International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 5, Sep-Oct-2026, ISSN (Online): 2395-566X

How Recommendation Systems Work At Netflix and Youtube: A Review of Candidate Generation, Ranking, Multitask Objectives and Responsible Personalization

Tambe Siddhi Mahesh[1] Prof. S. S. Bhor[2] Prof. M. C. Kolse[3] Prof. S. K. Said[4]
Department Of Artificial Intelligence And Data Science Jaihind College Of Engineering Pune, Maharashtra, India

Abstract—Streaming platforms such as Netflix and YouTube offer catalogues far larger than any user can browse, and most of what people watch is chosen from personalized recommendations. This paper explains how these industrial recommendation systems work, based on published research from the platforms and the wider community. It traces the evolution from the Netflix Prize and matrix factorization to deep learning, and describes the typical two-stage pipeline: candidate generation, which retrieves a few hundred relevant items from millions using embeddings learned from watch histories, and ranking, which scores candidates with richer features to predict clicks, watch time and satisfaction. It reviews multitask ranking that balances engagement and satisfaction objectives and corrects position bias, sampling-bias correction in two-tower retrieval models, reinforcement learning that optimizes long-term user value, and the personalization of rows, artwork and the home page at Netflix. Evaluation through offline metrics and online A/B testing is discussed, as are concerns about filter bubbles, calibration of recommendations to user interests and audits of problematic content pathways. A five-layer architecture is presented with research gaps, objectives, a comparison of techniques and challenges. The review concludes that modern recommenders are complex, multi-objective systems in which responsible design is as important as accuracy.

Index Terms—Recommendation systems, Netflix, YouTube, candidate generation, ranking, matrix factorization, deep learning, A/B testing

I. INTRODUCTION

1.1 Motivation

When a user opens Netflix or YouTube, the home page is assembled specifically for them within a fraction of a second. Gomez-Uribe and Hunt [1] reported that most hours streamed on Netflix come from recommendations and estimated that personalization saves the company a large sum each year by reducing cancellations. On YouTube, recommendations drive a large share of watch time [2].

These systems face a scale problem (billions of users and items), a freshness problem (new videos every minute) and an objective problem (what users click is not always what they value). This paper explains the main components and techniques used to address these problems.

1.2 Objectives Of The Review

- To explain the evolution of recommenders from the Netflix Prize to deep learning.
- To describe the candidate-generation and ranking pipeline.
- To review multitask, bias-corrected and reinforcement-learning recommenders.
- To present a five-layer streaming recommendation architecture.
- To discuss evaluation and responsible personalization.

II. OVERVIEW OF KEY CONCEPTS

2.1 Families Of Recommendation Methods

Fig. 1 organizes recommendation techniques into collaborative filtering, content-based, deep learning and reinforcement learning approaches.

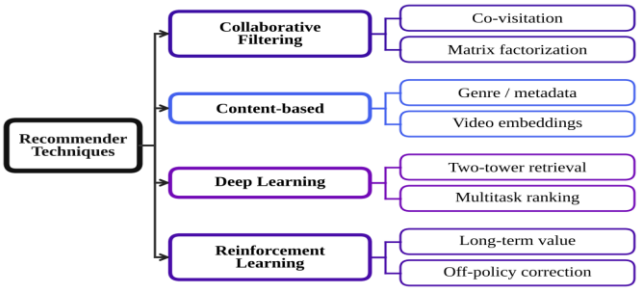


Fig. 1. Families of recommendation techniques used by streaming platforms.

© 2026 IJSRET

2.2 Implicit Feedback

Streaming platforms rely mainly on implicit signals such as plays, watch time, skips, likes and searches, rather than explicit ratings. These signals are abundant but noisy and biased by what was shown.

2.3 Review Of Existing Work

Netflix. The Netflix Prize [3] invited researchers to improve rating prediction by 10 percent and popularized matrix factorization [4], which represents users and items as latent vectors. Gomez-Uribe and Hunt [1] described the modern Netflix system as a set of algorithms for personalized video ranking, top-N rows, trending, continue watching and similar titles, evaluated by A/B tests on retention. Steck [5], also from Netflix, proposed calibrated recommendations so that the mix of genres matches each user's interests rather than being dominated by their main interest.

