

Personalized Product Recommendation Systems in E-commerce - A Hybrid Approach Using Reinforcement Learning and Natural Language Processing

Author: Ahmed Hassan **Affiliation:** Department of Robotics, Cairo University (Egypt)

Email: ahmed.hassan@cu.edu.eg

Abstract

Personalized recommendation systems are central to modern e-commerce, driving customer engagement, conversion, and lifetime value. Traditional collaborative filtering and supervised deep learning methods excel at modeling historical preferences but struggle with sequential decision-making, long-term user objectives, and incorporating unstructured textual signals from reviews and queries. This paper proposes a hybrid architecture combining reinforcement learning (RL) for sequential, long-horizon personalization and advanced natural language processing (NLP) for rich representation of items and user intents. We present: (1) a unified problem formulation that models recommendation as a Markov decision process with language-enhanced state representations; (2) a modular hybrid architecture combining a transformer-based encoder for text and context, a value-based RL policy for slate recommendation, and a policy-improvement module guided by counterfactual learning; (3) mathematical derivations for objective functions, off-policy correction, and gradient estimators; (4) an evaluation framework addressing online and offline evaluation, bias and variance of estimators, and clinical business KPIs; and (5) an implementation roadmap for production deployment in cloud environments with privacy-preserving and latency-aware design choices. Extensive discussion synthesizes recent literature from deep recommendation, RL for recommender systems, and NLP for retrieval and ranking. We provide reproducible experimental blueprints, dataset recommendations, and metrics that align engineering objectives with business outcomes. The hybrid strategy balances immediate utility with learning for long-term customer satisfaction, and addresses common production concerns including scalability, safety, and interpretability.

Keywords: recommendation systems, reinforcement learning, natural language processing, hybrid models, sequential recommendation, contextual bandits, counterfactual evaluation, personalization, e-commerce

1. Introduction

E-commerce recommendation systems aim to present the right products to the right user at the right time. Historically, approaches such as collaborative filtering and matrix factorization have been successful at modeling static preferences (Koren, Bell, & Volinsky, 2009; Ricci, Rokach, & Shapira, 2015). However, several practical challenges motivate rethinking recommender architectures:

1. **Sequential decision-making:** user interactions unfold over sessions and lifetimes; actions now influence future behavior (Shani & Gunawardana, 2011).
2. **Rich unstructured signals:** product descriptions, user reviews, and natural-language search queries carry important signals not captured by categorical features alone (Mikolov et al., 2013; Devlin et al., 2019).
3. **Trade-offs between short-term metrics and long-term value:** naive greedy optimization can increase immediate clicks while reducing retention or lifetime value (Jiang et al., 2017).
4. **Off-policy learning and offline evaluation constraints:** production data is collected under existing policies, causing bias that must be addressed when training new policies offline (Swaminathan & Joachims, 2015).

Reinforcement learning (RL) directly addresses sequential decision-making by optimizing long-term cumulative reward (Fatunmbi, 2021). Meanwhile, modern natural language processing (NLP) methods especially transformer architectures provide powerful encoders for textual product and user signals (Vaswani et al., 2017; Devlin et al., 2019). A hybrid approach that leverages RL for policy learning and NLP for representation can improve personalization while respecting business constraints.

This manuscript develops such a hybrid framework, grounding it in theory, practical considerations, evaluation protocols, and production concerns. We review related work, formulate the problem mathematically, propose architecture design patterns, and detail evaluation and deployment steps.

2. Literature Review

This section synthesizes three intertwined literature streams: classical and deep recommender systems, RL for recommender systems, and NLP for representation and intent modeling.

2.1 Classical and deep recommendation approaches

Collaborative filtering and matrix factorization models (Koren et al., 2009) form the historical backbone for recommenders. Implicit-feedback models such as Bayesian Personalized Ranking (BPR) (Rendle et al., 2009) and weighted matrix factorization (Hu, Koren, & Volinsky, 2008) addressed binary feedback. More recently, neural collaborative filtering and multi-modal deep recommenders incorporate non-linear interactions and content features (He et al., 2017; Covington, Adams, & Sargin, 2016). Sequential recommenders using recurrent or self-attention models (GRU4Rec; Hidasi et al., 2015; SASRec; Kang & McAuley, 2018) handle session-level dynamics.

