



THE ENTERPRISE LEADER'S GUIDE TO

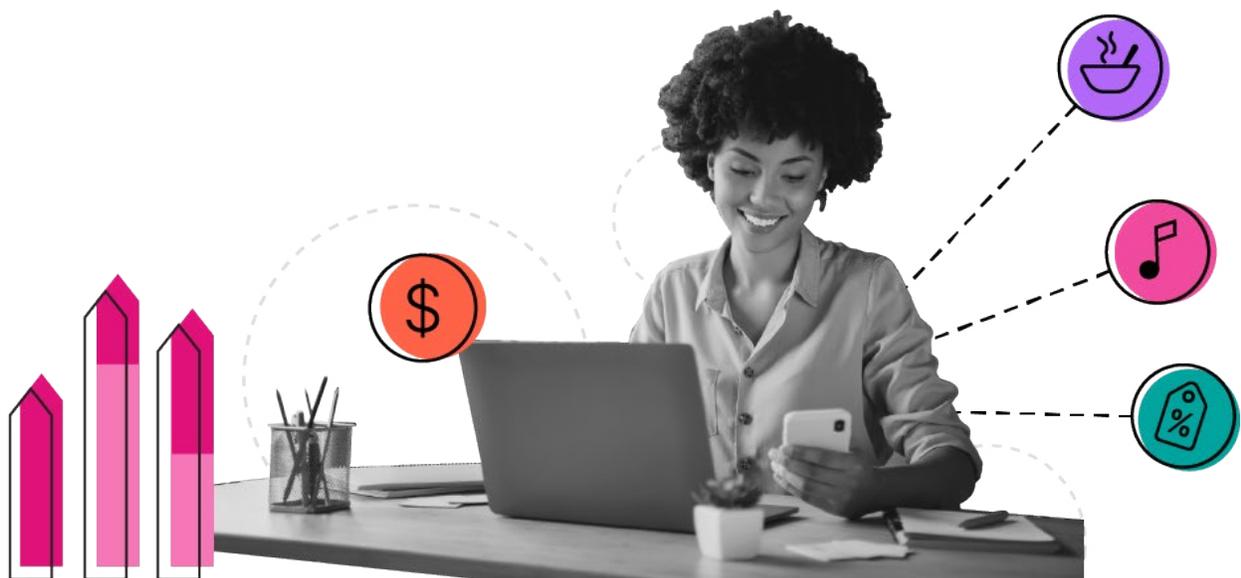
Understanding & Implementing Recommendation Engines

How to determine the right
approach for your scalability
and success

In any given week, millions of consumers make use of recommendation engines (also frequently called recommender systems or personalization systems) to watch movies and TV shows, discover new music, order food online, and see tailored content on social networks. Organizations of all sizes and industries are taking notice of the impact these systems, powered by machine learning, can have in matching their consumers with what they're seeking. **In fact, the global recommendation engine market size has the potential to grow by \$3.57 billion from 2020-2024.**

Amazon, for its part, generates over a third of its total revenue (35%) from its recommendation engine, according to McKinsey, while Netflix, another massive leader in recommendation systems, reports that its recommendation engine drives 75% of viewership. Given their success, it's understandable that many businesses look to Amazon and Netflix as companies to emulate when it comes to implementing and scaling recommendation systems. However, this can breed the misconception that only large-scale enterprises can make good use of recommendation systems. The reality is that organizations, both large and small, can use these systems to create a better experience for consumers while also helping employees make more-informed decisions — a true competitive advantage.

The problem is most leaders often don't know where to begin. In fairness, the task can be overwhelming, especially as you consider the technical undertaking. But it's necessary to find the right recommendation system for your particular enterprise, team structure, and data needs, and implement it well in order to drive impact both for your users and at every level of your business.





01- The Foundation of Recommendation Engines



We are leaving the age of information and entering the age of recommendation.

– Chris Anderson, Kitty Hawk, COO



Why knowing the parts of a recommendation tool is essential to selection

- Makes clear what you need from your users for a successful recommendation engine
- Can help to simplify the eventual implementation of your system
- Helps to better understand your users and determine what roadblocks you may face

What Are the Basic Pieces and Functions of a Recommendation Engine?

