Correcting Exposure Bias for Link Recommendation

Shantanu Gupta, Hao Wang, Zachary Lipton, Yuyang Wang

2021

Exposure bias and link recommendation

- Link recommender systems (RS) are applied to graph structured data.
  - Nodes represent entities like papers or persons.
  - Edges represent links between nodes (e.g. citations or connections).
- They recommend other nodes that a given node should link to based on node attributes.
- The observed graph used for training can exhibit exposure bias when users are systematically underexposed to certain items.
  - For example, authors might be more likely to encounter papers from their own field and thus cite them preferentially.
- Such systems can inherit this bias and relevant low-exposure nodes may not be recommended.

Exposure bias and link recommendation

- The dataset is in the form of a directed graph $\mathcal{G}(V, E)$.
- Link probability y_{ij}: probability that v_i links to v_j.
- Propensity score π_{ij}: probability that v_i is exposed to v_j.
- Due to the exposure a_{ij}, some true positive links are observed as negative links resulting in exposure bias.

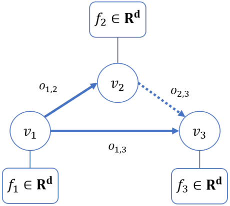

$$o'_{ij} \sim \mathrm{Ber}(y_{ij}) \quad \text{(True link)},$$
$$a_{ij} \sim \mathrm{Ber}(\pi_{ij}) \quad \text{(Exposure)},$$
$$o_{ij} = o'_{ij} a_{ij} \quad \text{(Observed link)}.$$

True Risk

- True risk is the risk of the predictions $\widehat{y}$ on the graph that would have been generated if all nodes were exposed to all other nodes:

$$R(\widehat{y}) = \mathbb{E}_{o'}\left[\frac{1}{N}\sum_{(i,j)}\delta(o'_{ij}, \widehat{y}_{ij})\right]$$

  for some loss function δ (e.g., log-loss).
- The performance of a link RS should be evaluated on its true risk.
- True risk is different from risk on the observed graph because some negative links are false negatives.

Naive estimator of true risk

- Naively estimating the risk on observed data will result in bias, i.e.,

$$\widehat{R}_{\text{naive}}(\widehat{y}) = \frac{1}{N} \sum_{(i,j)} \delta(o_{ij}, \widehat{y}_{ij})$$

  is a biased estimate of $R(\widehat{y})$.
- Thus directly evaluating a link RS on the observed graph can be a misleading measure of its performance.

Estimators for mitigating exposure bias

- We propose three estimators of the true risk — denoted by $\widehat{R}_w$, $\widehat{R}_{\text{PU}}$, and $\widehat{R}_{\text{AP}}$ — that use estimated propensity scores ($\widehat{\pi}$).
- We denote the first estimator by $\widehat{R}_w$:

$$\widehat{R}_w(\widehat{y}, \widehat{\pi}) = \frac{1}{N} \sum_{(i,j)} w_{ij} \delta(o_{ij}, \widehat{o}_{ij}), \text{ where}$$

$$w_{ij} = \frac{o_{ij}}{\widehat{\pi}_{ij}} + (1 - o_{ij})\psi_{ij}, \ \psi_{ij} = \frac{1 - \widehat{y}_{ij}}{1 - \widehat{\pi}_{ij}\widehat{y}_{ij}} \leq 1.$$

- The positive examples are up-weighted according to the inverse propensity. The negative examples are down-weighted.
- This estimator is unbiased if $\forall (i,j)$, $\widehat{\pi}_{ij} = \pi_{ij}$ and $\widehat{y}_{ij} = y_{ij}$.

Estimators for mitigating exposure bias

- The second estimator $\widehat{R}_{\text{PU}}$ is inspired by estimators from the positive-and-unlabeled setting:

$$\widehat{R}_{\text{PU}}(\widehat{y}, \widehat{\pi}) = \frac{1}{N} \sum_{(i,j)} \left[w_{ij} \delta(o_{ij}, \widehat{o}_{ij}) + w'_{ij} \delta(0, \widehat{o}_{ij}) \right],$$

$$\text{where } w_{ij} = \frac{o_{ij}}{\widehat{\pi}_{ij}} + (1 - o_{ij}),\ w'_{ij} = o_{ij} \left(1 - \frac{1}{\widehat{\pi}_{ij}} \right).$$

- We remove an appropriate number of negative examples for each positive example.

