Towards **User-oriented** privacy for recommender system data: A **personalization**-based approach to gender **obfuscation** for user profiles

Manel Slokom

E-mail: m.slokom@tudelft.nl
Twitter: @ManelSlokom

Delft University of Technology, The Netherlands

The Sim4IR Workshop at SIGIR 2021

July 15, 2021

Framework

Items

Users

	i_1	i_2	i_3	i_4	i_5	i_6	i_7
u_1	5	0	5	0	3	0	0
u_2	4	0	3	0	5	0	1
u_3	2	5	0	4	0	0	3
u_4	5	0	4	0	0	4	0
u_5	0	0	1	4	3	0	2

Framework

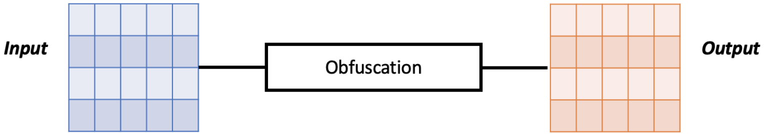

Framework

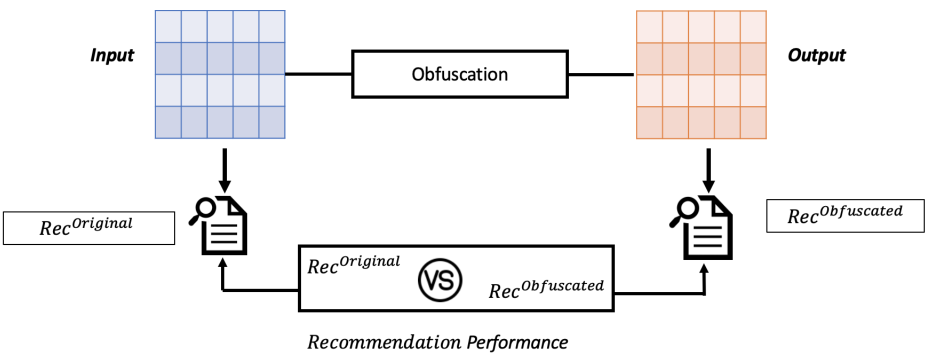

Input
Obfuscation
Output
$Rec^{Original}$
$Rec^{Obfuscated}$
$Rec^{Original}$
VS
$Rec^{Obfuscated}$
Recommendation Performance

Framework

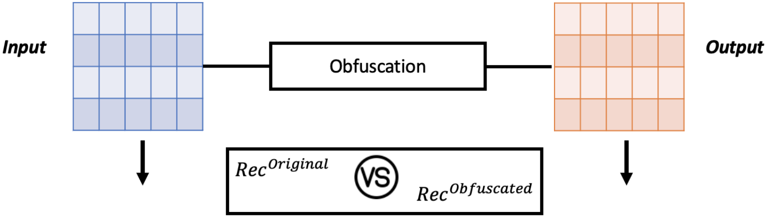

Recommendation Performance

	nDCG	***HR@10***
Original	0.1634	0.1712
PerBlur	0.1637	0.1704

In the table: we used ML1M data set. **PerBlur** is created with *addition* from the **personalized** lists of indicative items.

Framework

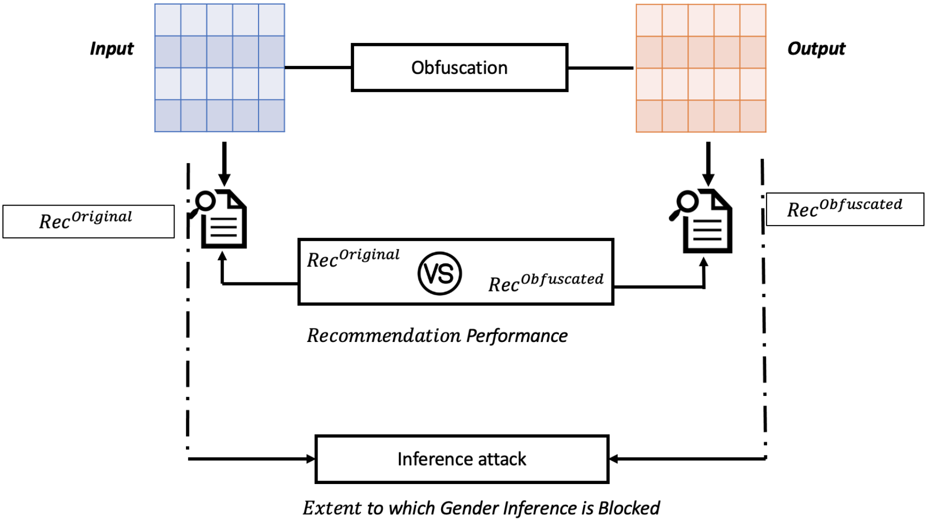

Input
Obfuscation
Output
$Rec^{Original}$
$Rec^{Obfuscated}$
$Rec^{Original}$ VS $Rec^{Obfuscated}$
Recommendation Performance
Inference attack
Extent to which Gender Inference is Blocked

Framework

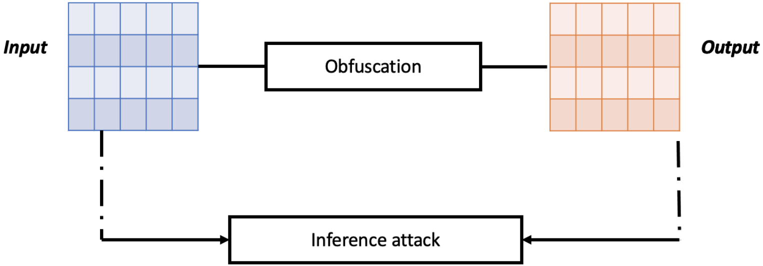

Extent to which Gender Inference is Blocked

	Extra ratings			
	0%	*1%*	*2%*	*5%*
PerBlur	0.87	0.66	**0.53**	0.26

In the table: we used ML1M data set. **PerBlur** is created with *addition* from the **personalized** lists of indicative items. Logistic regression classifier.

