WHITE PAPER
SERIES

2025

Definitive Guide to Recommendation Engines

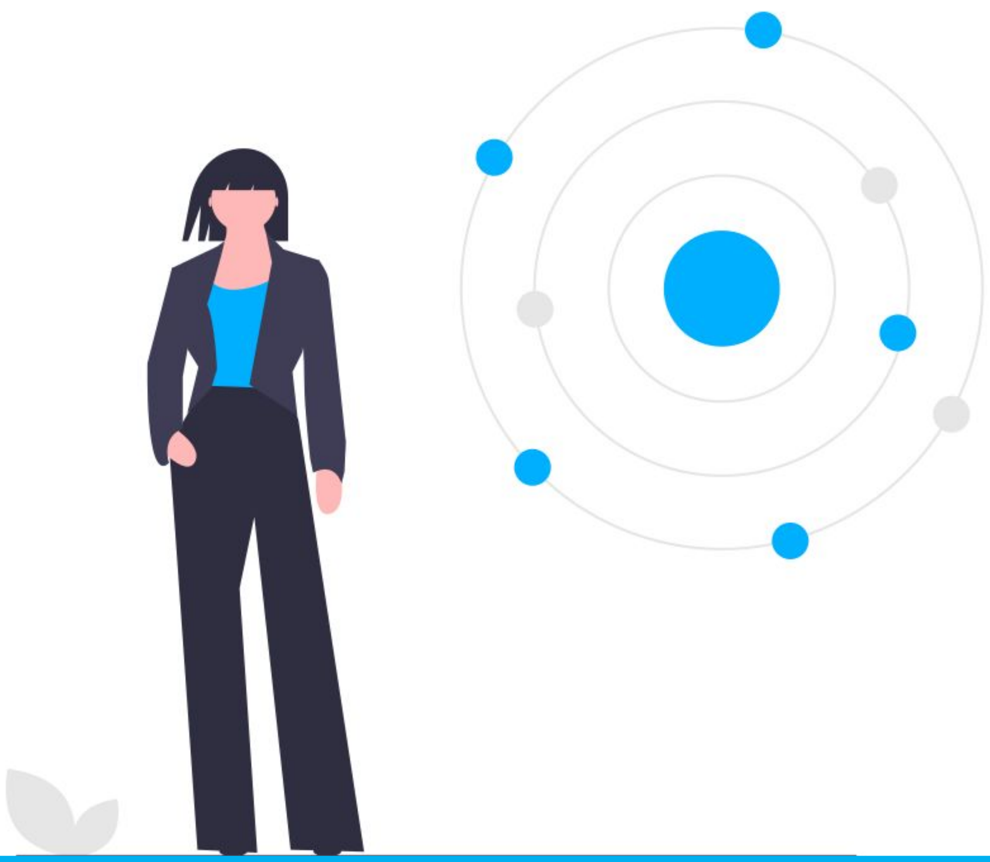

adsp.ai

Introduction

What are they? Why are they used? And why are they so popular & More!

The good news: the internet is full of products and services that may be exactly what you are looking for.

The bad news: the internet is full of products and services that may be exactly what you are looking for.

How can you make the right choices without exhaustively searching through all options?

This is where recommendation engines come into play. Recommendation engines are data science tools that predict which products or services are best suited for each individual customer.

What are the important things to know before successfully adopting recommendation engines in your business?

We will see that over the next ten chapters of this guide.

By the end of it, you will be able to effectively communicate about why and how a recommendation engine should be added to your customer interface.

We will also see which companies and applications champion the recommendation engine scene.

David Foster
Partner
Applied Data Science Partners
(ADSP)

Focus Areas

This whitepaper series is a practical guide to developing your understanding of Recommendation Engines.

It is organised around 10 key topics:

Chapter One:
What is Product Recommendation?

Chapter Two:
What Is A Recommendation Engine?

Chapter Three:
What Are The Types of Recommendation Engines?

Chapter Four:
How Does A Recommendation Engine Work?

Chapter Five:
How Recommendation Engines Are Used?

Chapter Six:
Challenges of Recommendation Engines

Chapter Seven:
Applications Of Recommendation Engines

Chapter Eight:
How to Solve Recommendation Engines Problems

Chapter Nine:
Why Are Recommendation Engines Becoming Popular?

Chapter Ten:
Recommendation Engine FAQs

1 What is Product Recommendation?

A product recommendation is an item that you proactively present to your customer because you believe they are likely to buy.

