The trade-off between **data minimization and fairness** in collaborative filtering

Nasim Sonboli, Sipei Li, Mehdi Elahi, Asia Biega

The 7th FAccTRec Workshop: Responsible Recommendation
In Conjunction with the 18th ACM Conference on Recommender Systems

BROWN
Data Science Institute

CENTER FOR TECHNOLOGICAL RESPONSIBILITY,
RE-IMAGINATION, AND REDESIGN

General Data Protection Regulation (GDPR)

- US State-Level Laws:
  - California Consumer Privacy Act (CCPA) (and CPRA) effective in 2020
  - Virginia Consumer Data Protection Act (CDPA) effective in 2023
  - The Colorado Privacy Act (CPA), which will be fully enforced in 2025, etc.
- Sectoral laws:
  - Health Insurance Portability and Accountability Act (HIPAA), Gramm-Leach-Bliley Act (GLBA), Children's Online Privacy Protection Act (COPPA), Federal Trade Commission Act (FTC Act)

Is it Possible to comply by GDPR Regulations simultaneously?

- What is the relationship between
  - Fairness & Accuracy
  - Data Minimization & Accuracy
  - **Data minimization & Fairness**

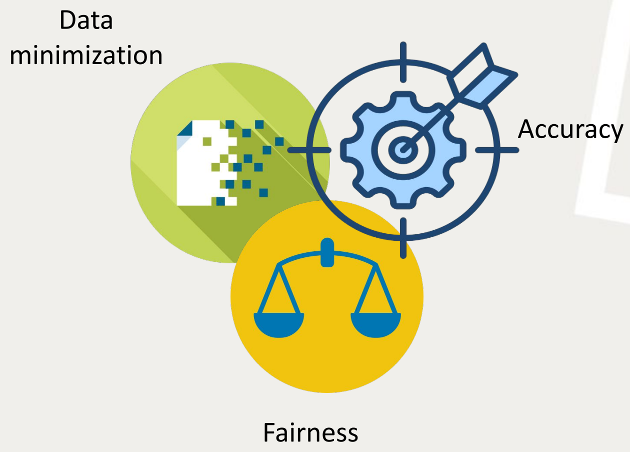

Data Minimization

& related principles

Data minimization — personal data must be "adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed".

Purpose limitation — Process data for the legitimate purposes specified explicitly to the data subject when you collected it.

Fairness — Data & Processing data must be fair to the data subject.