Take home message

- A simple, yet effective **personalized**-based approach to gender **obfuscation** for user profiles
- A recommender system trained on the obfuscated data is able to reach performance **comparable** to what is attained when trained on the original data
- A classifier can **no longer** reliably predict the gender of users
- The ability to recommend **diverse** items.

PerBlur - Personalized Blurring

- PerBlur creates the *personalized lists* of indicative items by intersecting:
  - Two lists of indicative items: L_m and L_f
  - A personalized list of items ranked in order of the probability that the user will have rated them.
- *Standard PerBlur*
  - Obfuscation by **adding** extra items from the personalized lists of indicative items
  - Level of obfuscation.
- *PerBlur with removal*
  - Similar to *Standar PerBlur* but we also **remove** certain items.

Gender inference

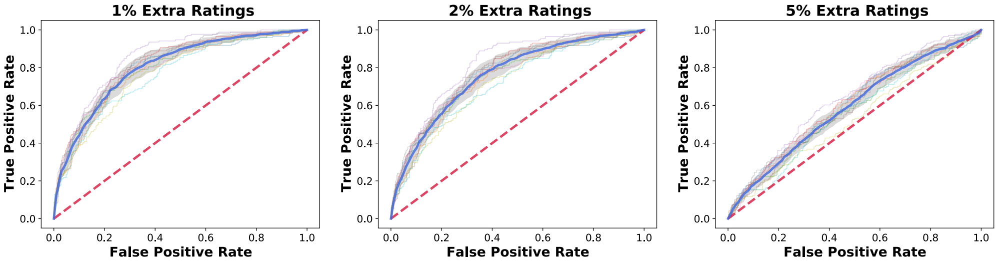

Gender inference

- Obfuscation inhibits the inference of the gender
- PerBlur requires less obfuscation
- Transferability

	Personalization	***Logistic regression***				***SVM***			
		0%	*1%*	*2%*	*5%*	*0%*	*1%*	*2%*	*5%*
BlurMe	*None*	0.87	0.76	0.69	**0.48**	0.82	0.74	0.67	**0.42**
PerBlur	*Personalized*	0.87	0.66	**0.53**	0.26	0.82	0.61	**0.46**	0.16

In the table: we report the AUC scores on ML1M data set. **BlurMe** is created with *addition* from L_m or L_f. **PerBlur** is created with *addition* from the **personalized** lists of indicative items.

Recommendation performance

	nDCG	*HR@10*
Original	0.1634	0.1712
BlurMe	0.1536	0.1633
PerBlur	0.1637	0.1704

- The recommendation performance comes close to what is achieved on original data
- PerBlur, thanks to its **personalization**, approaches the original performance more closely and more consistently than BlurMe

In the table: we used BPRMF algorithm on ML1M data set. **BlurMe** is created with *addition* from L_m or L_f. **PerBlur** is created with *addition* from the **personalized** lists of indicative items.

Achieving diverse results

- The proportion of correctly recommended items that are stereotypical for gender
- Three different cutoff levels (10, 20, 50)

	Obfuscation Strategy		**Gender-steretypical items**					
	Personalization	*Removal*	*top10F*	*top10M*	*top20F*	*top20M*	*top50F*	*top50M*
Original	*None*	*None*	0.0020	0.0045	0.0038	0.0069	0.0082	0.0128
PerBlur	*Personalized*	*Greedy*	**0.0003**	**0.0005**	**0.0014**	**0.0020**	**0.0051**	**0.0073**

- PerBlur is effective in lowering the proportion of TopN gender-steretypical items

In the table: we used ML1M data set. **PerBlur** is created with *addition* from the **personalized** lists of indicative items and *removal* from L_m or L_f.

Outlook and future work

1. Data obfuscation for recommender systems
2. Step toward **data sharing** without privacy concerns
3. From privacy to **fairness** and **diversity**
4. From **partially** to **fully** synthetic data

Thank You

References

Manel Slokom, Martha Larson and Alan Hanjalic (2021)
Towards User-Oriented Privacy for Recommender System Data: A Personalization-based Approach to Gender Obfuscation for User Profiles.
Under review.

Manel Slokom, Martha Larson, and Alan Hanjalic (2019)
Data Masking for Recommender Systems: Prediction Performance and Rating Hiding.
In: Late-Breaking Results RecSys'19. pp. 21-25.

Christopher Strucks, Manel Slokom, and Martha Larson (2019)
BlurM(or)e: Revisiting gender obfuscation in the user-item matrix.
Recommendation in Multistakeholder Environments in conjunction with the 13th ACM Conference on Recommender Systems (RecSys'19).

Manel Slokom (2018)
Comparing recommender systems using synthetic data
Doctoral consortium in conjunction with the 12th ACM Conference on Recommender Systems (RecSys '18). pp. 548–552.