A Comparative Study of Recommender Systems in E-Commerce

Kimberly Lim, Khushi Modi, Riley Durbin, Manjila Singh

Abstract— **Recommender systems are being widely used by modern e-commerce companies to make data-driven strategies to boost user satisfaction and retention. In this paper we focus on comparing the different recommender system approaches that e-commerce applications, specifically Amazon and Netflix, use to shape user experience and engagement. Amazon uses a combination of collaborative filtering, content-based filtering and deep learning techniques to better understand and make accurate predictions on user preferences. It focuses on extensive user engagement and purchase patterns to drive sales. On the other hand, Netflix utilizes a hybrid approach that blends a variety of algorithms together with machine learning techniques, popularity scales and similarity metrics. It aims to increase viewer satisfaction and watch time. Both platforms share similarities when it comes to product recommendations as they both keep track of interactions with products and services based on prior behavior patterns. The results highlight the contrast between Netflix and Amazon to meet specific commercial goals while ensuring optimal user experience.**

Keywords— recommender systems; collaborative filtering; content-based filtering; hybrid recommendation models; deep learning; cold-start problem

I. INTRODUCTION

With the growth of e-commerce, the range of products and digital contents have grown exponentially in recent years. While this abundance offers greater choice, it has introduced a problem as customers spend hours and hours browsing online for a decent item or content. To resolve this problem, most modern e-commerce applications use different recommender systems to give suitable suggestions to their customers through data-driven strategies. Big companies like Amazon, eBay, Spotify, and Netflix have successfully implemented such systems to help users find what they want easily. Recommender systems use various methods to come up with personalized recommendations for users. These systems look at similarities between items and even users before providing helpful suggestions. Some algorithms highlight popular or trending items, while others predict what a user would like based on their previous behavior patterns like their ratings, reviews for a product, purchases or skip patterns. Certain algorithms identify user groups who share similar preferences and recommend products that those users have liked in the past. These recommender systems help companies to retain their consumers for a longer period of time by preventing users from feeling overloaded with choices.

Most of the recommender systems use one of the two filtering methods: content-based and collaborative filtering. Collaborative filtering takes user similarities into account before making any predictions. It makes use of user interactions such as product ratings or past purchase history to predict future decisions by assuming that users will continue to follow similar behavior patterns in the future. On the contrary, content-based filtering uses product's feature information and users' past behavior to recommend items that are comparable to other items that the user has previously shown interest in. For example, a user who often watches action movies will be recommended similar titles within the same genre. In practice though, many e-commerce platforms have combined both content-based filtering and collaborative filtering to implement hybrid models that improve accuracy and optimization.

However, recommendation systems are not always perfect. One major issue with recommendation systems is that consumer preferences cannot be accurately predicted when there is not enough prior data from users. This problem is called the cold-start problem. This problem can take two different forms: the new item cold-start problem or the new user cold-start problem. The new item cold-start problem happens when a new item is recently added with no user interaction history. On the other hand, the new user cold-start problem occurs when a system fails to give recommendations to newly added user with no historical data. Other issues include working with huge amounts of data, scalability, and the need for real-time recommendations while ensuring the best user experience.

This study majorly focuses on comparing the different recommendation systems being used by two of the fastest-growing e-commerce companies: Amazon and Netflix. By looking at how they use collaborative filtering, content-based filtering and hybrid systems, the study aims to gain useful insights into making the current recommendation systems more effective in the field of e-commerce.

II. AMAZON CASE STUDY

A. Models and Algorithm

Amazon's recommendation engine uses a mix of machine learning techniques, with collaborative filtering at its core. In fact, Amazon pioneered item-to-item collaborative filtering in the early 2000s to address scalability issues with traditional user-based approaches. Instead of finding similar users, item-to-item collaborative filtering computes the similarity between products. For each item a user has interacted with, it finds related items, then aggregates those to recommend to the user. This method proved both highly scalable and effective, as it produces real-time recommendations and handles Amazon's massive catalog with millions of items [1][2]. Over time, Amazon augmented collaborative filtering with content-based filtering and hybrid models that blend multiple techniques. Content-based filtering utilizes item attributes like category, genre, or specifications to further enhance the recommendations by providing more semantic meaning to a product. For example, content-based methods allow Amazon to recommend items with similar attributes, like similar book genres or product features, complementing the behavior-based collaborative filtering approach [3].

B. Advanced Techniques

As data and computing power grew, Amazon invested in deep learning to improve recommendations. They moved beyond simple matrix factorization and developed neural network models. This included autoencoders and neural collaborative filtering to learn latent representations of users and items. These neural models capture complex patterns, such as sequential viewing behavior or item co-occurrence, that older methods missed. Amazon even open-sourced its in-house deep learning library DSSTNE (Deep Scalable Sparse Tensor Network Engine) to handle extremely large recommendation models. DSSTNE supports networks with hundreds of thousands to millions of input and output nodes, using model-parallel GPU training to fit these wide networks. This allows Amazon to train on huge sparse datasets (many customers × many products) that result in trillions of model parameters. After training, generating recommendations for each customer is parallel. The system can distribute prediction tasks across many GPUs to compute personalized lists for millions of users quickly [4].

Amazon's recommendation architecture also incorporates bandit algorithms and reinforcement learning as part of hybrid models. Multi-armed bandits help the system explore new or less-known products while still exploiting known user preferences, which is useful for introducing serendipity and for promoting new items. For example, a bandit-based component might decide in real time which recommendation strategy or content to show, learning from user clicks to maximize engagement. Amazon scientists have also leveraged causal inference to improve recommendations, aiming to identify which user–item relationships are truly meaningful versus just correlated [3]. This helps avoid spurious suggestions by accounting for confounding factors, such as making sure the algorithm recommends an item because the user genuinely would like it, not just because it's popular.

Recently, Amazon has been exploring transformer-based models and generative techniques for recommendations. One research innovation is GPT-4Rec, a framework that uses NLP transformers to generate "queries" from a user's history and retrieve relevant items, effectively learning user and item embeddings in a language space. Such approaches aim to improve recommendation relevance, diversity, and interpretability by borrowing advances from natural language processing [5].

