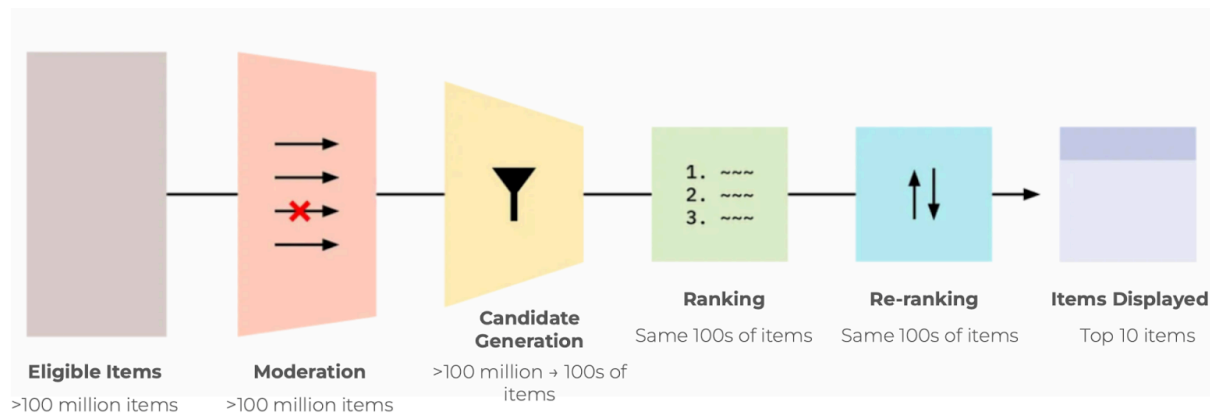


# Recommender Systems 101

The process of recommending content proceeds from mapping the universe of items a user could plausibly be interested in to identifying and finally recommending the items most advantageous to a platform's goals. This process usually occurs in four stages:

1. **Moderation:** The platform applies its moderation policies to the universe of available items, removing items from the pool that violate the platform's policies.
2. **Candidate generation:** The platform selects high-potential items as candidates from among the universe of available items. This universe can be very large, sometimes on the order of billions of items, so selecting candidate items is usually a very lightweight computational process that does not involve in-depth analysis.
3. **Ranking:** Each candidate item is then assigned a numeric score intended to capture the value of showing it to a particular user in a particular context. This score determines the order in which candidate items are output from the ranking stage.
4. **Re-ranking:** Finally, the order of candidate items is changed according to other ancillary goals, for example, to avoid repetitiveness in terms of the content type or source. These ancillary goals may be considered important because their absence undermines long-term engagement – for example, less variety may be more engaging in the short term, but may diminish the user's experience on a platform over the long-term.



*A typical recommender system pipeline, along with the approximate number of items retained at each stage for a large platform.*