At its most basic level, a recommendation engine is a “black box” that relates two typically unlike entities, forming a relationship between the two entities. The domain of operation for these systems or recommendation engines is generally quite broad with a lot of different terminology (‘engine’ vs. ‘system’) and different outcomes businesses are looking for in terms of achievement. The crux of these systems though, is going through a process of collection, storage, analysis and filtering of user data to form a relationship typically between a user and a product.

In RevUnit’s experience, this specifically relates to work in retail and transportation logistics, but we want to make it clear that these systems can be applied to almost any organizational goal. For example, RevUnit also did work with a virtual fortune cookie message app whose aim was to provide personalized, positive mantras to people. Relating people and fortunes instead of people and products, RevUnit generated fortunes and rated them on attributes. For every person who registered, they would answer questions and narrow down their attribute profile (uplifting, adventurous, aspirational content, etc.) to provide relevant fortunes to them.



A recommendation engine could even be built to predict which houses in a neighborhood have dogs based on some features of dogs and houses. It really doesn't matter so long as the two axes (sometimes there can even be a third, more on that later) have some latent characterization and there is a natural relationship between them. That example might be a little niche, but it hopefully communicates how broadly recommendation engines can be applied.

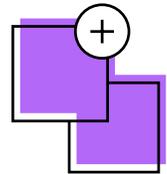
What You Need to Know – Your Products and Your Customers

When it comes to the initial steps of creating a recommendation engine, the general collection of attributes around your two entities has always been the traditional approach. We mentioned this before, but it's possible for there to be a third axis in addition to your consumer and the product you are trying to connect them with. What does that look like?

In RevUnit's work with a moving company, they needed help in understanding their user demographics, the type of layout and furniture their users have, and what their user would consider an ideal neighborhood in order to help them find the best one to move to. We developed a digital tool that could logically suggest one of the elements when the other two were present. When a user selected homes they liked, the app would detail neighborhoods where those types of homes exist. When users knew their ideal neighborhood, it would recommend homes that were laid out and sized right for their furniture, and so on. When the user provided two aspects, the third was presented to them.

Still, recommendation engines most commonly apply to two entities, especially within the retail industry, whose focus is relating customers to products. As a business you need to know something about the customer and something about the products from the start. This "something" you know could be any attribute like gender, age, location, etc. With products, it can also be attribute-based by category, color, size – anything. With this in place, recommendation engines can say, "This user has a similar collection of attributes, let's direct them to this group of attributes".

The challenge with this is that it can feel invasive to your users to give up their demographic information, especially if you haven't established much of a relationship with them. You need to be able to devise a way to create a connection between customers and products that is more organic. This is collaborative filtering, which we'll get into later.



02- Getting Started (Getting Your Pilot Off The Ground)



A great experience at a brick-and-mortar store happens when an employee knows you as a person and knows what you like in terms of products, and then can recommend other things you might like. This same idea can be translated into digital products when organizations know you, know what you like, and can recommend things you may never have found on your own. That's the power of recommendation engines.

– Colin Shaw, Director of Machine Learning, RevUnit



Why starting on the right foot is critical to the success of your recommendation tool

- Allows you to better measure and understand consumer data
- Can help to determine what success of your specific system looks like
- Helps illuminate what approach to recommendation engines may best suit you

Where Do You Start?

The process of devising a recommendation tool can feel daunting at first, for large enterprises and smaller businesses alike, so it's important that you start simply, and work smart in regards to your data. Ask the customer what things interest them and start putting together an idea of their profile (e.g. "are you interested in...?"), framed up to look like a personalized experience.

With your users, a few starter questions will help you determine what exactly they're looking for – house products or food products, for example. Then, as they buy and review products you can continue to refine your recommendations. Sometimes this is called a "learning-all-the-time system" or "online learning system" where the person and product are updated all the time as new information is gleaned from them.

This takes you from just simply giving a person random items to having a better understanding of your users' preferences so you can tailor their experience to what they're most interested in.



General Advantages

To talk about the benefits of recommendation engines, you first need to understand what success hinges upon. Typically, they're dependent upon understanding your domain well, ensuring relationship scoring makes sense, and using a modern approach for scalability, speed, and online learning capability. In this way it can be difficult to quantify the benefits of recommendation engines, but it's certainly easy to say so qualitatively.