C. Scalability and Personalization at Scale

Scalability is a critical concern, as Amazon serves hundreds of millions of users and a catalog of millions of products. The item-to-item algorithm was a key innovation because its online computation scales independently of the total number of users, focusing on a user's own history and similar items. This means Amazon can pre-compute item similarity matrices offline, then quickly look up related items for any user request. The algorithms are engineered to return results in under a few hundred milliseconds to keep the site responsive [2]. Amazon's infrastructure uses distributed data processing (like Apache Spark on AWS) to crunch huge volumes of interaction data and update models. Training of deep learning models is distributed across GPU clusters (using AWS services like ECS with auto-scaling) so that even extremely large models can be fit within reasonable time. For example, Amazon can partition the training workload for a neural recommender across many GPUs in parallel. This doubles the number of GPUs roughly halves the training or inference time [4]. This ensures that personalization models can be refreshed frequently with the latest data. The system ingests each new interaction, such as a click or purchase and can update recommendations dynamically so that they reflect a user's very latest behavior. In practical terms, each Amazon user's experience is personalized in real time. As Jeff Wilke, the former CEO of Amazon's consumer division, noted, the site "changes radically based on customer interests," showing different products to an engineer versus a new mother [3]. This level of personalization at massive scale is a major engineering feat, combining efficient algorithms, parallel computation, and fast data pipelines.

D. Data Sources for Recommendations

Amazon's recommender draws on a rich variety of data. Implicit behavioral data is primary. Every user's purchase history and product page views are key signals of interest. Additionally, Amazon tracks browsing behavior and interaction history like wish lists, ratings given, and even mouse-over or scrolling patterns in some cases. Browsing behavior data points may include search queries, items you have clicked on or added to cart, time spent, and much more. The engine also uses the "wisdom of the crowd", which utilizes patterns like items frequently bought together or commonly co-viewed by similar users. User–item relationships are mined in many ways which come from analyzing billions of transactions. Some common examples are "users who bought X also bought Y", "users who viewed A ultimately purchased B", and so on. Product-product relationships are also leveraged. For example, recognizing that two books are similar because many customers purchased both, even if the items are in different categories. While Amazon historically relied more on behavior data than on item descriptions, it does utilize product metadata as well. This includes, but is not limited to, category tags, genre, brand, and price range. It can then feed content-based recommendation components as well. This is especially helpful for new items with little interaction history, known as the cold-start problem. Amazon can also incorporate user profile information if available. This includes demographics such as location, age or Amazon Prime membership status, all of which might slightly influence recommendations. For example, a user's region could affect language-specific content or seasonally relevant suggestions [3]. However, Amazon typically emphasizes what you do on the site over who you are, meaning behavioral cues drive most of the personalization. In newer domains like Alexa and voice shopping, even customer voice interactions and queries could become data signals for recommendation in the future. Across all these sources, Amazon's data pipeline continuously aggregates explicit feedback and implicit feedback. Explicit feedback looks like star ratings, "likes", or helpful votes and implicit feedback look like purchases, views, and skips.

E. Further Innovations in the Recommendation System

Amazon's scale and diversity (retail, video, music, etc.) push its recommender system to evolve constantly. The company employs dedicated research teams (Amazon Science) to experiment with next-generation recommender algorithms. We've already mentioned the use of deep neural networks and transformers for session-based recommendations. Another innovation is the use of knowledge graphs – Amazon has projects to build product knowledge graphs that encode relationships (e.g. "isAccessoryOf" or "similar style as") which can give the recommendation engine a form of common sense beyond pure statistics [6]. For instance, a knowledge graph could help recommend a camera tripod to someone viewing a DSLR camera, even if that combination hasn't appeared frequently in past data, by understanding the complementary relationship. On the exploratory front, Amazon has looked into reinforcement learning for longer-term optimization of recommendations. A reinforcement learning agent could, for example, choose a sequence of products to show a user with the goal of maximizing total engagement or revenue over a session, balancing immediate clicks with novel item discovery. Such approaches try to overcome the myopic nature of one-shot recommendations by considering the long-term reward, similar to how video streaming services optimize a user's entire session.

F. Enhancing Customer Engagement

Amazon's recommendation system is not just a back-end algorithm; it directly shapes the customer experience on the site. By showing users items that align with their interests and needs, it keeps shoppers engaged longer and browsing more deeply. For example, a customer who just bought a camera might see suggested lenses, tripods, or memory cards. These are all useful add-ons that encourage them to continue exploring. These personalized suggestions feel like curated shopping advice rather than generic ads, which improves click-through rates and engagement. In fact, the click-through and conversion rates of Amazon's personalized recommendations vastly exceed those of untargeted content like banner ads or top-seller lists. Customers are more likely to click on a recommended item that "speaks" to their interests than on a random product. This means users spend more time on Amazon's platform, discover more products, and often find what they are looking for. Users may also find things they did not realize they wanted. It creates a cycle. The more a customer interacts, the more data the system gets to refine future recommendations, which in turn further improves engagement [2].

G. Driving Sales and Revenue

The ultimate business goal of recommendations is to increase sales, and Amazon's results have been remarkable. A McKinsey study estimated that about 35% of Amazon's product sales are driven by its recommendation engine [3]. In other words, well over a third of what customers buy on Amazon comes from those "You might also like" and "Frequently bought together" suggestions. This statistic underscores how critical the recommender system is to Amazon's growth. By smartly upselling and cross-selling, Amazon raises the average order value (customers buy additional items they were recommended) and improves conversion rates (customers who see relevant suggestions are more likely to find something they want to purchase). Internal quotes from Amazon have highlighted that the recommendation algorithms personalize each customer's storefront, which directly boosts sales and customer satisfaction [3]. Over the years, Amazon's meteoric revenue growth (from a $10B company to hundreds of billions in annual sales) has been aided in no small part by the ever-improving recommendation engine. One study noted that customers who click on recommendations are 4 times more likely to add those items to cart and complete a purchase than customers who don't interact with recommendations [7]. This lift in conversion translates to massive revenue when scaled across millions of shopping sessions. The recommender system also helps retain customers by creating a personalized experience that is hard to leave. It is like a salesperson who knows your tastes. This increases customer lifetime value and loyalty, contributing to Amazon's business growth. As evidence of how integral this system is, Amazon integrates recommendations into nearly every part of the purchasing process, and analysts attribute a significant portion of Amazon's year-over-year sales increases to its recommendation integration.