Goal

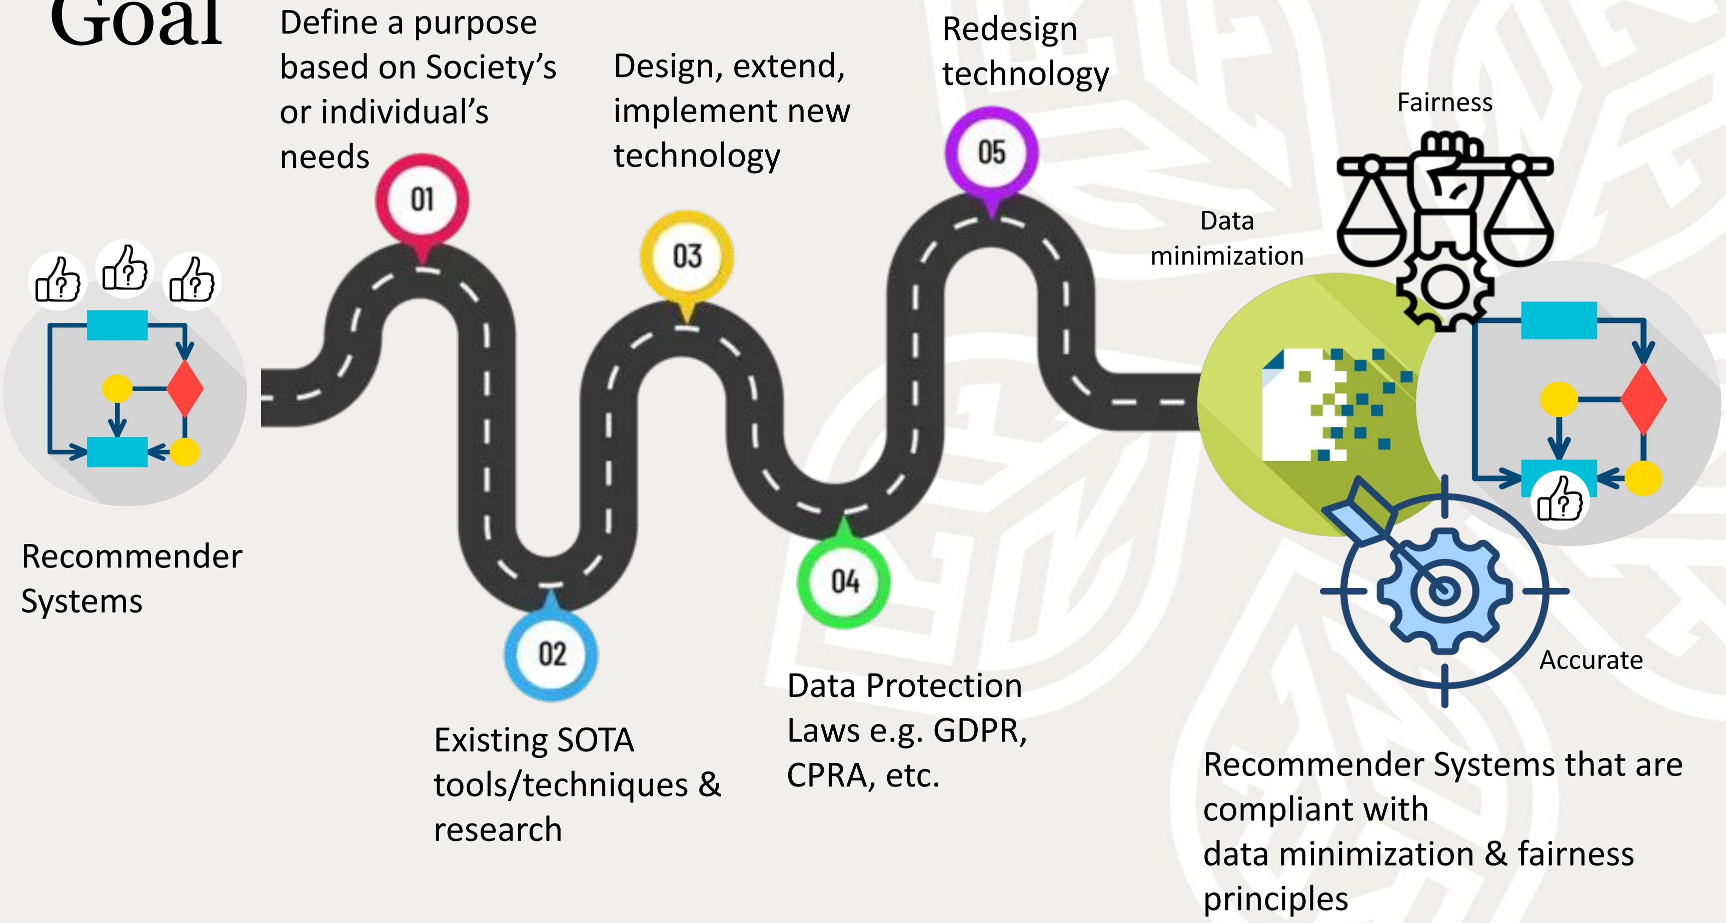
Define a purpose based on Society’s or individual’s needs
Design, extend, implement new technology
Redesign technology
Fairness
Data minimization
01
03
05
Recommender Systems
02
04
Existing SOTA tools/techniques & research
Data Protection Laws e.g. GDPR, CPRA, etc.
Accurate
Recommender Systems that are compliant with data minimization & fairness principles

Clarifying Data Minimization Definition

Data minimization *— personal data must be "<u>adequate, relevant</u> and limited to what is necessary in <u>relation to the purposes</u> for which they are processed".*

The core requirements of data minimization:

- **Adequacy ~ Amount**
- **Relevance ~ Quality**
- **Purpose-limited** (adequacy & relevance should be defined w.r.t purpose)
  - Purpose: personalization
  - E.g. certain amount of quality data is required for Recsys to improve its performance.
  - Lack of data prevents the system from completing its task as promised.

Goal: To find a balance between minimizing the amount of data (adequacy) and increasing (or maintaining) the performance of a recsys model (relevance).

Can we minimize & **learn accurately**? (Lit. Rev.)

Biega et al. is the 1st to study and demonstrate empirically the feasibility of integration of data minimization in recommender systems.

Shanmugan et al. uses the algorithm's performance curve for automatically determining and enforcing accurate stopping criteria for the data collection during training.

Clavell et al. using a qualitative methodology, investigate the tension between data minimization, performance, and fairness. They show it's possible to maintain accuracy while adhering to the GDPR data minimization.

Adequacy
(amount of data)

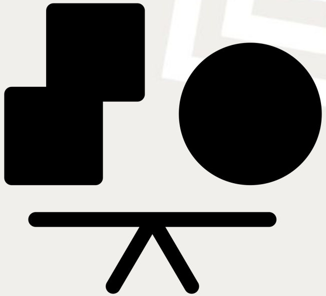

Relevance
(accuracy)

w.r.t purpose

We can minimize but… (Lit. Rev.)

Biega et al. is the 1st to study and demonstrate empirically the feasibility of integration of data minimization in recommender systems.

- **DM impacts individuals differently, potentially harming under-represented groups with higher accuracy losses.**

Shanmugan et al. uses the algorithm’s performance curve for automatically determining and enforcing accurate stopping criteria for the data collection during training.

- **Accumulating more data doesn’t always increase the per-user accuracy. If the collected data is not representative or is disparate, the data collection can hurt user performance**

Clavell et al. using a qualitative methodology, investigate the tension between data minimization, performance, and fairness. They show it’s possible to maintain accuracy while adhering to the GDPR data minimization.