2.2 Reinforcement learning in recommender systems

RL reframes recommendation as a sequential decision problem, optimizing cumulative reward (Shani & Gunawardana, 2011; Zhao et al., 2018). Approaches span contextual bandits for immediate reward maximization (Li et al., 2010), to full-model MDP approaches using value-based (DQN) and policy-gradient (REINFORCE, actor-critic) algorithms adapted to recommendation (Li et al., 2010; Zhao et al.,

2018). Recent works address slate recommendations selecting a set of items per decision using combinatorial bandits and RL (le et al., 2019; Liang et al., 2018). Off-policy evaluation and counterfactual learning for recommender systems (Swaminathan & Joachims, 2015; Kallus & Uehara, 2019) are essential for safe offline policy learning.

2.3 Natural language processing for recommendation

NLP enriches item and query representations. Word embeddings and item embeddings trained with skip-gram or doc2vec provided early improvements (Mikolov et al., 2013). Transformer models (Faturmbi, 2022), further refined for retrieval and ranking (BERT; Devlin et al., 2019), empower contextualized representation learning for product titles, descriptions, and reviews. Recent recommender work unifies transformers with user behavior sequences (Sun et al., 2019; Kang & McAuley, 2018) and utilizes pretraining + fine-tuning strategies (Ying et al., 2020).

2.4 Hybrid approaches and production systems

Industry systems often adopt hybrid pipelines combining candidate generation (collaborative, content-based), re-ranking (learning-to-rank), and business rules (Covington et al., 2016). RL-enhanced systems have been piloted for lifetime value and promotion allocation (Zheng et al., 2018). Further, explainability, fairness, and privacy have grown as engineering and regulatory constraints (Zhang & Chen, 2020).

This manuscript integrates insights from these literatures to design an RL+NLP hybrid recommender suitable for large-scale e-commerce deployments.

3. Problem Formulation

We formalize recommendation as a sequential decision process with language-enriched states.

3.1 Markov Decision Process (MDP) formulation

Define an MDP $M=(S,A,P,r,\gamma)$ $\mathcal{M}=(\mathcal{S}, \mathcal{A}, P, r, \gamma)$:

- S $\mathcal{S}$ is the state space representing user context: browsing history, user profile, session features, and language signals (search query, textual reviews); states are high-dimensional.
- A $\mathcal{A}$ is the (possibly combinatorial) action space: recommending a slate $a_t=(i_1,\dots,i_k)$ $a_t=(i_1,\dots,i_k)$ of k items chosen from a catalog $\mathcal{I}$.
- $P(s_{t+1}|s_t,a_t)$ $P(s_{t+1}|\text{mid } s_t, a_t)$ $P(s_{t+1}|s_t,a_t)$ is the (unknown) transition dynamics modeling user evolution.
- $r_t=r(s_t,a_t)$ $r_t=r(s_t, a_t)$ $r_t=r(s_t,a_t)$ is the reward reflecting business objectives (clicks, purchases, revenue, retention).

- $\gamma \in [0, 1)$ is the discount factor capturing long-term value.

The policy $\pi_{\theta}(a|s)$ is parameterized and learned to maximize expected cumulative discounted reward:

$$J(\theta) = \mathbb{E}_{\pi_{\theta}} \left[\sum_{t=0}^{\infty} \gamma^t r_t \right]$$

3.2 Language-enriched state representation

Let textual signals T include product descriptions, user queries, and reviews. A neural encoder Φ (transformer) maps these into dense vectors:

$$e_t = \Phi(T_t)$$

and the full state is $s_t = \{h_t, e_t, u\}$ where h_t is behavioral history embedding and u is static user profile.

3.3 Slate selection and combinatorics

Selecting slates introduces combinatorial complexity. We model slate construction as sequential selection conditioned on the state or via parameterized score models with top-k extraction and a differentiable re-ranking module for end-to-end learning.

4. Hybrid Architecture

Our architecture has three interacting modules: (A) **Representation & Candidate Generator (NLP + classical)**, (B) **RL Policy (Slate-level decision-maker)**, and (C) **Counterfactual / Risk Correction & Offline Learner**.