H. Impact on Business Metrics

The recommendation system's contribution to revenue is significant and measurable. Beyond the headline figure of one-third of sales, it also improves key metrics like average revenue per user, frequency of repeat purchases, and the size of shopping carts. By encouraging add-on purchases (like attachments, warranties, complementary goods), Amazon increases the basket size in a single session. By suggesting new products aligned with a user's taste, it brings users back for additional orders in the future. The system also helps inventory turnover – slow-moving items can find buyers if recommended to the right person, and popular items get even more exposure. In fact, Amazon can use the engine strategically to promote strategic products (for example, Amazon's own private-label products or new marketplace sellers) by injecting them as recommendations where appropriate. Overall, the recommender is a major part of why Amazon's marketing is so efficient; rather than spend on broad advertising, they leverage their on-site personalization which essentially markets products directly to the most likely buyers. This efficiency translates to better margins and growth.

I. Personalization and Customer Satisfaction

From a customer perspective, Amazon's recommendations make the shopping experience feel tailored and convenient. Instead of wading through millions of products, users are presented with options likely to fit their needs. This level of personalization increases customer satisfaction, as shoppers feel understood by Amazon's platform. A famous quote from Amazon's early days is that their goal is to "delight customers by helping them find the right product", not just any product. Personal recommendations achieve that by surfacing relevant items. For example, a new mystery novel from an author you've bought before, or a compatible accessory for a device you own. It can also help customers discover new brands or products they wouldn't have found via search alone, enhancing the sense of a bespoke shopping experience. Amazon reports that customers respond very positively to personalized content. This is backed by industry studies that show 80% of consumers are more likely to buy from a brand that personalizes the experience [5]. By changing the store for each user, such as showing baby toys to the new parent or programming books to the engineer, Amazon makes each customer's journey unique [3]. This not only drives sales but builds trust and loyalty. Shoppers come to expect that visiting Amazon will yield useful suggestions, saving them time and effort in finding items. Amazon has also extended personalization to email recommendations and push notifications, reminding customers of products they might like. The consistency of relevant suggestions across channels reinforces customer engagement with the Amazon ecosystem as a whole.

J. On-Site Personalized Product Recommendations

The most visible use case is the personalized product suggestions on Amazon's retail website. Nearly every page has some recommendation module tailored to the user. On the homepage or the "Your Recommendations" page, customers see a selection of items "Recommended for You" based on their past browsing and purchases. These might span across various categories, which demonstrates how Amazon tries to cover the range of a user's interests. For example, suggesting both electronics and books if you have shown interest in both. On product detail pages, Amazon includes carousels like "Customers who bought this item also bought" and "Customers who viewed this item also viewed", which is illustrated in Figure 2. These leverage collaborative filtering on a per-item basis, which are essentially crowdsourced suggestions from users with similar behavior. Such modules help with product discovery. For instance, perhaps you are looking at a kitchen blender, and the "viewers also viewed" shows you a food processor that others considered, giving you alternatives. They also act as gentle recommendations to consider related products. Another module is "Inspired by your browsing history," which lists items similar to what you have recently looked at, reminding you of items you showed interest in but perhaps did not purchase. This is illustrated in Figure 1. This can re-engage users to revisit items or see close alternatives [7].

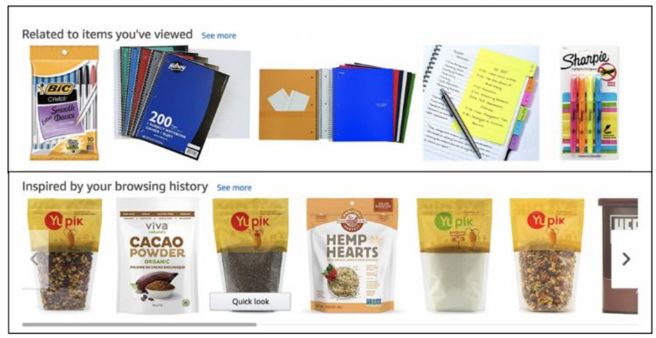

Fig. 1. Example of two of Amazon's recommendation modules that can be seen on their site.

Fig. 2. Example of a carousel on Amazon's site that leverages collaborative filtering.

K. Cross-Selling and Upselling

Amazon's engine doesn't just recommend items in isolation. It often suggests combinations of items to increase basket size. The classic example is the "Frequently Bought Together" section on a product page. This module will display the item you are viewing alongside one or two complementary products, often with a combined price and a one-click option to add all to cart. This is illustrated in Figure 3. It uses association rules mined from transaction data. If many customers who bought item A also bought item B, Amazon will recommend B when you view A. The goal here is cross-selling. Cross-selling is the practice of encouraging the purchase of related items. For example, on a camera page, "Frequently Bought Together" might show a camera bag and an SD memory card as complements. This both enhances customer convenience (you might indeed need those accessories) and increases the average order value by bundling items. Cross-selling recommendations like these have a direct impact on sales. They act as reminders of items the customer may have forgotten and as inspiration for additional purchases. Amazon also uses emails for cross-selling. After a purchase, you might receive follow-up emails like “You bought X, you might also like Y” [8].

Fig. 3. Example of the ”Frequently bought together” recommendation module for Amazon to cross-sell and up-sell.

Cross-selling recommendations like these have a direct impact on sales. They act as reminders of items the customer may have forgotten (batteries for an electronic device, or a case for a phone) and as inspiration for additional purchases (a user buying a gaming console might be tempted by a suggested popular game for that console). Amazon also uses emails for cross-selling: after a purchase, you might receive follow-up emails like “You bought X, you might also like Y.” These recommendation emails have reported high conversion rates. In fact, analysts noted Amazon’s email recommendations can be even more effective than on-site recommendations in some cases, possibly because they reach the customer when they are not actively browsing and draw them back to the site [8].

All these on-site recommendations contribute to making the Amazon website highly “sticky.” As users click around, the site continuously updates with new suggestions (“Since you liked X, maybe Y?”), mimicking the helpful guidance of a salesperson. Search results pages on Amazon also use recommendations: for instance, after you search for a product, Amazon might show a banner of recommended items related to your query or items “Trending in [category]” that align with your interests. Overall, personalized product recommendations on-site serve both to personalize the browsing experience and to drive cross-sell/up-sell.