Let's start with the obvious thought – hyper personalization. If you can give people options that are relevant and personalized to them, they're simply more likely to purchase them compared with random items. Amazon and Netflix in particular have been game changers in their industry and it seems like part of that is due to this type of recommendation.

We did mention it can be hard to make a quantitative case for this —how do you know they wouldn't have searched using keywords and found the same item? We often don't know the definitive reason why someone purchases something, but there's a strong case for personalization given the staggering amounts of revenue large-scale companies like Amazon and Netflix are reporting from using recommendation engines.

Also, to the same extent that someone working beside you in a brick-and-mortar recommending relevant items would be considered good customer service, this same idea is simply translated online with recommendation engines.



[CASE STUDY]

TRANSPORTATION & LOGISTICS

A leading transportation and logistics company worked alongside RevUnit to pilot a carrier and lanes recommendation engine. Each lane is distinct based on origin and destination (LA to Memphis is one lane, Memphis to LA is a different lane), and while our client was the main carrier the tool was developed for, they work with many single-owner carriers and small businesses with two to three trucks that also pick up routes. Certain trucking companies had a few specific lanes they liked to run and often didn't venture out of those zones.

RevUnit developed a recommendation system for new lanes that drivers may not typically take, but they might consider if it was presented to them. As an example, with RevUnit's recommendation engine, a trucker located on the southside of Los Angeles is shown a requested route from San Diego to Memphis – this trucker likely never would have searched for this route independently, but they may consider it when it makes sense logistically and financially for the carrier.

While RevUnit found success with this particular recommendation engine, driver feedback was strictly in regards to UX (cleanliness, amount of rest stops along route, etc.), so it became clear this 'best route' approach needed to be geared specifically to the carrier business operator (focused on revenue and profit) for the appropriate scoring system. Determining what exactly is a route and carrier proved more difficult than defining a product and a customer, which brought up an important truth of recommendation systems – the further you get from the product-customer dynamic, the more likely the need for consideration in your system.



03- Approaches to Recommendation Engines - Big vs. Small



We want our recommendations to be fresh and responsive to the actions a member takes, such as watching a show, adding to their list, or rating; but we also want some stability so that people are familiar with their homepage and can easily find videos they've been recommended in the recent past.

– Chris Alvino and Justin Basilico, Netflix



How determining the right recommendation engine benefits organizations and their users

- Answer how to solve issues like the cold start problem for new users or products.
- Answer the question of whether you can scale with your current system
- Determine whether ad hoc methods work or modernization is required

Choosing the right recommendation engine for your business

When considering what the right recommendation engine looks like for your business, you can easily begin to build an understanding between things when you start to connect people based on their attributes. For example, you can present your users with products that resemble other products with a collection of attributes that fall in line with their preferred attributes. This approach is particularly efficient and effective for smaller businesses who may have the time to come up with all the attributes for all of their products, but certainly not for larger organizations.

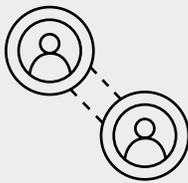
The traditional approach to recommendation engines that enjoys wide use among some of the largest and most widely-known companies (Amazon, Netflix, LinkedIn, Spotify, and Youtube, to name a few.) is **collaborative filtering**. Collaborative filtering functions on the assumption that people who agreed in the past will agree in the



future, and that they will like similar kinds of items as they liked in the past. It's best to think of collaborative filtering like a matrix where you have something that can relate to a person and a product. For example, with Netflix, you are related to other shows and movies by your ratings (ex: "I thought this movie was 4 out of 5 stars.") which allows Netflix to have a clearer understanding between a person and the products they seem to prefer.

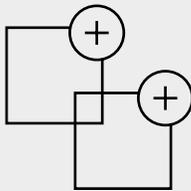
A key advantage of the collaborative filtering approach is that it does not rely on machine analyzable content and therefore it is capable of accurately recommending complex items such as shows or movies without requiring an "understanding" of the item itself. The collaborative filtering algorithm is not dependent on any additional information other than user behavior. Below we've broken down some of the types of collaborative filtering:

TYPES OF COLLABORATIVE FILTERING



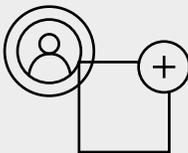
User-User

This searches for strong similarity between users. Based on a determined similarity score, the engine then picks out the users it identified as being the most similar and recommends products which these similar users have shown an interest in or purchased previously. Using Netflix as an example, this form of filtering finds the similarity between each user based on the ratings they have previously given to different shows and movies.