- **Collecting personal information becomes essential if its absence results in inaccuracies, or unfairness, or if the data is required for auditing and accountability purposes. So, data minimization should not be applied unless other legal principles of GDPR such as fairness are considered.**

Data Minimization might cause unfairness!

- How to implement Data Minimization?
- How to measure unfairness?

How to minimize & learn accurately?
(using existing tools/techniques)

Proposed method:

Active Learning Methods

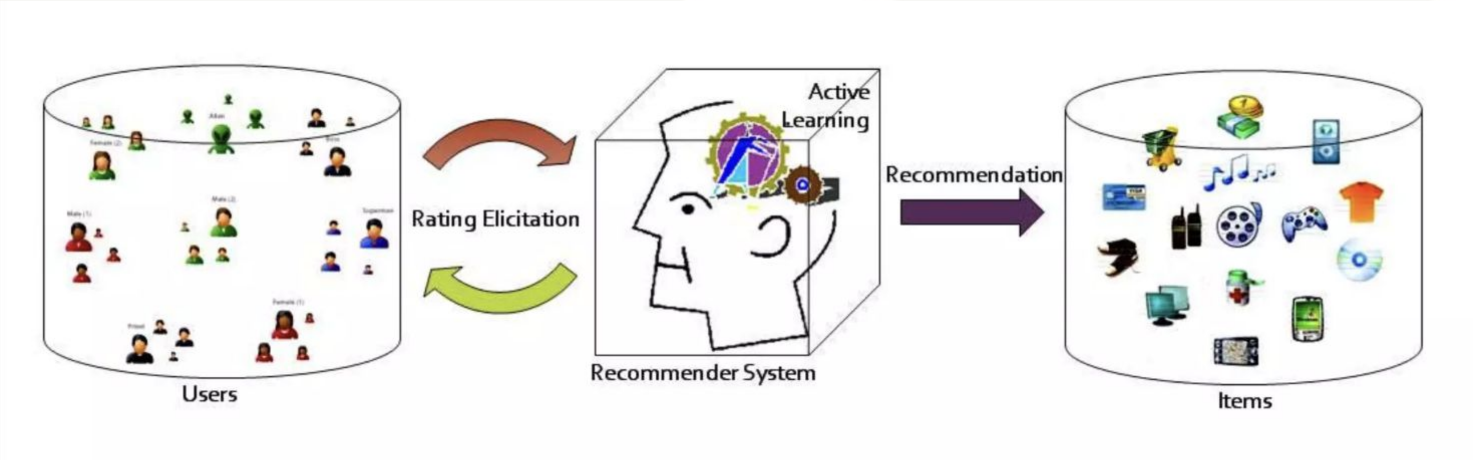

Active Learning Strategies

- Unpersonalized
  - Variance (uncertainty reduction)
  - Greedy Extend (Error-reduction)
  - Popularity (attention-based)
  - Popularity*Variance (hybrid)
  - Random
- Personalized
  - MaxRating
  - MinRating
  - MixedRating
  - KNN (neighborhood-based)
  - Random

Algorithm 1 The testing algorithm for a strategy S

Require: Dataset R, strategy S, base recommendation model M, gender mapping G, query size q

Ensure: RMSE of female users after each query $RMSE_f$, RMSE of male users after each query

$X, T \leftarrow$ userfixed_split(R).
$K \leftarrow$ randomly sampled 0.2% of X.
$X \leftarrow X \setminus K$.
$T_f, T_m \leftarrow T$ partitioned based on G.
for each user u **do**
  $I_u \leftarrow \{i | k_{ui} = NULL\}$.
end for
$RMSE_f \leftarrow$ empty list.
$RMSE_m \leftarrow$ empty list.
while $\exists I_u \neq \emptyset$ **do**
  for each user u s.t. $I_u \neq \emptyset$ **do**
    $L \leftarrow S(u, q, K, I_u)$
    $L_e \leftarrow \{i \in L | x_{ui} \neq NULL\}$.
    for $i \in L_e$ **do**
      $k_{ui} \leftarrow x_{ui}$.
      $X \leftarrow X \setminus x_{ui}$
    end for
    $I_u \leftarrow I_u \setminus L$.
  end for
  Train M on K.
  $RMSE_f.append(RMSE(T_f, M(T_f)))$
  $RMSE_m.append(RMSE(T_m, M(T_m)))$
end while

Experimental Setup

- Dataset: MovieLens-1M, (and ML-100k)
  - 5-core (6,040 users, 3,377 items, density of %8)
  - 5 fold cross validation
  - 80% train and 20% test (userfixed technique)
- Recommendation glgorithm: SVD (Surprise library)
  - 100 latent factors, a learning rate of 0.005, and regularization term of 0.1
- Evaluation metrics: Root Mean Squared Error (RMSE) @10
- Protected group: women
  - the protected group due to their lower count and smaller profile sizes (4,331 men and 1,709 women)