Estimators for mitigating exposure bias

- The third estimator, $\widehat{R}_{\text{AP}}$, adds positive examples for each negative example:

$$\widehat{R}_{\text{AP}}(\widehat{y}, \widehat{\pi}) = \frac{1}{N} \sum_{(i,j)} \left[w_{ij} \delta(o_{ij}, \widehat{o}_{ij}) + w'_{ij} \delta(1, \widehat{o}_{ij}) \right],$$

$$\text{where } w_{ij} = o_{ij} + (1 - o_{ij}) \psi_{ij}, w'_{ij} = (1 - o_{ij}) \tau_{ij},$$

$$\tau_{ij} = \left(\frac{\widehat{y}_{ij}(1 - \widehat{\pi}_{ij})}{1 - \widehat{\pi}_{ij} \widehat{y}_{ij}} \right).$$

- The positive examples are up-weighted according to the inverse propensity. The negative examples are down-weighted.

Comparison of the proposed and naive estimators

- We provide sufficient conditions for when the bias of the proposed estimators is lower than that of $\widehat{R}_{\text{naive}}$.
- (Informal) We show that if the $\widehat{\pi}$ are not too-underestimated and $\widehat{y}$ are not too-overestimated, the proposed estimators will have lower bias than the naive estimator.
- Thus our proposed estimators reduce bias as long as the propensities and link probabilities are learned sufficiently well.
- For all values of $\widehat{\pi}, \widehat{y}$, we have $\text{Var}(\widehat{R}_{\text{AP}}) < \text{Var}(\widehat{R}_{\text{naive}})$ and $\text{Var}(\widehat{R}_{\text{AP}}) < \text{Var}(\widehat{R}_{w}) < \text{Var}(\widehat{R}_{\text{PU}})$.

Learning propensities and link probabilities

Theorem (Generalization bound)

Let $\mathcal{F}$ be a class of functions $(\widehat{\pi}, \widehat{y})$. Let $\delta(o_{ij}, \widehat{y}_{ij}) \leq \eta \ \forall (i,j)$ and $\widehat{\pi}_{ij} \geq \epsilon > 0 \ \forall (i,j)$. Then, for $\widehat{R} \in \left\{ \widehat{R}_w, \widehat{R}_{PU}, \widehat{R}_{AP} \right\}$, with probability at least $1 - \delta$, we have

$$R(\widehat{y}) \leq \widehat{R}(\widehat{y}, \widehat{\pi}) + B(\widehat{R}) + 2\mathcal{G}(\mathcal{F}, \widehat{R}) + \mathcal{O}\left(\sqrt{\frac{2}{\delta}}\right),$$

where $\mathcal{G}$ is the Rademacher complexity and $B(\widehat{R})$ is the bias.

- The bound shows that w.h.p., if $\widehat{R}(\widehat{y}, \widehat{\pi})$ is small and the bias $B(\widehat{R})$ is small, then the true risk is also low.

Learning propensities and link probabilities

- We learn the propensities ($\widehat{\pi}$) and link probabilities ($\widehat{y}$) by minimizing the following objective:

$$l(\widehat{\pi}, \widehat{y}) = \lambda_L \mathcal{L}(o|\widehat{\pi}, \widehat{y}) + \lambda_R \widehat{R}(\widehat{\pi}, \widehat{y}), \tag{1}$$

  where $\mathcal{L}(o|\widehat{y}, \widehat{\pi})$ is the log-likelihood and $\widehat{R} \in \left\{ \widehat{R}_w, \widehat{R}_{\mathrm{PU}}, \widehat{R}_{\mathrm{AP}} \right\}$.
- The log-likelihood should ensure that the learned values are faithful to the observed data.
- The risk estimator should ensure that the true risk is small.
- We model $\widehat{\pi}$ and $\widehat{y}$ using neural networks and optimize the objective using gradient descent.

Experiments

- We use citation data from the Microsoft Academic Graph (MAG) dataset.
- It is a real-world citation dataset that contains the citation graph and paper attributes (like text and field-of-study).
- We test our methods on a semi-synthetic dataset with $42,000$ papers and two real datasets with more than 2 million and 1 million papers, respectively.

Experiments

- We use the paper title and abstract to predict the link probabilities $\widehat{y}$.
- For each paper p_i, we generate a text embedding $h_i \in \mathrm{R}^{768}$ using a pre-trained SciBERT model.
- Then $\widehat{y}_{ij}$ is modeled as a linear predictor on those embeddings:

$$\widehat{y}_{ij} = \text{Sigmoid}(\widehat{w}^\top (h_i \otimes h_j) + \widehat{b}),$$

  where $\widehat{w}$ and $\widehat{b}$ are trainable parameters.
- For simplicity, we only use the field-of-study of the papers to predict $\widehat{\pi}$.

Results on semi-synthetic data

- In the real dataset, we do not have access to true exposure values.
- So we generate a semi-synthetic dataset with real paper text but simulated exposure and outcome values.
- Our methods significantly outperform *No Propensity* (which does not correct for exposure bias) when evaluated on true links.

Table: Evaluation metrics on the test set of the semi-synthetic data computed against known ground truth citation links.