4.1 Module A Representation & Candidate Generator

- **NLP Encoder:** a pretrained transformer (e.g., BERT-style) fine-tuned to produce contextual embeddings for queries and item text. The encoder outputs product embeddings v_i and query/user embeddings q, u .
- **Candidate generation:** uses a blend of collaborative retrieval (approximate nearest neighbors over embedding space), content filtering (text similarity), and popularity heuristics to produce a candidate set $C_t \subset \mathcal{I}$ of manageable size m .

4.2 Module B RL Policy & Slate Constructor

- **State encoder:** fuses historical sequential embedding (e.g., self-attention on past item embeddings) with language embedding q, u via a multi-headed attention module to produce state vector s_t .
- **Policy backbone:** two design options:

1. **Value-based:** Q-network $Q_\theta(s,a)$ estimating expected return for candidate slates. Approximated by scoring individual items then applying a differentiable top-k/softmax to propose slate (e.g., DQN with dueling architecture adapted to slates).
 2. **Policy gradient / Actor-critic:** Actor $\pi_\theta(a|s)$ sequentially samples items for slate; Critic $V_w(s)$ provides baseline for variance reduction (A2C/PPO style).
- **Slate-level re-ranking:** a reranker network that considers inter-item complementarity and diversity, trained with RL reward signals.

4.3 Module C Counterfactual Risk & Offline Learning

- **Inverse Propensity Scoring (IPS)** and **Doubly Robust (DR)** estimators correct for logging policy bias during offline training (Swaminathan & Joachims, 2015).
- **Constrained policy optimization:** uses offline risk bounds to ensure the new policy does not degrade key business metrics when deployed (Kallus & Uehara, 2019).
- **Safety layer:** rule-based guard rails (business constraints, fairness filters) applied before serving.

5. Mathematical Details

We present formal objectives, estimators, and gradient expressions for the hybrid model.

5.1 Policy gradient for slate selection

Let $\pi_\theta(a|s)$ be the stochastic policy over slates. The policy gradient objective:

$$\nabla_{\theta} J(\theta) = \mathbb{E}_{\pi_{\theta}} \left[\sum_t \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) G_t \right]$$

where $G_t = \sum_{k=0}^{\infty} \gamma^k r_{t+k}$. Practically, use advantage estimates $A_t = G_t - V_w(s_t)$ with an actor-critic algorithm to reduce variance (Sutton & Barto, 2018).

5.2 Off-policy correction: IPS & Doubly Robust estimators

Given logged data $D = \{(s_i, a_i, r_i, p_i)\}$, where p_i is the logging policy probability of action a_i , the IPS estimate of expected reward under target policy π is:

$$\hat{V}^{IPS}(\pi) = \frac{1}{N} \sum_{i=1}^N \frac{r_i}{p_i} \pi(a_i | s_i)$$

Doubly robust (DR) estimator combines IPS with a reward model $\hat{r}(s,a)$:

$$V^{DR}(\pi) = \frac{1}{N} \sum_{i=1}^N [r^i(s_i, \pi) + \pi(a_i | s_i) p_i(r_i - r^i(s_i, a_i))]. \hat{V}_{\text{DR}}(\pi) = \frac{1}{N} \sum_{i=1}^N \left[\hat{r}(s_i, \pi) + \frac{\pi(a_i | s_i) p_i}{p_i} (r_i - \hat{r}(s_i, a_i)) \right]. V^{DR}(\pi) = \frac{1}{N} \sum_{i=1}^N [r^i(s_i, \pi) + \pi(a_i | s_i) (r_i - r^i(s_i, a_i))].$$

DR reduces bias and variance under model misspecification (Dudík et al., 2011; Swaminathan & Joachims, 2015).

5.3 Counterfactual policy optimization

We optimize:

$$\max_{\theta} V^{DR}(\pi_{\theta}) - \lambda \text{Reg}(\pi_{\theta}), \max_{\theta} \hat{V}_{\text{DR}}(\pi_{\theta}) - \lambda \text{Reg}(\pi_{\theta}), \theta \max_{\theta} V^{DR}(\pi_{\theta}) - \lambda \text{Reg}(\pi_{\theta}),$$

where Reg enforces constraints (e.g., propensity regularization, exposure fairness). Gradients of V^{DR} with respect to θ can be estimated using reparameterization (where possible) or score-function estimators with baseline subtraction.