Results (MovieLens-1M)

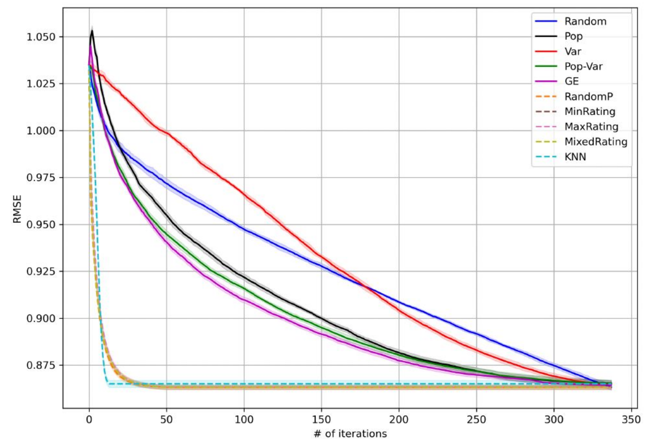

Figure 1: Experiment#1: RMSE trend of personalized & non-personalized active learning strategies in MovieLens dataset over 340 iterations, with a sliding window of 10 items

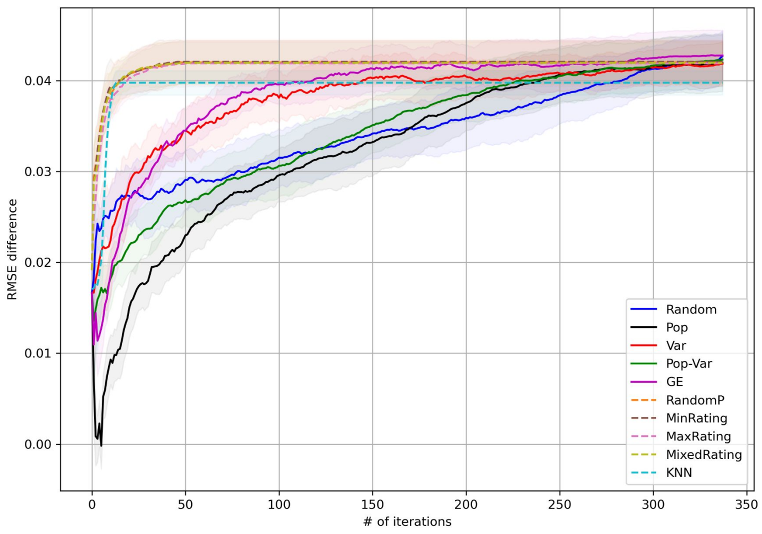

Active learning strategies behave differently and affect the accuracy and data collection differently. They affect the RMSE of the protected & unprotected groups differently.

Is RMSE difference because of data imbalance?

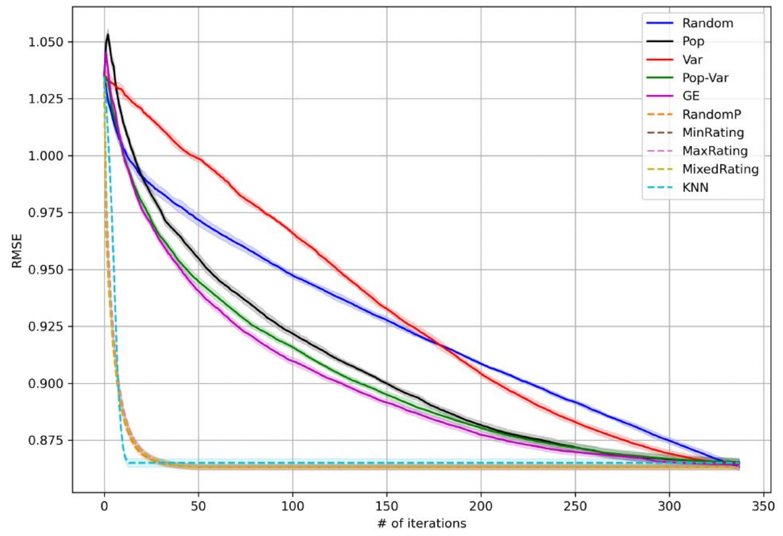

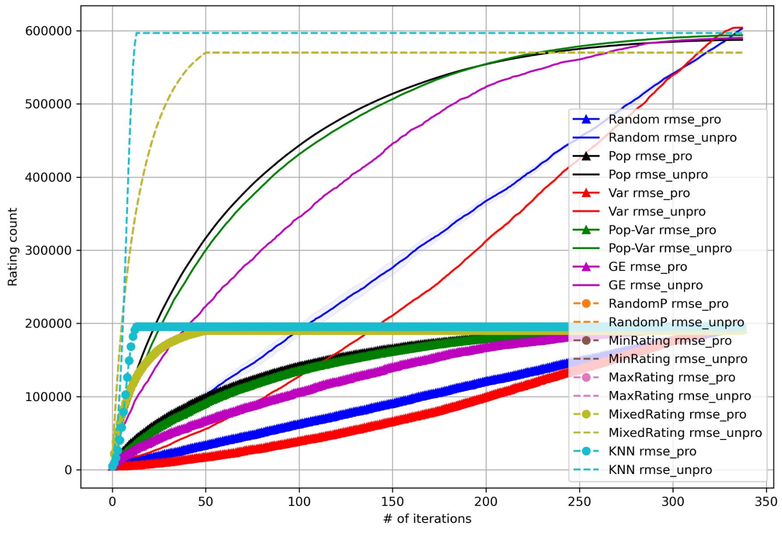

AL data collection for pro & unpro groups is different. This can lead to unfairness.