Product recommendations do not need to always lead to a purchase.

They are successful if they recommend interesting products.

In fact, having 100% successful recommendations is bad; it means that you are playing it too safe and are missing out on potential preferences of your customers.

2 What Is A Recommendation Engine?

A recommendation engine is a software system that outputs product recommendations.

It does so by processing data generated by the customers, building a predictive model of their behaviour and outputting the items most likely to be desirable.

Recommendations are not restricted to products.

They can refer to services, content such as movies on Netflix and connections on social media.

One example of a recommendation engine is Amazon's "Customers Who Bought This Item Also Bought" feature.

This feature shows customers items that other customers who bought the same item also bought.

3 What Are The Types of Recommendation Engines?

There are three main types of recommendation engines: collaborative filtering, content-based filtering – and a hybrid of the two.

Collaborative Filtering

Collaborative filtering makes recommendations to a user based on their similarity to other users.

Its main element is the preference matrix, where each user is a row and each product is a column.

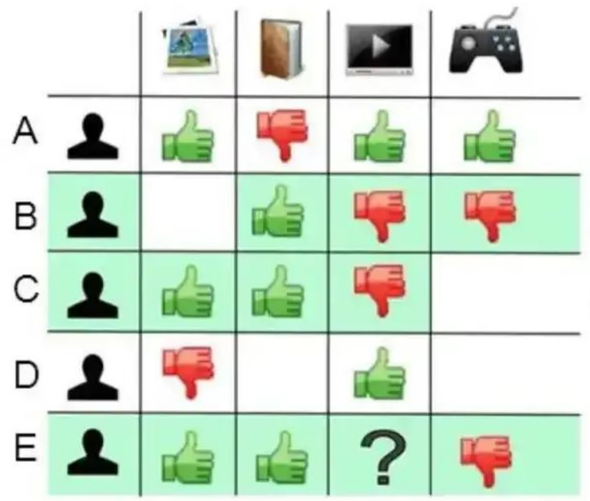

If you are user E, then probably you don't want a TV, because users B and C are the most similar to you and they also don't want a TV.

An advantage of collaborative filtering is that it doesn't need to analyse or understand the content (products, films, books).

It simply picks items to recommend based on what they know about the user.

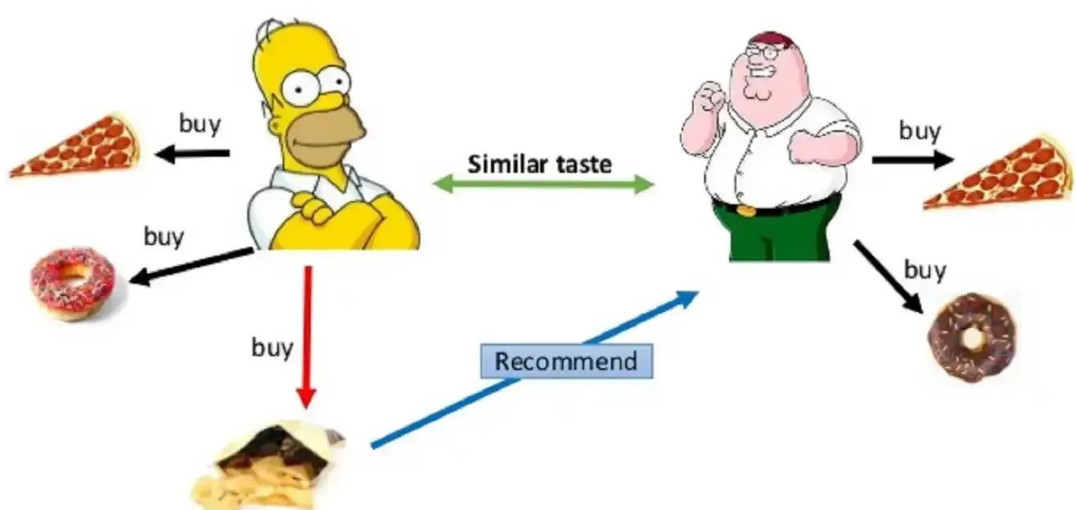

Content-based filtering

Collaborative filtering makes recommendations to a user based on their past preferences. If you liked the Squid Game series, then you will probably also like the Parasite movie.

To make recommendations, algorithms use a profile of the customer's preferences and a description of an item (genre, product type, colour, word length) to work out the similarity of items.