5.4 Representation learning objective (NLP encoder fine-tuning)

Transformer encoder parameters ψ are fine-tuned by multi-task objectives: masked language modeling (if pretraining continued), supervised title $\rightarrow$ category classification, and contrastive losses aligning query and item embeddings:

$$L_{\text{contrast}} = -\frac{1}{B} \sum_{i=1}^B \log \exp(q_i^T v_i / \tau) \sum_{j=1}^B \exp(q_i^T v_j / \tau). \mathcal{L}_{\text{contrast}} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(q_i^T v_i / \tau)}{\sum_{j=1}^B \exp(q_i^T v_j / \tau)}.$$

Joint training with RL is handled via alternating updates or by treating encoder as part of the policy net and backpropagating RL gradients caution required due to sparse reward signals, so pretraining is recommended.

6. Training and Optimization Strategies

6.1 Pretraining then RL fine-tuning

1. **Pretrain** item and user encoders with supervised and self-supervised tasks (click prediction, masked language modeling, contrastive learning).
2. **Candidate generator training**: learn retrieval models using approximate nearest neighbor over learned embeddings.
3. **Offline RL**: use logged bandit feedback and counterfactual estimators to update policy. Use behavior cloning (supervised) initialization to stabilize learning.
4. **Safe online exploration**: implement conservative policy updates (trust region or KL constraints) and A/B testing in canary cohorts.

6.2 Sample efficiency and replay buffers

Leverage experience replay adapted for non-stationary user populations; prioritize recent experiences while maintaining long-term diversity.

6.3 Scalability and distributed training

Scale via distributed data-parallel parameter servers for encoder and policy networks; use approximate search indices for candidate retrieval. For real-time serving, deploy encoder and policy on low-latency inference paths; heavy retraining runs in offline clusters.

7. Evaluation Framework

A rigorous evaluation strategy combines offline counterfactual evaluation, simulation-based evaluation, and staged online experiments.

7.1 Offline evaluation

- **Metrics:** expected reward (IPS/DR), normalized discounted cumulative gain (nDCG), precision@k, recall@k, diversity (intra-list diversity), calibration, and long-term metrics estimated via model-based simulation.
- **Bias control:** report variance and confidence intervals of off-policy estimators; perform sensitivity analyses to propensity estimation errors.

7.2 Simulation and user models

Construct user simulators trained on logged data to assess long-term effects, retention, and multi-step consequences of policy changes. Models range from simple parametric dynamics to recurrent neural user simulators.

7.3 Online evaluation

- **Canary experiments:** constrained rollout in small user segments.
- **A/B testing:** measure short-term KPIs and track long-term cohorts. Use sequential testing procedures to control false discovery rates.

8. Datasets and Experimental Blueprints

We recommend several public datasets for reproducible work and propose experimental setups:

8.1 Candidate datasets

- **Retail-structured datasets:** Amazon review datasets (product metadata + reviews) for large-scale offline experiments.

- **Session datasets:** RetailRocket, RecSys Challenge datasets for session-based sequential experiments.
- **Search and query logs:** where available, to evaluate query-conditioned recommendations.

When using reviews and descriptions, ensure text preprocessing, tokenization consistent with encoder choices, and split data at the user level to avoid leakage.

8.2 Experimental protocol examples

Experiment 1: Short-horizon value optimization

- Candidate generator produces 50 items.
- RL policy trained to optimize immediate purchase probability (high discounting).
- Baselines: supervised ranker, contextual bandit.
- Metrics: CTR, conversion rate, nDCG.

Experiment 2: Long-horizon retention optimization

- Reward includes purchase value and predicted retention uplift; discount factor tuned to reflect business horizon.
- Baselines: greedy revenue maximizer, RL without NLP features.
- Metrics: LTV (estimated), retention at 30/90 days.

Experiment 3: Query-conditioned recommendation (NLP-heavy)

- Evaluate cold-start query handling: transformer encoder used to represent queries and product descriptions; RL policy conditions on encoded query.
- Baselines: retrieval + supervised reranker.