YouTube. Davidson *et al.* [2] described an early YouTube system based on co-visitation of videos. Covington *et al.* [6] introduced a deep two-stage architecture with a candidate-generation network and a ranking network trained to predict expected watch time. Yi *et al.* [7] proposed sampling-bias correction for two-tower retrieval models trained on streaming data. Zhao *et al.* [8] built a multitask ranking model with a mixture of experts for engagement and satisfaction objectives and a shallow tower to reduce position bias. Chen *et al.* [9] applied off-policy reinforcement learning with REINFORCE to optimize long-term user engagement.

Responsibility. Ribeiro *et al.* [10] audited YouTube and found pathways by which users moved toward more extreme content, highlighting the need to consider societal effects. Table I summarizes these works.

TABLE I Summary of Literature Survey

Sr. No.	Ref.	Year	Methodology	Findings	Research Gap
1	[3]	2007	Netflix Prize	Spurred CF research	Rating prediction only
2	[4]	2009	Matrix factorization	Latent factor models	Cold start
3	[2]	2010	Co-visitation at YouTube	Early scalable system	Shallow signals
4	[1]	2015	Netflix system overview	Business value of personalization	Descriptive
5	[6]	2016	Deep two-stage YouTube model	Watch-time ranking	Engagement-focused
6	[5]	2018	Calibrated recommendations	Balanced interest mix	Extra re-ranking
7	[7]	2019	Bias-corrected two-tower	Better streaming retrieval	Frequency estimation
8	[8]	2019	Multitask ranking	Engagement + satisfaction	Complex tuning
9	[10]	2020	Radicalization audit	Problematic pathways	Observational

2.4 Identified Research Gaps

Engagement objectives may not reflect long-term satisfaction or wellbeing.

Feedback loops reinforce popular items and narrow exposure.

New items and new users suffer from cold start.

Offline metrics often disagree with online A/B results.

Public understanding of how recommendations are made is limited.

2.5 System Architecture

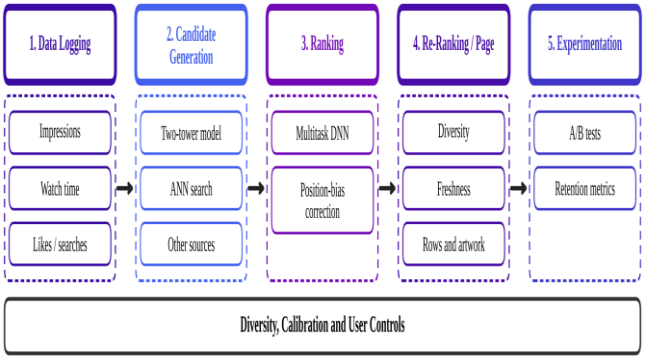


Fig. 2 shows a reference architecture.

© 2026 IJSRET

Fig. 2. Reference five-layer streaming recommendation architecture.

2.6 Data Logging Layer

Impressions, plays, watch time, likes, searches and device context are logged with the position at which each item was shown.

2.7 Candidate Generation Layer

A two-tower model embeds users and items; approximate nearest-neighbour search retrieves hundreds of candidates from millions, combined with sources such as subscriptions and trending items.

2.8 Ranking Layer

A multitask deep model predicts click, watch time, like and satisfaction; scores are combined into a single ranking value.

2.9 Re-Ranking And Page Construction Layer

Business rules, diversity, freshness and calibration adjust the list; rows and artwork are personalized for the home page.

2.10 Experimentation Layer

Changes are tested with online A/B experiments measuring long-term engagement and retention.

III. IMPLEMENTATION METHODOLOGY

Fig. 3 describes the serving pipeline. In matrix factorization, the predicted preference of user u for item i is

$$\hat{r}_{ui} = \mu + b_u + b_i + p_u^T q_i \quad (1)$$

In two-tower retrieval, candidates maximize the dot product of user and item embeddings. The ranking score combines task predictions,

$$score = \Sigma_k \, w_k \cdot \hat{y}_k \quad (2)$$

where $\hat{y}_k$ are predicted engagement and satisfaction signals and w_k their weights.