Experiment #2: Balancing the pro/unpro Ratio

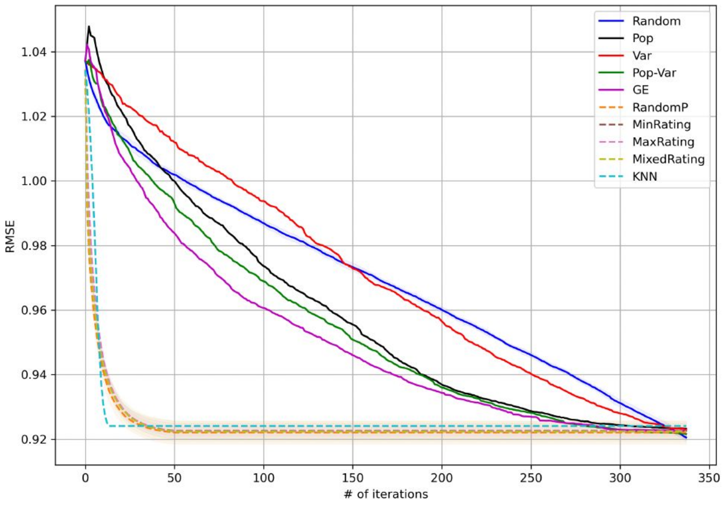

Figure 2: Experiment#2: RMSE trend of personalized & non-personalized active learning strategies in MovieLens dataset over 340 iterations, with a sliding window of 10 items