Model	Prec.	Rec.	AUC	MAP
No Propensity	67.24	54.81	84.45	41.87
MLE	81.04	60.19	93.12	56.77
$\widehat{R}_w$	**83.28**	63.73	**96.42**	56.96
$\widehat{R}_{\mathrm{PU}}$	82.16	63.07	94.28	58.01
$\widehat{R}_{\mathrm{AP}}$	83.01	**65.54**	95.38	**59.90**

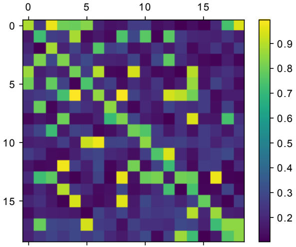

(a) Ground Truth

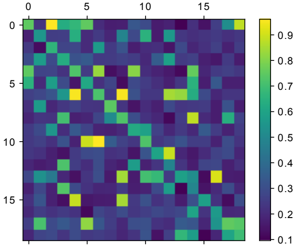

(b) Estimated

Figure: The estimated propensities propensities are close to the true simulated values when learned using $\widehat{R}_w$.

- The estimated propensities are close to the true (simulated) propensities.

Results on semi-synthetic data

- Our proposed estimators are better estimators of the true risk.

Table: RMSE of the estimated risk with respect to the true risk computed using our proposed estimators. The first column shows the risk used in the loss function in Eq. 1 to learn $\widehat{\pi}$ and $\widehat{y}$.

Trained Using	Estimator Used			
	$\widehat{R}_{\text{NAIVE}}$	$\widehat{R}_w$	$\widehat{R}_{\text{PU}}$	$\widehat{R}_{\text{AP}}$
No Prop.	1.50	-	-	-
MLE	0.67	**0.23**	0.24	0.32
$\widehat{R}_w$	0.43	**0.04**	0.10	0.11
$\widehat{R}_{\text{PU}}$	0.38	0.05	0.11	**0.04**
$\widehat{R}_{\text{AP}}$	0.41	0.06	0.08	**0.03**

Results on real data

- We test our methods on two distinct subgraphs of the MAG graph with $\tilde{2}$.4 million and 1 million papers.
- Since true exposure values are not available, we evaluate the performance against *observed* risk.
- Performance does not substantially drop even when evaluated against the observed citation graph.

Results on Dataset 1

- Performance remains comparable to *No Propensity* even when evaluated against the observed citation graph.
- The last column is a measure of diversity in the recommended papers' fields-of-study. Our methods recommend more papers from different fields.

Table: Evaluation metrics for various models computed on the test sets of a real-world citation dataset.

Model	Prec.	Rec.	AUC	MAP	FOS Entropy
Dataset 1					
No Prop.	29.45	78.30	84.44	24.10	1.65
MLE	30.24	77.84	84.41	24.60	1.73
$\widehat{R}_w$	31.46	78.02	84.74	25.60	**1.74**
$\widehat{R}_{\mathrm{PU}}$	30.98	**78.94**	**85.24**	25.11	1.73
$\widehat{R}_{\mathrm{AP}}$	**36.07**	76.08	84.67	**28.58**	1.71

Results on Dataset 2

Table: Evaluation metrics for various models computed on the test sets of a real-world citation dataset.

Model	Prec.	Rec.	AUC	MAP	FOS Entropy
Dataset 2					
No Prop.	44.86	70.85	83.22	33.19	1.06
MLE	44.43	74.66	84.97	34.39	1.08
$\widehat{R}_w$	**48.70**	71.62	83.90	**36.25**	**1.12**
$\widehat{R}_{\mathrm{PU}}$	42.17	**76.15**	**85.43**	33.26	1.08
$\widehat{R}_{\mathrm{AP}}$	47.22	71.84	83.89	35.27	1.10

Feedback loops

- We analyze the setting when a RS is repeatedly trained on data generated by users interacting with its recommendations.
- In this setting, the users are only exposed to items that are recommended and only form links with those items.
- We show that feedback loops arise which worsen exposure bias over time.
- Items with low propensity are recommended less often as time goes on.

Feedback loops

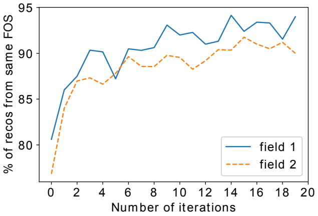

(a) No Propensity

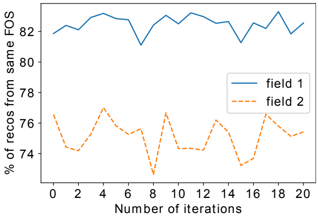

(b) With Propensity ($\widehat{R}_w$)

Figure: Feedback loops can exacerbate exposure bias.

- We run a simulation on citation data with only two fields-of-study.
- When we do not correct for exposure bias, the fraction of papers recommended from the same field increases over time (Figure (a)).
- When we correct for exposure bias, the fraction of papers recommended from the same field remains stable over time (Figure (b)).

Thank You