



Recommender Systems



- Systems that suggest **content** to **users**
- Get user feedback to **learn preferences**

Users are strategic

Preferences are **not fixed**. Users change. Some adapt to get **what they want** from platforms.

- Not clicking a link to avoid related ads
- Browsing content in incognito mode
- Trying not to "ruin" the algorithm


All of these break classical assumptions.

Key issue: (Dis-)trust



Making things formal

Instead of fixed preferences, we model recommendation as **alternating two-player game**.



Platform

Actions: recommenders $f \in \mathcal{F}$
Utility: $U_p(f, b)$ (e.g., engagement)



User

Actions: behavior $b \in \mathcal{B}$
Utility: $U_u(f, b)$ (e.g., enjoyment)

Strategy s maps **history** H to **action**.


In response to platform strategy, users can be:

Optimal: $s = \arg \max_{s_u} \sum_{\tau=1}^T U_u(s_p(H_\tau), s_u(H_\tau))$


Greedy (truthful): $s(H) = \arg \max_b U_u(s_p(H_\tau), b)$

A user **trusts** the platform is their optimal strategy is to be truthful!


Ex: Privacy-Conscious User



Item x



User features θ



$$\theta^T x + \mathcal{N}(0, 1)$$

User interest y

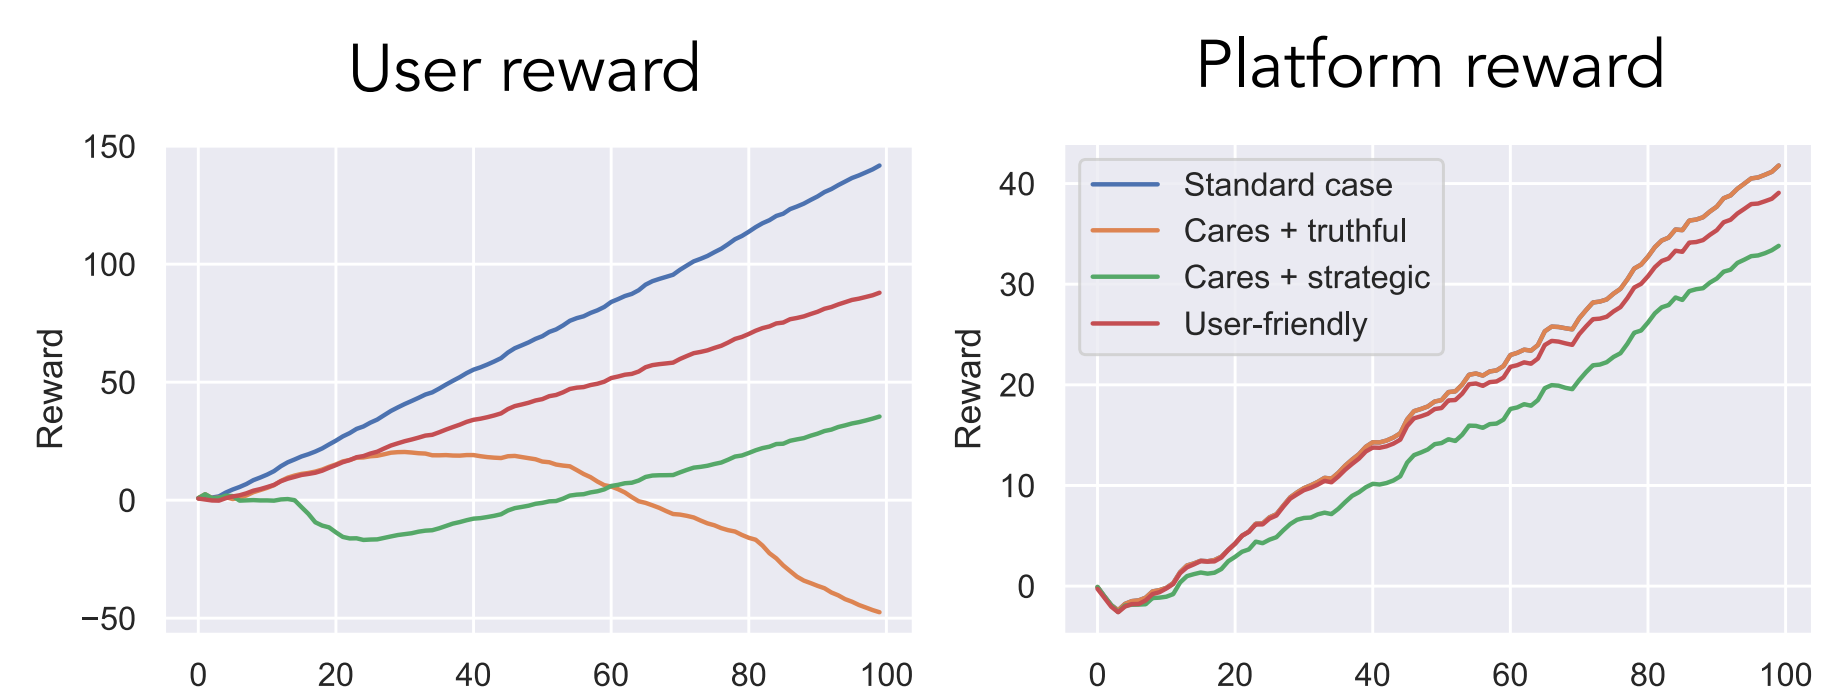
Non-strategic user:

Truthfully reports interest y


Strategic (privacy-conscious) user:

Reports interest while *hiding* a feature θ_i


- Regret-minimizing platform works for non-strategic user but is **misled** by strategic user.
- Both **user** and **platform** outcomes improve when the platform cooperates & builds trust.



Ex: Bimodal Preferences



User



One of two moods

Prefers horror movies

Prefers comedies

Non-strategic user:

Truthfully reports current interest θ_t .

Strategic user:

Stops reporting when in less common mood.

- If platform assumes user is **unimodal**, will get **feedback loops** or learn **average** of moods.
- Even if platforms is bimodal, have to **guess** user's mood to serve correct content.
- Instead, platform should **build trust** by **soliciting** the user's mood.

Takeaways

- Users of modern recommendation systems are **strategic** and have their own objectives.
- Can lead to **feedback loops** and **suboptimality**.
- We should develop systems with trust in mind!