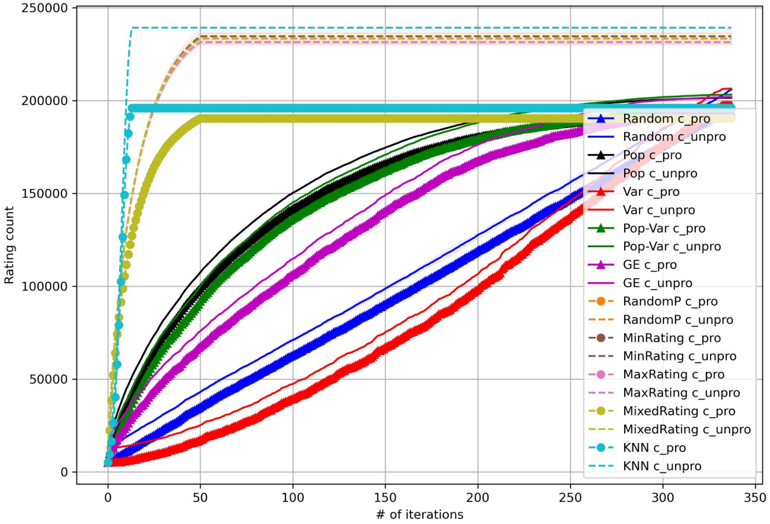

Experiment #2: Balancing the pro/unpro Ratio

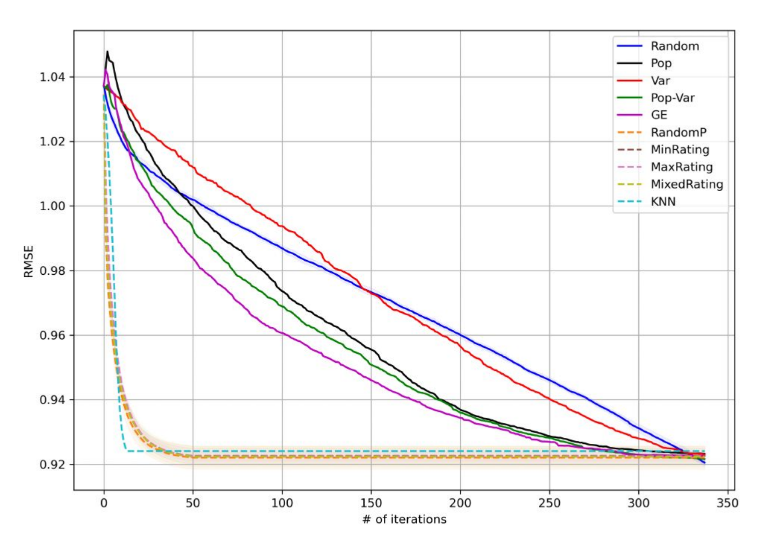

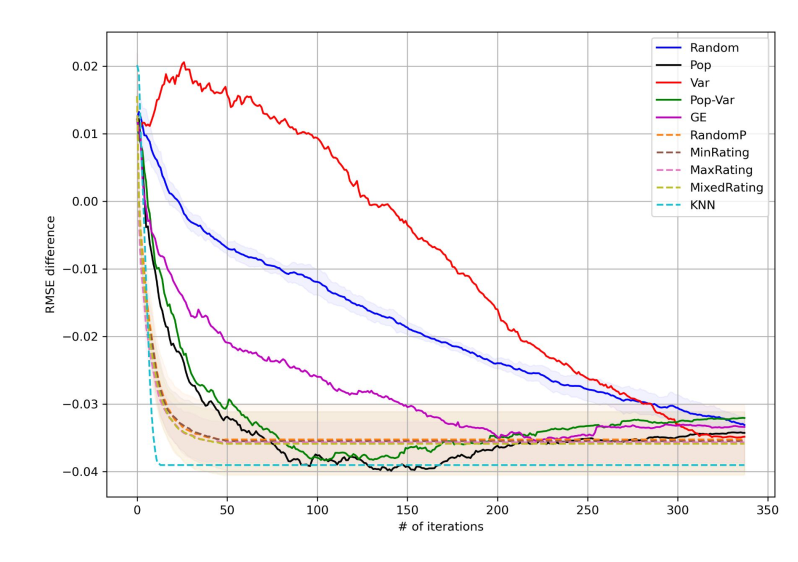

The RMSE of unprotected is worse then the protected group sometimes!

It’s **not** about the quantity!

The amount of collected data matters, but the quality of data matters more!