L. Ethical Considerations and Challenges

While Amazon’s recommendation engine is powerful, it’s not without challenges and criticisms. One concern is the “filter bubble” effect, which is the idea that personalization can narrow a user’s exposure to new or diverse content. If the algorithm only ever shows you items similar to those you’ve already seen or bought, you might miss out on exploring other categories or be “stuck” in a certain profile. An analysis of Amazon’s recommender found evidence that it can reinforce user preferences to the point that it becomes harder for people to change or broaden their tastes over time [9]. For example, if you have historically bought mystery novels, Amazon might keep recommending mysteries and never show you science fiction, even if you would enjoy it. Thus, this creates a self-reinforcing loop. Amazon must balance relevance with diversity. Accordingly, they should likely include some degree of novelty or diversity in recommendations (occasionally showing less predictable items) to mitigate this. Another ethical aspect is preference manipulation. Critics wonder if the recommender is truly serving the customer’s interests or nudging customers to serve Amazon’s short-term sales goals. There is a fine line between helpful suggestions and manipulative upselling. Amazon has to ensure the recommendations maintain customer trust (e.g. not exclusively pushing higher-priced items or Amazon’s own brands in a way that feels biased).

M. Data Privacy

With great personalization comes great responsibility for data privacy. Amazon collects a vast amount of data on each customer’s shopping behavior, raising the question of how that data is used and protected. Amazon’s official stance is that customer data is used responsibly to improve the user experience, and customers retain control over their information. They have implemented tools for transparency. For instance, users can view and delete their browsing history and can manage which past purchases inform their recommendations. This allows customers to remove items (say, a one-time gift purchase) so they don’t keep getting related suggestions, thereby giving some control over the personalization process. Amazon also abides by strict privacy policies and regulations (like GDPR), and they emphasize data security measures to protect personal information. Importantly, Amazon does not share individual-level browsing or purchase data with third parties for the purpose of recommendations on other sites. The data is primarily used in-house to better your Amazon experience. Nonetheless, some users may feel uneasy knowing that “Big Amazon” knows so much about their shopping habits. The company addresses this by being transparent (providing privacy notices and explanations) and highlighting the benefits customers get in return (e.g. “We can show you items inspired by your past purchases” as a feature [10].

III. NETFLIX CASE STUDY

Netflix's recommender system is integral to their success as a streaming service. User experience is one of the most important aspects for Netflix, and their high-quality recommender system improves their customer retention rate and increases customer time on the service. According to a 2020 "Deep Dive into Netflix's Recommender System", "80% of stream time is achieved through Netflix's recommender system" [11].

A. Frontend Design

The frontend design is simple and clear for the user. There are a series of rows that each represent a theme or category. Within each of these rows, the strongest recommendations are on the left, and the strongest recommendations for the rows are at the top. As a result, the top left corner of the grid (where the user's cursor is located when they open the application) is where the overall strongest recommendations are located. This is an intuitive frontend design that enables users to quickly access the items that Netflix recommends to them. This is illustrated in Figure 4.

Fig. 4. Netflix's frontend interface, designed to place the "best" recommendations in the top left corner [11].

B. Hybrid Recommendation Approach

The actual recommender system that Netflix uses can be described as a hybrid approach. It blends a variety of algorithms together to combine some machine learning techniques with popularity scales and similarity metrics. Logically, it makes sense for Netflix to use a mix of collaborative filtering, which uses similarities in streaming habits between users to find recommendations, and content-based recommendations, which uses the previously watched or interacted with items to recommend similar videos. However, Netflix faced a difficult challenge to come up with the best way to leverage these different approaches.

C. Content-Based Recommendations

The content-based recommendations aspect of the algorithm is clear in the rows labeled "Because you watched". This row is referred to as the "Video-Video Similarity Ranker" [11]. The items in this row are not related to the user's history or preferences, but they are simply items that are the most similar to a film or TV show that the user previously watched. In this way, the display of this row on the user's homepage (and the index of this row in the lists of rows) is a personalized algorithm that attempts to predict what the user would want to see, but the actual films in this list are strictly based on content-based filtering. An example of this row is displayed below in Figure 5.

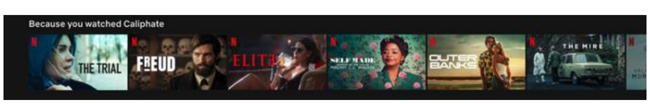

Fig. 5. An example of the "Because you watched" row based on the content-based recommendation system [11].

D. Collaborative Filtering Recommendations

An example of the collaborative filtering algorithm aspect is the "Trending Now" row. This algorithm uses data from other Netflix users to recommend items that are currently popular for the user. This is classified as a collaborative algorithm instead of a content-based algorithm because it recommends items based on other users' actions instead of the content it is recommending. Netflix views this as a strong recommendation, so this "Trending Now" row is usually near the top for a user. An example of this row is displayed in Figure 6.

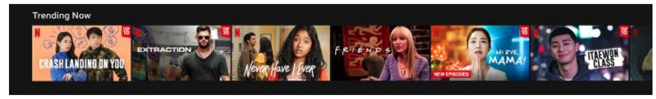

Fig. 6. An example of the "Trending Now" row based on the collaborative filtering recommendation system [11].

Since Netflix uses a variety of ranking systems, there are some other types of rows, like a "Top Picks" row or media category rows. These are generally based on a mix of user activity data and item popularity. Specifically, these other three algorithms are referred to as "Personalised Video Ranking (PVR)", "Top-N", and "Continue Watching" [11]. The PVR algorithm is a general algorithm that finds items that fit a certain category. The "Top-N" algorithm is similar, but it looks at the top recommendations regardless of category. The "Continue Watching" algorithm identifies items that the user has started but not finished. This algorithm ranks these based on how likely it thinks that the user is to continue watching the item.

E. Page Generation and Item Ordering

Assembling these rows is a task of its own, and then the Netflix application must decide which rows need to be shown to the user and in which order. The goal of this page generation algorithm is to satisfy the frontend interface goal of having the strongest recommendations in the most easily accessible location for the user. The article about the Netflix recommender system illustrates the process that the algorithm takes to generate the page of recommendations that the user sees. The goal of this process is to provide the user with similar items, unfinished items, diverse items, and trendy items, all at the same time. This is a difficult task when the available library of media is as large as Netflix's is. The users also hold a certain standard in their mind that Netflix needs to reach. The overall algorithm starts with the total library of content, uses the previously mentioned algorithms to sort these items into rows, and then creates a ranking of them to display to the user. This process is illustrated in Figure 7.

