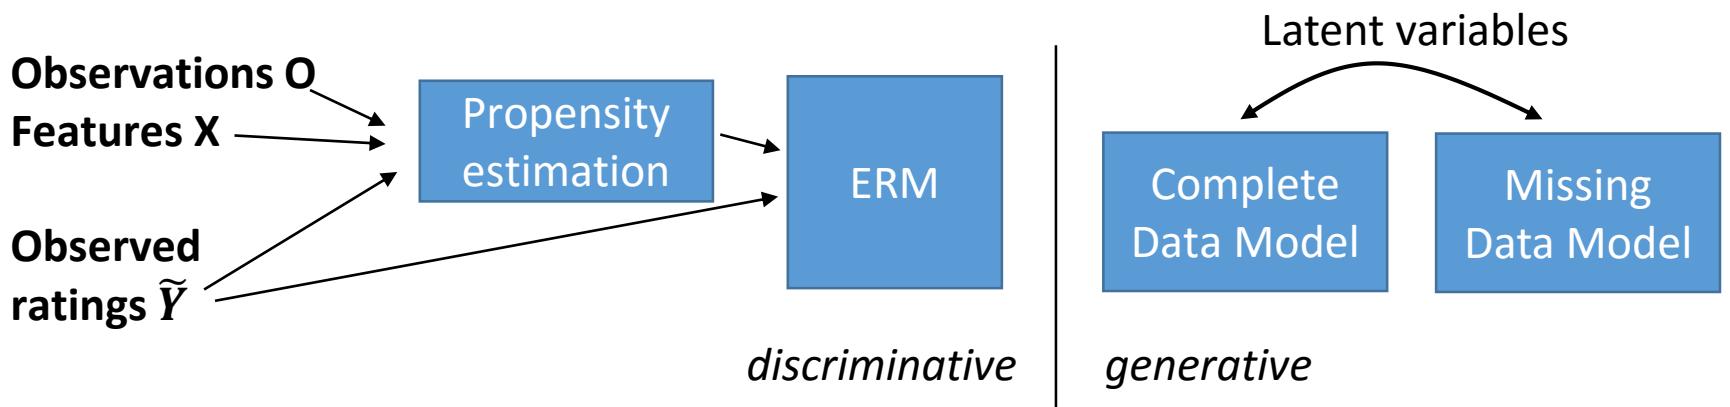
Bias

$$+ \frac{\Delta}{U \cdot I} \sqrt{\frac{\log(2|\mathcal{H}|\eta)}{2}} \sqrt{\sum_{u,i} \frac{1}{\hat{P}_{u,i}^2}}$$

Variance

Propensity-scored ERM

- Approach is **modular** and **discriminative**:
 1. Pick and estimate propensity model
 2. Use estimated propensities in ERM objective



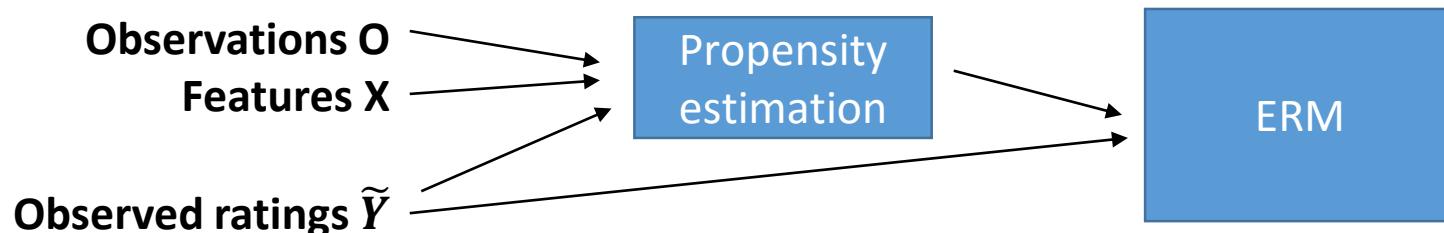
(Marlin et al., 2007; Steck, 2011;
Hernández-Lobato et al., 2014)

Debiasing Learning

- Results on two **real-world datasets**:
 - COAT: Shopping dataset (300 users; newly collected)
 - YAHOO: Song rating dataset (15400 users; Marlin & Zemel, 2009)
- Report performance on **MAR test data**:
 - HL: Latest generative approach (Hernández-Lobato et al., 2014)

	YAHOO		COAT	
	MAE	MSE	MAE	MSE
<i>MF-IPS</i>	0.810	0.989	0.860	1.093
<i>MF-Naive</i>	1.154	1.891	0.920	1.202
HL MNAR	1.177	2.175	0.884	1.214
HL MAR	1.179	2.166	0.892	1.220

Conclusions



- **Discriminative propensity scoring:**
 - Modular
 - Directly optimizes target loss
 - No latent variables
 - Scalable
- **Data and code:**
 - <http://www.cs.cornell.edu/~schnabts/mnar/>

