

Long-Term User Value

The concept of long-term value to users encapsulates the objective of aligning recommender system design to outcomes that prioritize users' deliberative, forward-looking aspirations or preferences – in the “System 2” thinking sense – while remaining generic enough to allow different platforms to design their own way of achieving it.

For example, platforms that are optimized to support long-term value to users:

- may ask users directly to state their explicit preferences;
- may rely on surveys, quality indicators selected by the user, or predictions of each;
- may rely on signals that are deliberative, clear, or onerous; or
- may combine aspects of these or other approaches.

On many platforms, the most accurate data about long-term user value comes from users directly expressing preferences through user controls and settings. This should be viewed as the most robust approach to understanding long-term user value, where this data is available.

However, user controls are typically only adopted by a fraction of users, and for some recommender systems it may not be workable for platforms to offer user controls. In the absence of explicit preference data from users, the list of accounts that users choose to follow or subscribe to (on platforms that support such functionality) could be considered an indicator of long-term preference. Some platforms survey users to identify long-term value, and while survey populations can be representative of the user base, they typically engage a small subset of users.

If explicit preferences or survey responses are not available, approaches that predict or extrapolate from user preference or survey data to the rest of the user population may be a reasonable alternative. Unlike optimizing for short-term predicted engagement, these approaches are based on explicit information supplied by users, albeit not all users. Platforms may also consider adopting deliberative processes where small subsets of users engage in intensive processes to identify how recommender systems can be optimized for long-term value, and applying those results across the user base.

One way of knowing when a platform's design fails to serve its users' long-term preferences is when meaningful numbers of users regret their experiences on the platform or report a loss of self-control. External surveys of user regret can serve as a means to validate that platforms are indeed optimizing for long-term preference. These outcomes indicate design patterns that undercut what users want from their experience on a platform over the long run.

Optimizing recommender systems to maximize predicted short-term engagement does *not* typically promote long-term value. The mere fact that recommended items succeed in engaging users is not sufficient to establish that these recommendations align with long-term value to users. Rather, predictions of long-term value must be supported by evidence of explicit, expressed desires held by individual users or representative subsets of users, not ambiguous behaviors that may correlate with inferred “preferences.”