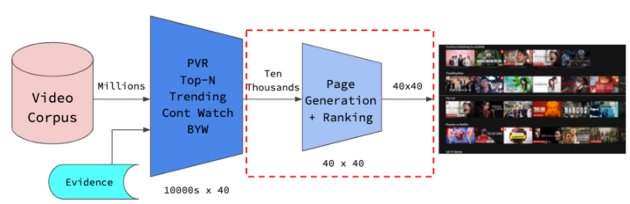

Fig. 7. Netflix's Model Workflow, which selects items from their content library to recommend to the user [11].

F. Machine Learning for Recommendation Page

One problem with this approach is the potential for the algorithms to provide similar results, and although this might match the user's interests, it would result in a page of repeated/similar recommendations. To fix this, Netflix uses a machine learning approach, which is common in data science and recommender systems. This machine learning model is trained with historical recommendation pages that have been created and the interactions from the user. This includes what recommendations they viewed and what items they played [11]. To test and score this model, the model creates a recommendation page and is told which items the user would have selected, interacted with, or watched. An illustration of this machine learning process is displayed in Figure 8.

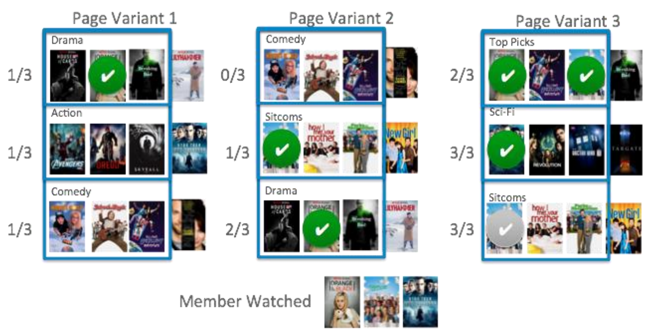

Fig. 8. An illustration of a machine learning model's recommendation pages and the items interacted with by the user [11].

This is a strong way to create recommendations, as it combines a variety of approaches to ensure that the user is provided with items that are similar to their recent watch habits, popular at the time or in their area, and are different in case they want to try something new. The goal of this approach is to be strong enough that the user spends a minimum amount of time selecting an item to increase the amount of time watching it. As previously mentioned, Netflix uses a collaborative filtering approach to compare a user with users in their area or users in the same time frame. This is primarily used to find items that are currently trending. Netflix also uses a content-based recommendation system to provide the user with item recommendations that are similar to previous items they have watched. These recommendation systems are used to generate the rows, and then a mix of scoring algorithms and machine learning models are used to display the rows in an optimal order on the user's homepage.

G. Multi-Tiered Architecture

Netflix's complex recommender systems can further be separated into a multi-tiered architecture. This architecture is designed to be scalable and efficient, as it needs to handle over two hundred million global users [12]. Netflix's recommender system can be split into three different layers. These are the offline, nearline, and online layers. This complex architecture design is illustrated in Figure 9.

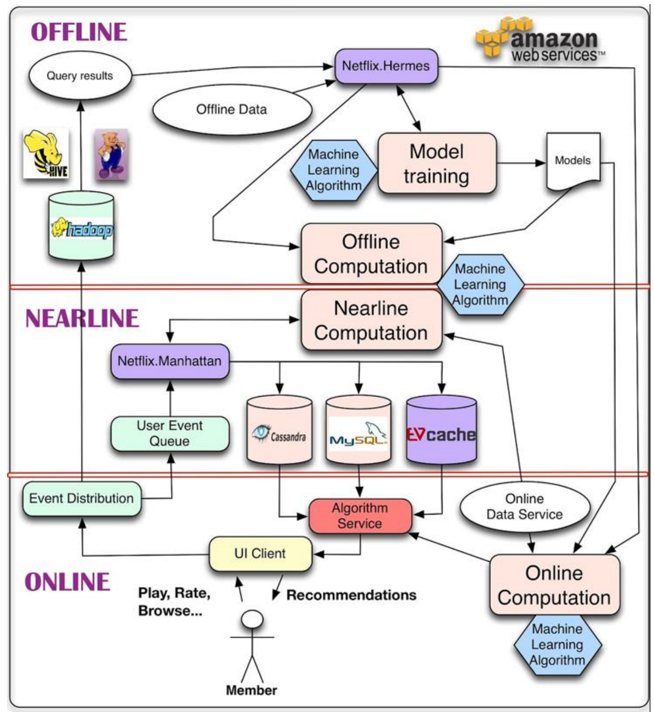

Fig. 9. An illustration of Netflix's detailed architecture diagram, highlighting three of their data computation strategies [11].

Each of these layers have their own role for handling data computation and generating recommendations for Netflix users. The offline layer uses historical user data with less time constraints to generate recommendations. This allows Netflix to run more computationally expensive models, but it has the challenge of becoming "stale" [12] because it does not include real-time interactions from the user. Since Netflix wants the best personalization possible, it combines offline computation with online computation. However, this is challenging to combine. The online computation includes recent user interactions in real-time, but this can be costly and difficult to scale. Since it also has to handle latency issues, it cannot be fully relied on. This is the reason that Netflix also includes nearline computation, which can be viewed as a balance between offline and online computation. This allows Netflix to process user actions quickly and asynchronously. This is a compromise between online and offline, as it performs computations mimicking online computation, but it does not need to occur in real-time. One example of this is that recommendations can be updated as soon as an item has been watched by the user by using the nearline computation. This allows the recommendation list to update without needing to fully retrain a model. This helps Netflix to keep their recommendations personal, relevant, and efficient.

H. Data Collection

Netflix's recommender system relies on data collection from the user to understand what the user is interested in and what the most relevant items to recommend are. Netflix collects implicit and explicit taste data [13]. Explicit data refers to item ratings, demographic data, and thumbs-up or thumbs-down ratings from the user. Implicit data refers to user events like clicks, watches, durations, pauses, and searches. Netflix also considers data like device type, location, and time of day to create an extensive profile for each user. This data collection is the key to their personalization algorithms, and this data is considered in each item recommended and row ordered.

I. Personalized Artwork

Beyond the items recommended to the user, Netflix goes as far as displaying different artwork for each item based on recommender systems. Netflix performed A/B testing to test different artworks for items to see which ones resulted in better click rates and play durations. Their testing revealed that they could "widen the audience and increase engagement by using different artwork" [14]. Over time, Netflix began to factor this aspect of personalization into their algorithms, using user data to dictate which art was used for each item. Figure 10 shows an example of different users receiving different personalized artworks.