The downside of content-based filtering is that the system is limited to recommending products or content similar to what the person is already buying or using.

It can't go beyond this to recommend other types of products or content. For example, it couldn't recommend products beyond homeware if the customer had only brought homeware.

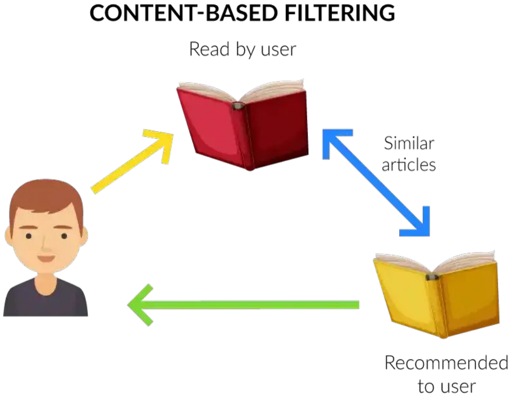

Hybrid Model

A hybrid recommendation engine makes recommendations based on both the users' similarity to others and their past preferences.

For example, one can generate a recommendation with both techniques and choose to show one of the two or both of them.

Netflix is the perfect example of a hybrid recommendation engine.

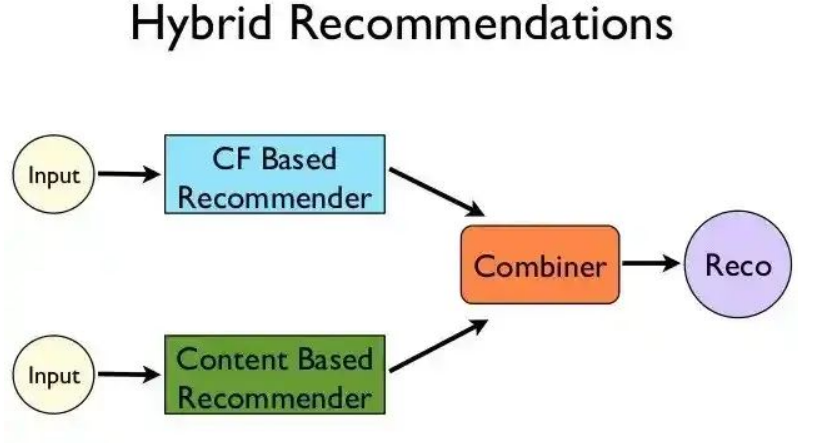

4 How Does A Recommendation Engine Work?

A recommendation engine works using a combination of data and machine learning technology.

The more data it has, the more efficient and effective it will be in making relevant revenue-generating suggestions.

Recommendation engines complete a standard four-step process:

Step 1: Data Collection

There are two main types of data to be collected:

Implicit Data:

This includes information collected from activities such as web search history, clicks, cart events, search log, and order history.

Explicit Data:

This is information gathered from customer input, such as reviews and ratings, likes and dislikes, and product comments.

If you are using collaborative filtering, you will also need customer attribute data, such as demographics (age, gender) and psychographics (interests, values).

If you are using content-based filtering, you will instead need product attribute data, such as genre and colour.

Step 2: Data Storage

Once the data is gathered, it needs to be stored. Over time, the amount of data will grow to be vast.

This means ample, scalable storage must be available.

Depending on the type of data you collect, different types of storage are available.

Step 3: Data Analysis

Depending on how often you process your data you can refer to your data analysis as:

Real-time: new recommendations generated each time new data is produced

Batch processing: new recommendations generated e.g. every few hours
Near-real-time: new recommendations generated e.g. every few minutes

Step 4: Data Filtering

This step refers to the transformation of your data into recommendations.

The type of recommendation engine you picked in Chapter 3 will determine how this filtering is performed.

5 How Recommendation Engines Are Used?

What are valid reasons for recommending products?

Recommendations need to arise from specific needs of your customers.

For example, having a surplus of a product does not mean you should recommend it more often.

Here are some valid reasons for producing recommendations:

Providing cross-selling opportunities

Your customer came to buy a smart-phone, but they don't know you also offer the smart-phone case they may look for later. This type of recommendation is the most likely to hit home.

Addressing cart abandonment

I came, I saw, I forgot to finalise my order.

Unfortunately customers often unwillingly interrupt a purchase and then forget about it. Or maybe they discover that a product is missing a certain feature and willingly abandon the purchase.