Movielens-100k

Experiment #1

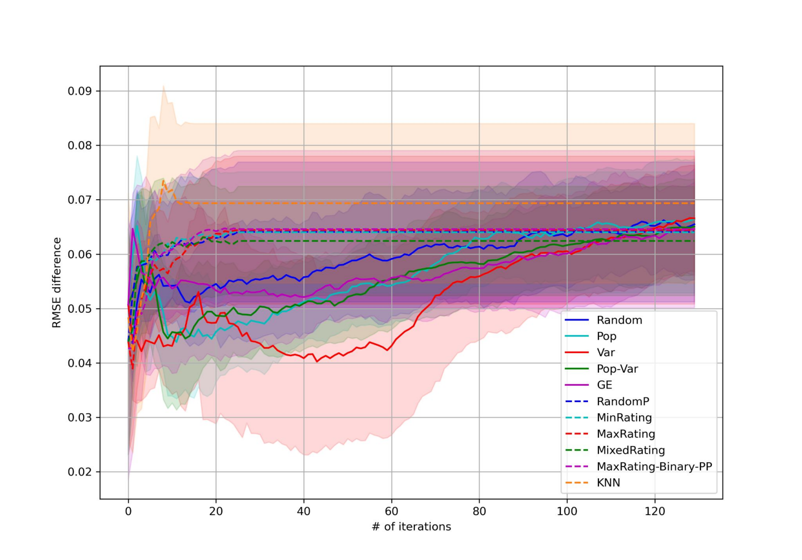

Experiment #2

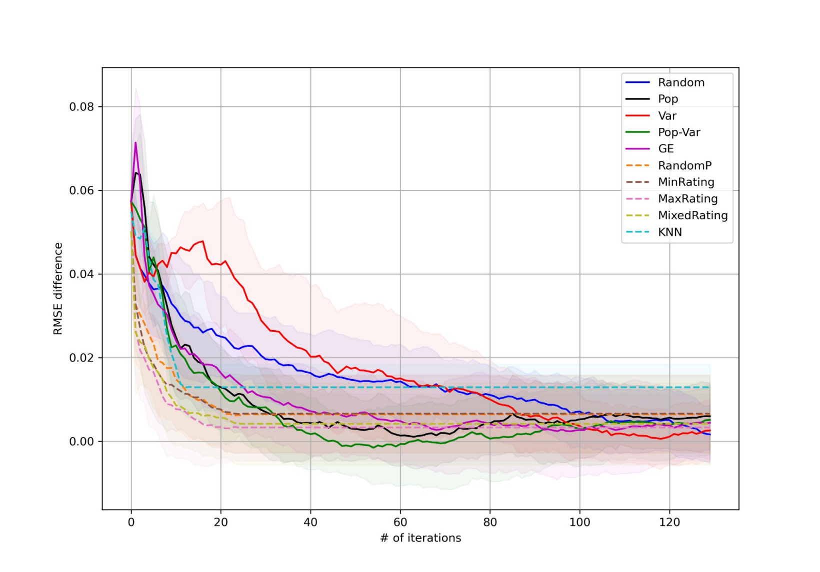

Limitations & Future Direction

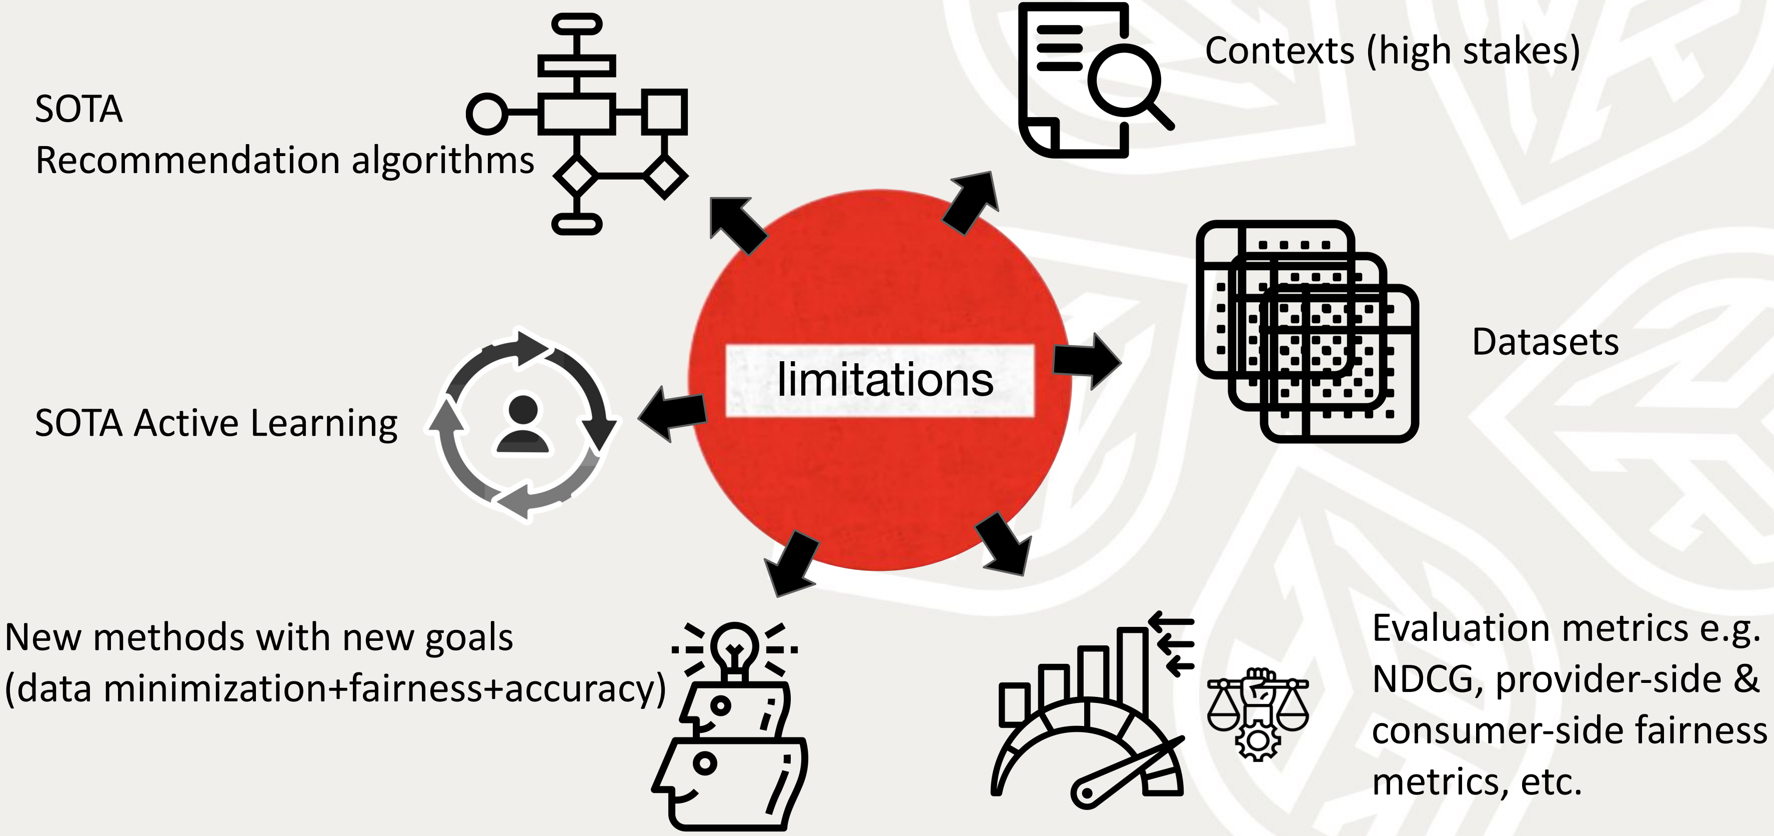

Takeaways

- Don’t interpret data minimization litarally. Its goal is to limit the data collection to the pre-specified purpose, avoid over-collection of data, or collection of irrelevant data.
- Active Learning strategies are one way of operationalizing Data Minimization.
- Data Minimization via active learning widens the RMSE gap between the protected & unprotected user groups, could lead to unfairness.