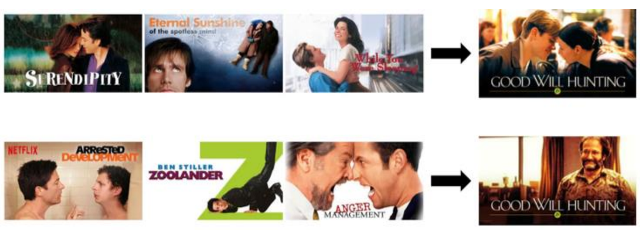

Fig. 10. An example of user actions leading to different art displayed for the same item [13].

In the example in Figure 10, the user on the top had watched *Serendipity*, *Eternal Sunshine of the Spotless Mind*, and *While You Were Sleeping*, which are all films that feature romance. As a result, the artwork that is displayed for their recommendation of *Good Will Hunting* features Matt Damon and Minnie Driver. In this case, Netflix believes that this user would be more likely to click on this film with this artwork displayed. In the bottom example, the user had watched *Arrested Development*, *Zoolander*, and *Anger Management*, which are all comedies. In this case, the artwork for *Good Will Hunting* features Robin Williams. This is a great example of the detailed personalization that Netflix prioritizes for their users. In their blog, a Netflix representative claimed that "if you don't capture a member's attention within 90 seconds, that member will likely lose interest and move onto another activity" [14]. It is a priority for Netflix to maximize user watch time and engagement, so it is clear why they invest in the research and development of their recommender system.

J. Business Impact of Personalization

By using detailed user data, Netflix is able to offer convenient content discovery for its users to keep them engaged and subscribed each month. Personalized recommendations help users make decisions quicker, which means they spend more time watching the movies and shows that Netflix offers. This recommender system is successful from a business perspective for a few main reasons. First, it has been proven to increase user engagement, as 80% of user stream time comes from personalized recommendations [11]. This system also helps to recommend niche items to users, which increases the return on investment for smaller movies and shows that may have a small audience [15]. This system also can help Netflix to make better decisions when acquiring and producing movies and television shows. Since they have the data to predict that an item would be recommended to and viewed by a large number of users, they have a quantitative reason to produce or acquire that item. Finally, personalization leads to higher customer satisfaction, which increases the length of users' subscriptions. These points support the reason behind the research and development, as Netflix has a significant financial incentive to provide an excellent recommender system.

K. Cold-Start Problem

Although this system is successful, there are still challenges that Netflix faces. Namely, these are the cold-start problem, privacy, scalability, and algorithmic bias.

The cold-start problem refers to the challenge of recommending content to new users that do not have existing user data when they first sign up. To address this, Netflix uses "rating elicitation" [16] when a user signs up to understand some of the user's preferences. An example of this is displayed in Figure 11.

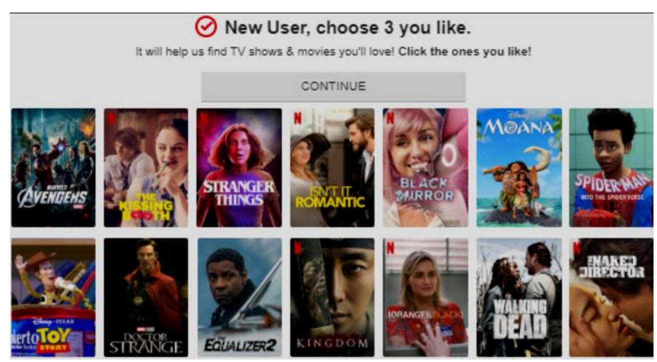

Fig. 11. An image of Netflix's rating elicitation for a new user [16].

A new Netflix user is presented with a set of items on Netflix so they can choose the ones that they are interested in. This helps Netflix with the cold-start problem because they can provide initial recommendations based on the small list of interests that the user provides. These initial recommendations will not be as strong as future recommendations made from the user's data, but they are a better starting point than a true cold-start. One of the challenges with this rating elicitation method is in the choice of items to show the user. Netflix has to display options that provide enough information about the user's preferences without showing too many options so that the user is not bothered by the interference. Netflix also uses some implicit data like location and device type to select initial recommendations.

The cold-start problem produces a difficult balance between the desire to immediately provide the user with relevant recommendations and the necessity to respect user privacy. Although there may be other ways to gain more information about a user, like the use of cookies or other data collection methods, it is important for Netflix to keep the trust of their users by avoiding invasive data gathering. It could lead to more relevant recommendations, but it can also be viewed as unethical. Since the most important metric for Netflix is user retention and repeated subscriptions, the trust of their users is more important. Instead of relying on third-party data, Netflix focuses on gathering first-party data like clicks and interactions to learn about a user and train models. By prompting a user for their preferences when they are initially signing up, they are able to create a personalized experience for users without compromising the user's privacy.

L. Scalability

Netflix also has to consider scalability when designing and implementing their recommender system. In a Netflix Technology Blog post from 2022 [17], two Netflix engineers that work on the recommender system detailed the main priorities for maintaining the recommender system on such a large scale. The main points from the article were to predict issues before they occur, develop tools to better understand why issues occur, and improve the process after each fix. This set of practices allows them to keep the system performing well for their large userbase, and it encourages innovation and improvement within the developer teams.

M. Ethical Concerns of the Recommender System

Despite the benefit and data supporting the personalized artwork for each user, there have been some challenges and ethical concerns with its implementation. In 2018, there were concerns raised about Netflix's algorithms displaying artwork featuring Black actors in minor roles as a form of "racial or demographic targeting" [18] [19] to get users to select or watch the item, even if the advertised actors are hardly in the film. Some Netflix users viewed this as an unethical "bait-and-switch". Netflix claimed that all personalization is purely user behavior-driven, but this controversy still highlights the importance of bias auditing in machine learning models and ethical guidelines for personalization. Not all users are comfortable with this data use and personalized marketing, so it is important for Netflix to balance data gathering and use with user privacy and trust.

N. Autoplay Previews

As Netflix has continued to improve their recommender system, they have introduced autoplay preview clips that play as users scroll through items on the homepage. If the user hovers over an item for an extended period of time, a short clip will play showing the user interesting parts of the film or television show. These clips are another example of their detailed recommender system, as the specific scenes and clips that are shown to the user are personalized based on the user's preferences and behavior. These previews are also optimized using A/B testing, like the personalized item artwork. This is another example of a feature designed to shorten a user's decision time and keep them subscribed to the service. An example of this autoplay preview feature is illustrated in Figure 12.