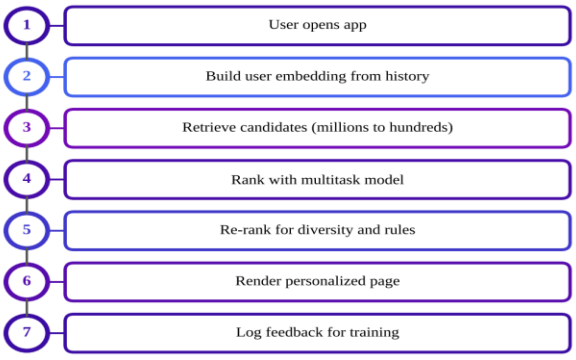


Fig. 3. Recommendation serving pipeline.

IV. PERFORMANCE COMPARISON

Table II compares recommendation techniques by scalability, cold-start handling and objective.

Technique	Scalability	Cold Start	Objective
Co-visitation CF	High	Poor	Relatedness
Matrix factorization	High	Poor	Rating / play
Content-based	High	Good	Similarity
Deep two-stage	Very high	Medium	Watch time
RL recommender	High	Medium	Long-term value

TABLE II Comparison of Recommendation Techniques

V. OPEN ISSUES AND FUTURE DIRECTIONS

Table III summarizes challenges. Future directions include generative and LLM-based recommenders, user controls over recommendations, wellbeing-aware objectives and transparent explanations.

TABLE III Challenges and Mitigation

Challenge	Mitigation
Filter bubbles	Diversity and calibration
Position and selection bias	Bias-corrected training
Cold start	Content features; exploration
Clickbait objectives	Satisfaction signals; surveys
Opacity	Explanations; user controls

VI. CONCLUSION

Netflix and YouTube recommenders combine large-scale retrieval, deep multitask ranking and continuous experimentation to personalize content for billions of sessions. The reviewed literature shows the progression from matrix factorization to deep and reinforcement-learning methods. Balancing engagement with satisfaction, diversity and societal impact remains the central challenge for responsible recommendation.

6.1 Acknowledgment

The authors express their sincere gratitude to the Department of Artificial Intelligence and Data Science, Jaihind College of Engineering, Pune, for providing the facilities, guidance

© 2026 IJSRET

and encouragement needed to carry out this work. The authors also thank the faculty members and colleagues whose valuable suggestions helped improve the quality of this review.

REFERENCES

1. C. A. Gomez-Uribe and N. Hunt, "The Netflix recommender system: Algorithms, business value, and innovation," ACM Trans. Manage. Inf. Syst., vol. 6, no. 4, Art. no. 13, 2015.
2. J. Davidson et al., "The YouTube video recommendation system," in Proc. 4th ACM Conf. Recommender Syst. (RecSys), 2010, pp. 293–296.
3. J. Bennett and S. Lanning, "The Netflix Prize," in Proc. KDD Cup and Workshop, 2007.
4. Y. Koren, R. Bell, and C. Volinsky, "Matrix factorization techniques for recommender systems," Computer, vol. 42, no. 8, pp. 30–37, 2009.
5. H. Steck, "Calibrated recommendations," in Proc. 12th ACM RecSys, 2018, pp. 154–162.
6. P. Covington, J. Adams, and E. Sargin, "Deep neural networks for YouTube recommendations," in Proc. 10th ACM Conf. Recommender Syst. (RecSys), 2016, pp. 191–198.
7. X. Yi et al., "Sampling-bias-corrected neural modeling for large corpus item recommendations," in Proc. 13th ACM RecSys, 2019, pp. 269–277.
8. Z. Zhao et al., "Recommending what video to watch next: A multitask ranking system," in Proc. 13th ACM RecSys, 2019, pp. 43–51.
9. M. Chen, A. Beutel, P. Covington, S. Jain, F. Belletti, and E. H. Chi, "Top-K off-policy correction for a REINFORCE recommender system," in Proc. 12th ACM WSDM, 2019, pp. 456–464.
10. M. H. Ribeiro, R. Ottoni, R. West, V. A. F. Almeida, and W. Meira Jr., "Auditing radicalization pathways on YouTube," in Proc. Conf. Fairness, Accountability, Transparency (FAT)*, 2020, pp. 131–141.

© 2026 IJSRET