In both cases, recommending products similar to what lingers in the abandoned cart may prove useful.

Offering alternatives

As a seller, you are the entity most qualified to discover alternatives to a customer's choices.

This is one of the main reasons why customers are attracted to online shopping: the combination of high efficiency and good market research is simply unbeatable when compared to traditional shopping practices.

6 Challenges of Recommendation Engines

Recommending can prove more challenging than other classical business-oriented data science tasks, such as customer segmentation and market basket analysis.

Here are some of the unique technical challenges of recommendation engines:

The Cold Start Problem

Every new item or user added to your database is a blank state.
This makes kickstarting recommendations for them particularly challenging.

Luckily one can leverage existing items or users that are similar to the newly added ones to avoid random recommendations.

This is why customer and product attribute data are very important: they can help build a profile even in the absence of behavioural data.

Data Sparsity Problem

Remember the preferences matrix we saw in Chapter Three?
That matrix was deceptively dense.

Real matrices generated by user data tend to be very sparse, as customers tend to only rate a small portion of items.

To put things into perspective, the rating sparsity of recommender systems is usually up to 99%.

This means that calculating user similarity is quite difficult.

Changing User Preferences Problem

Consumers do not stand still – they are constantly behaving and evolving both as people and customers.

Staying on top of these changes is a constant battle.

A strong recommendation engine will be able to identify changes (or signs of an impending changes) in customers' preferences and behaviour, and constantly auto-train themselves in real time in order to serve relevant recommendations.

7 Applications Of Recommendation Engines

Recommendation engines can prove useful in many industries.

1. E-Commerce and Retail

Online retailers can use recommendation engines to suggest products to customers based on their previous purchases.

This can help increase sales and customer satisfaction by providing personalised recommendations.

2. Social media

Social networks can use recommendation engines to suggest friends, groups, and pages to users based on their interests. This can help users find new friends and groups, and discover new content.

3. Banking

Recommendation engines can prove useful in online banks in different ways:

Offering personalised products and services:

Banks can use recommendation engines to offer personalised products and services to their customers.

For example, a customer who frequently uses their credit card for online shopping could be recommended a credit card with a higher credit limit or cashback rewards.

Cross-selling and upselling products:

Recommendation engines can also be used to cross-sell and upsell products to customers. For example, a customer who has a savings account with a bank could be recommended a term deposit or a home loan.

Improving customer retention:

Banks can use recommendation engines to improve customer retention. For example, a customer who is about to close their account could be recommended a different account with better interest rates or fees.

Identifying at-risk customers:

Recommendation engines can also be used to identify at-risk customers. For example, a customer who has missed several loan repayments could be recommended a debt consolidation loan.

4. Telecom

A telecom provider heavily relies on advertising campaigns to increase revenues and generate user loyalty; in other words, they aim to increase the average amount a user spends per service and per month and decrease churn rates. Using a recommendation system, a provider can send a notification of any marketing campaign to users.

These marketing notifications can be sent to customers when certain events are triggered, such as their balance reaching an amount that is lower than a specific limit or their average credit, or when entering into a particular area or location, like a new town or shopping mall.

8 How to Solve Recommendation Engines Problems

The challenges we saw in chapter 6 were just the tip of the iceberg. Your business may discover its own unique problems on its journey to adopting recommendation engines.

For example, privacy concerns and initial investments may prove prohibitive at the beginning.

How can you be certain that you can address any problem that may arise?

There is no one-size-fits-all answer to this question.

The best way to solve a recommendation engine's problem depends on the specific problem at hand.

However, some tips on how to solve a recommendation engine's problem include understanding the user's needs and goals, understanding the data, and using a variety of techniques to find the best solution.

9 Why Are Recommendation Engines Becoming Popular?

The Recommendation Engine market was valued at USD 2.12 billion in 2020, and it is expected to reach USD 15.13 billion by 2026, registering an impressive increase in its market share of 37.46% during the period of 2021-2026.

While this increase is impressive, it is not surprising.

Recommendation engines are tools that impact business in many direct and meaningful ways:

- They increase revenue
- They improve customer satisfaction
- They offer a personalised user experience
- They help users explore and discover
- They can be used as reporting tools

10 Recommendation Engine FAQs

1. What are recommendation engines?

A recommendation engine is a software system that outputs product recommendations. It does so by processing data generated by the customers, building a predictive model of their behaviour and outputting the items most likely to be desirable.