  (Any method that samples and minimizes data is prune to the issue of unfairness due to data imbalance)

- To design GDPR-compliant algorithms considering only one principle is not enough. One must consider the trade-offs of each principle with other principles. (e.g. fairness, accuracy, data minimization)
- Better data representation sometimes helps with the accuracy gap, however, the contributed information matters besides the amount of data. (adequate relevant data)

Looking into the Future

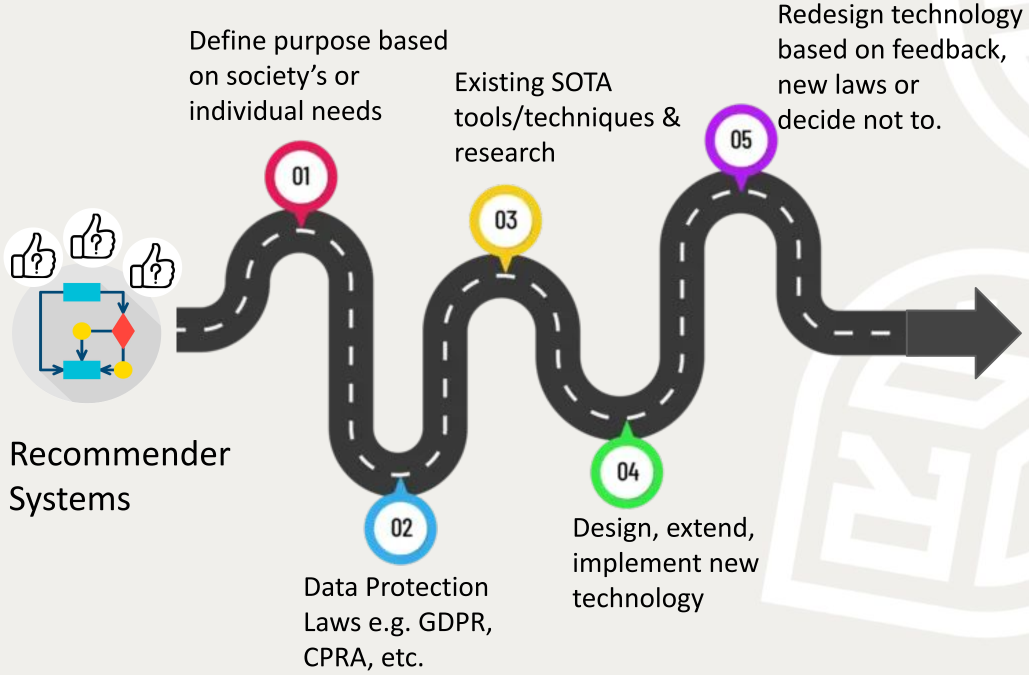

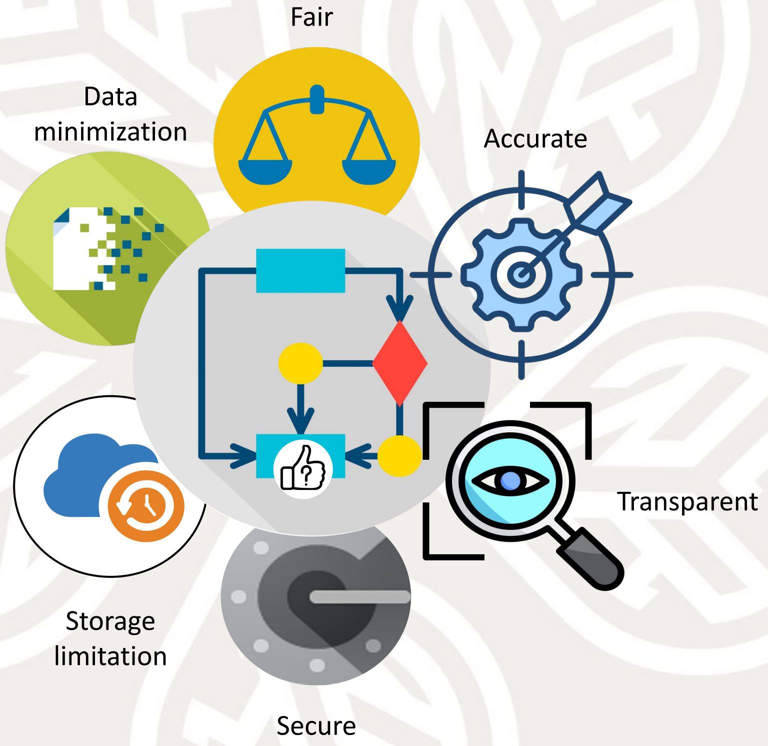

GDPR compliant Recommender Systems

Thank you for listening!

- Reviewers and organizers of the FAccTRec workshop
- My collaborators: Sipei Li, Mehdi Elahi, Asia Biega
- My postdoc PI Prof. Suresh Venkatasubramanian for his constructive feedback.
- My previous PI, phd advisor, Prof. Robin Burke for his constructive feedback.
- NSF for funding this project.

Paper link

Be happy to connect on LinkedIn

CCC Computing Community Consortium Catalyst
CRA Computing Research Association
NSF