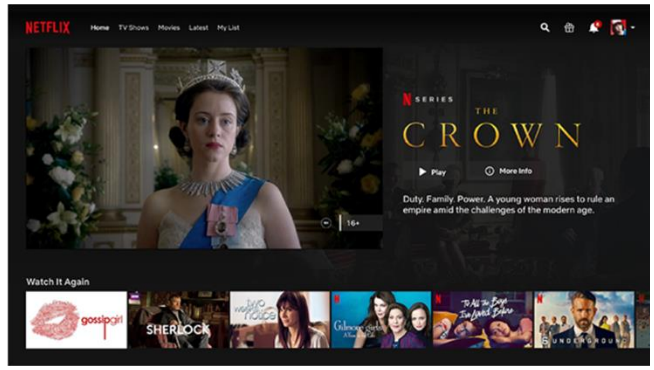

Fig. 12. An example of the Netflix autoplay preview feature to show the user a clip from the show "The Crown" [20].

Netflix's recommender system is a complex blend of collaborative filtering and content-based recommendations. They combine user data with content data to personalize the experience for each user individually. The application design and the item recommendations are built to minimize decision time and maximize watch time. Netflix integrates multiple algorithms to balance responsiveness and efficiency, and they are consistently researching and designing new improvements and innovations for the system. Despite some challenges that they have faced, their understanding that user retention is the most important metric is the reason that Netflix continues to invest in their recommender system. Prioritizing the user has been a key to their success as a streaming service.

IV. COMPARISON AND ANALYSIS

In the era of personalized digital experiences, recommendation systems have become central to user engagement across industries. Among the most prominent, Amazon and Netflix stand out for leveraging advanced machine learning algorithms to tailor experiences that drive loyalty, consumption, and business outcomes. Despite operating in different domains—e-commerce and video streaming—their recommendation systems share common ground in personalization strategies while diverging significantly in objectives, data use, and implementation.

Both Amazon and Netflix have developed sophisticated recommendation systems that track user interactions with products and videos to optimize their individual platforms. Amazon's Deep Scalable Sparse Tensor Network Engine and Netflix's Video-Video Similarity Ranker represent the core of their recommendation algorithms, designed to analyze and predict user preferences based on similar past interactions. Amazon’s Deep Scalable Sparse Tensor Network Engine (DSSTNE), for example, is an open-source deep learning library optimized for sparse data typically found in product interactions, allowing scalable training and prediction in large catalogs. Netflix's Video-Video Similarity Ranker, on the other hand, ranks unseen content by how closely it aligns with videos previously watched, using embeddings that capture both user and content features. This similarity underscores a shared reliance on data-driven strategies to personalize user experience, yet their implementation and end goals separate significantly to align with their business models.

Netflix's recommendation system is finely tuned to enhance user engagement by minimizing the time spent browsing and maximizing viewing time. By strategically placing the most recommended videos in the top left corner of the screen, where viewers' eyes are most likely to focus first, and employing a rotating carousel for seamless browsing, Netflix ensures that users find appealing content quickly. This user-center approach is from machine learning techniques that prevent the redundancy of recommendations, which maintains a fresh and engaging user interface. Netflix also uses a collaborative filtering method that not only considers individual preferences but also joins trends among users in similar regions or time frames, which adds to the recommendation pool.

In contrast, Amazon's system prioritizes increasing user interaction and transaction volumes. While it shares the use of a rotating carousel with Netflix, Amazon's goals are geared towards keeping users engaged in browsing through a vast catalog of products, encouraging them to add more items to their shopping cart. Amazon's recommendation engine, paired with item-to-item collaborative filtering, has evolved to include content-based filtering and hybrid models that use deep learning to handle its massive product assortment efficiently. These algorithms are capable of suggesting items in real-time, using not just user behavior but also item attributes like genre or category, enhancing the relevance of recommendations.

Furthermore, Amazon enhances its recommendations by incorporating a wide array of data, including implicit and explicit user behaviors such as purchase history, search queries, and cursor movements. This dataset is processed through Amazon's neural network models, like autoencoders and neural collaborative filtering, which capture nuanced patterns in user behavior that simple matrix factorizations or older models might miss. This capability allows Amazon to personalize recommendations on their user scale, making the shopping experience highly tailored and responsive to individual user needs.

Ultimately, while both Amazon and Netflix use advanced machine learning and data analysis to power their recommendation systems, the core difference lies in their application: Netflix aims to quickly satisfy content discovery to increase viewer satisfaction and watch time, whereas Amazon focuses on extensive user engagement and purchase behavior to drive sales. Each approach reflects the company's strategic priorities and the specific demands of their service platforms—video streaming for Netflix and e-commerce for Amazon. Both companies continuously refine their algorithms to better understand and predict user preferences, ensuring that their recommendations not only meet but predict the needs and desires of their respective audiences.

V. RESULTS AND DISCUSSION

In the area of streaming and e-commerce, recommendation systems play a huge role in shaping user experience and engagement. Amazon and Netflix, two highly notable technology companies utilize algorithms to tailor their user interfaces and recommendation strategies. Both employ algorithms that track user interaction with products or videos. These systems analyze user behavior to suggest similar items, which enhances the relevance of recommendations and improves user satisfaction.

Netflix's approach to user interface design is crafted to optimize viewer engagement without unnecessary browsing. By placing recommended videos in the top left corner, Netflix ensures that suggestions catch the viewer's eye immediately. The design choice, along with the use of a rotating carousel, facilitates a seamless selection process. This efficiency- first approach is different from Amazon's strategy, which encourages users to linger and browse more.

The core of Netflix's recommendation system lies in its use of machine learning to stop the problem of repetitive content suggestions. By continually refining its algorithms based on user interactions and viewing patterns, Netflix ensures that the recommendations remain fresh and relevant. This not only enhances user satisfaction but also subtly encourages longer viewing sessions. In contrast, Amazon's recommendation engine, while also based on user interaction, lacks this level of intervention. Amazon's model is designed to maximize product exposure and potential purchases by keeping users engaged in an ongoing discovery process.