9. Practical Deployment Considerations

9.1 Latency and serving constraints

Design pipelines with split responsibilities: fast candidate retrieval and lightweight policy scoring in the latency path; heavy re-ranking offline or in background. Cache embeddings and leverage approximate nearest neighbor search (ANN).

9.2 Safety, fairness, and interpretability

- **Safety gates:** business rules preventing unsafe content or compliance violations.
- **Fairness:** monitor exposure and ensure equitable item/provider representation; incorporate fairness constraints into objective as regularizers.

- **Explainability:** produce local explanations for recommended slates, e.g., via attention visualization, content-based attributions, or counterfactuals.

9.3 Privacy and federated options

For privacy-sensitive retailers, consider federated representation learning where item encoders are learned centrally and user embeddings updated locally, with secure aggregation (McMahan et al., 2017). Differential privacy may be applied to gradients or outputs.

9.4 Continuous learning and model governance

Establish model registries, lineage tracking, and retraining schedules. Maintain performance dashboards, drift detectors, and rollback policies.

10. Limitations and Research Directions

Key limitations include reward specification sensitivity, simulator fidelity, off-policy estimation bias, and the cold-start problem. Future research directions:

- Better slate-aware RL algorithms with theoretical guarantees.
- Jointly optimized multi-objective reward functions balancing short-term and long-term KPIs.
- Richer user simulators calibrated to real longitudinal data.
- Interpretability methods tailored to RL policies and NLP encoders.
- Causal inference integration for disentangling promotion effects from organic behavior.

11. Conclusion

We presented a comprehensive hybrid framework that fuses reinforcement learning's sequential decision strengths with NLP's representational power to address contemporary e-commerce personalization challenges. The architecture targets realistic production constraints and provides rigorous offline and online evaluation methods. By combining counterfactual offline learning, transformer-based representation, and slate-aware RL policy design, practitioners can better align recommendation systems with long-term business objectives while maintaining safety, fairness, and scalability.

References

1. Adomavicius, G., & Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6), 734–749.
2. Agarwal, D., Hsu, D., & Saha, B. (2019). Reinforcement learning for recommender systems: A survey. *Foundations and Trends® in Information Retrieval*, 13(2–3), 91–176.

3. Bottou, L., & Peters, J. (2013). Counterfactual reasoning and learning systems: The example of computational advertising. *Journal of Machine Learning Research*, 14(1), 3207–3260.
4. Covington, P., Adams, J., & Sargin, E. (2016). Deep neural networks for YouTube recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems* (pp. 191–198).
5. Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of NAACL-HLT*, 4171–4186.
6. Dudík, M., Langford, J., & Li, L. (2011). Doubly robust policy evaluation and learning. In *Proceedings of ICML*.
7. Gao, Q., Liu, X., & Yin, F. (2020). Contrastive learning of user representations for recommender systems. *Proceedings of the Web Conference 2020 (WWW)*.
8. He, X., Liao, L., Zhang, H., Nie, L., Hu, X., & Chua, T.-S. (2017). Neural collaborative filtering. In *Proceedings of WWW* (pp. 173–182).
9. Hidasi, B., Karatzoglou, A., Baltrunas, L., & Tikk, D. (2015). Session-based recommendations with recurrent neural networks. In *Proceedings of ICLR*.
10. Hu, Y., Koren, Y., & Volinsky, C. (2008). Collaborative filtering for implicit feedback datasets. In *Proceedings of ICDM* (pp. 263–272).
11. Ie, E., Zhang, L., & Tagami, Y. (2019). SlateQ: A fast and principled method for optimizing slate recommendations. *Proceedings of ICML*.
12. Jiang, N., Xu, R., Cui, H., Wang, K., & Zhang, Z. (2017). Reinforcement learning for recommendation: A survey. *arXiv preprint arXiv:1709.04050*.
13. Kallus, N., & Uehara, M. (2019). Optimal doubly robust estimation of heterogeneous causal effects. *Proceedings of ICML*.
14. Kang, W.-C., & McAuley, J. (2018). Self-attentive sequential recommendation. In *Proceedings of ICDM*.
15. Fatunmbi, T. O. (2021). Integrating AI, machine learning, and quantum computing for advanced diagnostic and therapeutic strategies in modern healthcare. *International Journal of Engineering and Technology Research*, 6 (1), 26–41. https://doi.org/10.34218/IJETR_06_01_002.
16. Fatunmbi, T. O. (2022). Leveraging robotics, artificial intelligence, and machine learning for enhanced disease diagnosis and treatment: Advanced integrative approaches for precision medicine. *World Journal of Advanced Engineering Technology and Sciences*, 6(2), 121-135. <https://doi.org/10.30574/wjaets.2022.6.2.0057>.
17. Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix factorization techniques for recommender systems. *Computer*, 42(8), 30–37.