2. What is an example of a recommendation engine?

One example of a recommendation engine is Amazon's "Customers Who Bought This Item Also Bought" feature. This feature shows customers items that other customers who bought the same item also bought.

3. What is a recommendation engine in machine learning?

Machine learning algorithms in recommender systems are typically classified into two categories — content based and collaborative filtering methods although modern recommenders combine both approaches. Content based methods are based on similarity of item attributes and collaborative methods calculate similarity from interactions.

4. What is the best recommender system?

There is no one-size-fits-all answer to this question.
The best way to solve recommendation engine problems depends on the specific problem at hand.

5. Where are recommendation engines used?

Recommendation engines can prove useful in many industries, such as e-commerce, e-banking, social media and telecom.

6. Challenges of Recommendation Engines

The major challenges of Recommendation Engines are the Cold Start problem, the Sparsity Problem and the Changing User Preference problem.

7. Why are Recommendation Problems becoming popular?

Recommendation engines are tools that impact business in many direct and meaningful ways, such as increasing revenue and customer satisfaction.

11 Bonus Chapter - Top Companies Using Recommendation Engines

Where can you turn for inspiration on your next project?

While your recommendation engine needs to be bespoke, you can learn a lot from companies that have set an example of how you can make the most out of your recommendation engine:

Amazon

If you were to leave this tutorial knowing one Recommendations leader, this should be Amazon.

Amazon is well-known for its extremely accurate selection of recommendations in the online store, such as the "Customers who also bought this item" feature. It uses technologies such as artificial intelligence algorithms and machine learning to improve customer experience and increase sales.

Neflix

The Netflix Recommendation Engine (NRE), is made up of algorithms which filter content based on each individual user profile. The engine filters over 3,000 titles at a time using 1,300 recommendation clusters based on user preferences. It's so accurate that 80% of Netflix viewer activity is driven by personalised recommendations from the engine. It's estimated that the NRE saves Netflix over $1 billion per year.

Spotify

At the centre of the Spotify recommendation system is the home screen which is peppered with many customised playlists and recommendations. The recommended playlists include Discover weekly, B Side, Release Radar, your mixtapes, and many more. Other sections of the cleverly arranged home screen are Jump Back In, Recently Played, or Recommended for today. The home screen is created and curated by an AI called BaRT ("Bandits for Recommendations as Treatments").

LinkedIn

The job market is all about recommendations.
LinkedIn uses content matching and collaborative filtering to recommend companies or jobs a user might be interested in. LinkedIn collects data on positions, education, summary, specialty, experience and skills of a user from the profile itself.

YouTube

Youtube users encounter recommendations at two places — one on the homepage as they enter YouTube which is a mix of content based on past viewing and subscriptions as well as latest news. Then they also see recommendations in the "Up Next" panel as they are watching a video.

YouTube takes signals from a mix of user behaviours which are strong indicators that they like a video. So clicks on videos, watch-time and shares are taken as good cues.

TikTok

Looking for recommendations engines that keep you hooked? TikTok's engine recommends content by ranking videos based on a combination of factors — starting from interests you express as a new user and adjusting for things you indicate you're not interested in, too.

Instagram

Instagram uses two different recommendation engines: the Home Feed Ranking System ranks the posts from the sources you follow based on factors like engagement, relevance, and freshness while the Explore Ranking System opens you up to many other public posts which might be relevant and engaging to you.

Over half of the Instagram community visits Instagram Explore every month to discover new photos, videos, and Stories relevant to their interests.

Facebook

Wondering how recommendations work at scale?

Facebook's average data set has 100 billion ratings, more than a billion users, and millions of items. They are dealing with numbers at least two orders of magnitude larger than anyone else is dealing with. Facebook uses an advanced distributed algorithm based on collaborative filtering and data science.

Tinder

Tinder's personalised recommendation engine is called TinVec. TinVec embeds users' preferences into vectors leveraging on the large amount of swipes by Tinder users. With 26 million matches per day and more than 20 billion matches made to date, Tinder is the world's most popular app for meeting new people.

Google
Where did it all start?

Google's PageRank (PR) is the grandfather of recommendation engines, originally developed to rank web pages in their search engine results.

Today Google uses its bespoke Recommendations AI platform, which AI draws on that experience and expertise in machine learning to deliver personalised recommendations that suit each customer's tastes and preferences across all your touchpoints.