Furthermore, Netflix utilizes a collaborative filtering approach that goes beyond individual user data. It considers trends among users in similar geographical areas or those watching during similar times, creating a more nuanced and community-oriented recommendation experience. This method allows Netflix to surface content that might appeal based on broader viewing trends, providing a dynamic viewing experience tailored to cultural or regional preferences. Amazon's strategy, however, remains focused on the individual, with recommendations tailored specifically to one's past interactions without integrating broader user trends or community preferences.

These contrasting strategies highlight the different business goals of Netflix and Amazon: Netflix prioritizes keeping viewers engaged with content, minimizing search time to maximize watching time. Amazon, on the other hand, aims to maximize user engagement and time spent browsing through an extensive catalog, enhancing the likelihood of purchases.

Each approach reflects a deep integration of business objectives with technology, showcasing how recommendation systems can be tailored to meet specific commercial aims while shaping the overall user experience.

V. CONCLUSION AND FUTURE WORK

Recommendation systems have entirely reshaped the customer experience for e-commerce and streaming platforms. Amazon and Netflix have profited from implementing different recommendation strategies to help users find what they are looking for easily. The findings show that both e-commerce platforms share dependency on data-driven strategies to improve customer satisfaction. Both platforms use massive behavioral datasets that require multi-tiered computation. On the contrast, their implementation and end goals differ significantly to meet their specific business needs. In this comparative analysis, we have learned that Amazon prioritizes real-time recommendations to increase daily orders and promote incremental sales. It uses hybrid model of filtering algorithms and deep learning to provide real-time product suggestions. On the other hand, Netflix uses a blend of filtering algorithms and machine learning to minimize decision fatigue and increase customer satisfaction. However, the findings of this study cannot be generalized to other industries because this paper focuses on analyzing the recommender systems for only two largest e-commerce platforms. This may not be representative for other e-commerce platforms due to the vast recommendation strategies being used in the industry.

Future research, which expands to impact of emerging technologies such as explainable AI and deep learning in recommender systems to a broader platform, is required to get a better understanding of recommender systems across the e-commerce industry.

REFERENCES

[1] L. Hardesty, “The history of Amazon’s recommendation algorithm,” Amazon Science, https://www.amazon.science/the-history-of-amazons-recommendation-algorithm.

[2] Amazon.com recommendations item-to-item collaborative ..., https://www.cs.umd.edu/~samir/498/Amazon-Recommendations.pdf.

[3] A. Krysik, “Amazon product recommendation system: How does amazon algorithm work?,” Stratoflow, https://stratoflow.com/amazon-recommendation-system/.

[4] Generating recommendations at Amazon Scale with apache spark and Amazon dsstne | AWS big data blog, https://aws.amazon.com/blogs/big-data/generating-recommendations-at-amazon-scale-with-apache-spark-and-amazon-dsstne/.

[5] J. Li et al., “GPT4Rec: A generative framework for personalized recommendation and user interests interpretation,” Amazon Science, https://www.amazon.science/publications/gpt4rec-a-generative-framework-for-personalized-recommendation-and-user-interests-interpretation.

[6] C. Yu and Z. Li, “Building commonsense knowledge graphs to aid product recommendation,” Amazon Science, https://www.amazon.science/blog/building-commonsense-knowledge-graphs-to-aid-product-recommendation.

[7] G. Muralidharan, “A complete study of Amazon’s recommendation system,” Argoid, https://argoid.findableis.com/blog/decoding-amazons-recommendation-system.html.

[8] M. Arsenault, “The Amazon recommendations secret to selling more online,” Rejoiner, https://www.rejoiner.com/resources/amazon-recommendations-secret-selling-online.

[9] J. H. Rystrøm, “Assessing preference manipulation in Amazon’s Recommender System,” HackerNoon, https://hackernoon.com/assessing-preference-manipulation-in-amazons-recommender-system.

[10] A. Staff, “Amazon is earning and Maintaining Customer Trust through privacy,” About Amazon, https://www.aboutamazon.com/news/how-amazon-works/amazon-is-earning-and-maintaining-customer-trust-through-privacy (accessed Apr. 23, 2025).

[11] D. Chong, "Deep Dive into Netflix's Recommender System," Apr. 30, 2020, *https://medium.com/towards-data-science/deep-dive-into-netflixs-recommender-system-341806ae3b48*.

[12] M. Iqbal, "Netflix Revenue and Usage Statistics (2025)," Feb. 21, 2025, *https://www.businessofapps.com/data/netflix-statistics/.*

[13] G. Biddle, "A Brief History of Netflix Personalization," Jun. 1, 2021, *https://gibsonbiddle.medium.com/a-brief-history-of-netflix-personalization-1f2debf010a1.*

[14] Netflix Technology Blog, "Selecting the best artwork for videos through A/B testing," May 3, 2016, *https://netflixtechblog.com/selecting-the-best-artwork-for-videos-through-a-b-testing-f6155c4595f6.*

[15] HelloPM, "Netflix Content Recommendation System - Product Analytics Case Study," Jul, 16, 2024, *https://hellopm.co/netflix-content-recommendation-system-product-analytics-case-study/*.

[16] W. Kweon, S. Kang, J. Hwang, and H. Yu, "Deep Rating Eliciation for New Users in Collaborative Filtering," Feb. 26, 2024, *https://arxiv.org/html/2402.16327v1*.

[17] E. Saberian and J. Basilico, "RecSysOps: Best Practices for Operating a Large-Scale Recommender System," Sep. 30, 2022, *https://netflixtechblog.medium.com/recsysops-best-practices-for-operating-a-large-scale-recommender-system-95bbe195a841*.

[18] L. Bradley, "Netflix Faces Backlash Over How It Targets Users with Artwork," Oct. 23, 2018, *https://www.vanityfair.com/hollywood/2018/10/netflix-artwork-personalization-backlash-racial-profiling#:~:text=In%20a%20statement%2C%20Netflix%20essentially,Netflix%20experience%2C”%20a%20representative%20told*.

[19] D. C. Chmielewski, "Netflix's Use Of Artwork Personalization Attracts Online Criticism," Oct. 22, 2018, h*ttps://deadline.com/2018/10/netflixs-artwork-personalization-attracts-online-criticism-1202487598/*.

[20] H. Labuschagne, "Netflix now lets you disable autoplay previews," Feb. 7, 2020, *https://mybroadband.co.za/news/broadcasting/337874-netflix-now-lets-you-disable-autoplay-previews.html*.