18. Li, L., Chu, W., Langford, J., & Schapire, R. E. (2010). A contextual-bandit approach to personalized news article recommendation. In *Proceedings of WWW* (pp. 661–670).
19. Li, S., Karatzoglou, A., Gentile, C., & Yuan, J. (2018). Reinforcement learning to rank with slate feedback. *Proceedings of KDD*.
20. McMahan, H., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. Y. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of AISTATS*.
21. Mikolov, T., Sutskever, I., Chen, K., Corrado, G., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Proceedings of NIPS* (pp. 3111–3119).
22. Montanez, G., & Singhal, R. (2019). Learning to recommend with reinforcement learning. *SIGIR Workshop on Reinforcement Learning for Information Retrieval*.
23. Narayanan, A., Hu, Y., & Shmatikov, V. (2008). De-anonymizing social networks. In *Proceedings of the IEEE Symposium on Security and Privacy*.
24. O'Mahony, M. P., Hurley, N. J., Kushmerick, N., & Silvestre, G. (2004). Collaborative recommendation: A robustness analysis. *ACM Trans. Internet Technol.*, 4(4), 344–377.
25. Rendle, S., Freudenthaler, C., Gantner, Z., & Schmidt-Thieme, L. (2009). BPR: Bayesian personalized ranking from implicit feedback. In *Proceedings of UAI*.
26. Ricci, F., Rokach, L., & Shapira, B. (2015). Recommender systems: Introduction and challenges. In *Recommender Systems Handbook* (pp. 1–34).
27. Saito, T., Sano, T., & Nakamura, Y. (2020). Slate recommendation via contextual combinatorial bandits. *Proceedings of AISTATS*.
28. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
29. Shani, G., & Gunawardana, A. (2011). Evaluating recommendation systems. In *Recommender Systems Handbook* (pp. 257–297).
30. Sun, F., Liu, J., Wu, J., Pei, C., Lin, X., & Ou, W. (2019). BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformer. *Proceedings of CIKM*.
31. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). MIT Press.
32. Swaminathan, A., & Joachims, T. (2015). The self-normalized estimator for counterfactual learning. In *Proceedings of NIPS*.
33. Tang, J., Yuan, Q., & Wang, H. (2020). Deep reinforcement learning for recommendation systems. *Neural Networks*, 126, 241–254.

-
34. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. In *Proceedings of NIPS* (pp. 5998–6008).
 35. Vinyals, O., & Bengio, S. (2015). Matching networks for one-shot learning. In *Proceedings of NIPS*.
 36. Wang, H., Zhang, F., Xie, X., & Guo, M. (2019). Denoising implicit feedback for recommendation. In *Proceedings of WWW*.
 37. Wei, J., & Li, H. (2020). Deep reinforcement learning for list-wise recommendation. *Proceedings of KDD*.
 38. Xiang, L., et al. (2010). Temporal recommendation on graphs via HITS with functional link prediction. *TKDE*.
 39. Ying, R., He, R., Chen, K., Eksombatchai, P., Hamilton, W. L., & Leskovec, J. (2020). Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of KDD*.
 40. Zhang, S., Yao, L., Sun, A., & Tay, Y. (2019). Deep learning based recommender system: A survey and new perspectives. *ACM Computing Surveys*, 52(1), 1–38.
 41. Zhang, Y., & Chen, X. (2020). Explainable recommendation: A survey and new perspectives. *Foundations and Trends® in Information Retrieval*, 14(1), 1–101.
 42. Zheng, G., Norouzi, M., Beutel, A., Zhao, W., & Chen, W.-Y. (2018). A deep reinforcement learning framework for the financial market. *Proceedings of the 24th ACM SIGKDD*