Item-Item

In our Netflix example above, we will find the similarity between each movie pair and based on that, we will recommend similar movies which are liked by the users in the past. This algorithm works similar to user-user collaborative filtering with just a little change – instead of taking the weighted sum of ratings of "user-neighbors", we take the weighted sum of ratings of "item-neighbors".



User-Item

The user-product relationship occurs when some users have an affinity or preference towards specific products that they need. Using a sports equipment e-commerce site as an example — a cricket player might have a preference for cricket-related items, thus the e-commerce website will build a user-product relation of player->cricket.



Another aspect of collaborative filtering systems is the ability to generate more personalized recommendations by analyzing information from the past activity of a specific user, or the history of other users considered to be similar to the tastes of other users. This form of user profiling helps sites recommend content on a user-by-user basis. The more a given user makes use of the system, the better the recommendations become, as the system gains data to improve its model of that user.

Essentially, a collaborative filtering system helps businesses learn more about their users and user preferences of their products in real-time, and more users and items can be added without trouble. This is the standard approach now with the majority of large organizations who encounter scaling issues, but it's not very well understood as a solution type.

Roadblocks and hazards of each approach

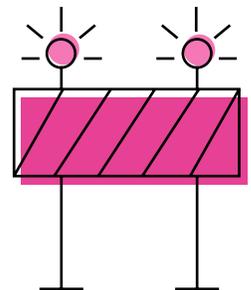
Making use of collaborative filtering puts a lot of predictive power into your recommendation engine, but it also has a few roadblocks and limitations in particular situations that we should go over.

In practice, many commercial recommendation engines are based on large datasets. As a result, the user-item matrix used for collaborative filtering can be incredibly large and sparse. As collaborative filtering methods recommend items based on users' past preferences, new users will need to rate a large amount of items for the system to capture their preferences in order to accurately provide reliable, personalized recommendations. This can also be solved by looking at proxies to ratings, such as which items new users actually looked at after searching.

New inventory suffers from the same problem. When items are added to the system, they need to be rated by a substantial number of users before they are recommended to users who have similar preferences to the ones who rated them. So if you don't have a considerable number of users, this will be a considerable problem. Similarly, for items that didn't receive much data, the model tends to give them less weight and creates a popularity bias by recommending more popular items.

Additionally, as the amount of users and items grow, collaborative filtering becomes more susceptible to serious scalability issues. With tens of millions of customers and millions of items, the complexity of your collaborative filtering algorithm can become too large to manage.

Unfortunately, the fact that systems need to react and make recommendations immediately for all users regardless of their purchases and ratings history also necessitates higher scalability. Large web companies such as Twitter use clusters of machines to help with scaling recommendations for their millions of users, with most computations happening in very large memory machines.





Wrapping Up

It's up to you and your data teams to determine what kind of approach to recommendation engines is right for your organization, but you also need to make an accurate assessment of whether your organization is set up to succeed with your particular choice of recommendation engine at all.

This is a truly important question organizations need to ask themselves. While we've detailed enough that a recommendation engine is a fantastic idea in theory, when rolled out haphazardly or without all of the necessary elements in practice, businesses are bound for frustration.

Will your recommendation engine address your needs and satisfy your customers? Do you have adequate information about your customers and products? Will you be able to create a functioning rating system so you can learn more of your users than whether they just bought a product or not? How do you account for returned products in this rating system? Organizations will find themselves set up for success far more easily when they have nuanced information about how much customer X likes products like Y. Good approaches exist for both large and small companies when it comes to recommendation engines, but to avoid issues like sparsity and the cold start issue that are inherent to certain approaches, you need to adequately determine which one is right for your business.



VIDEO SERIES

WHAT ARE DATA SILOS?

What are silos, where do they come from, and how do you remove them? In this video series, we answer these questions and more based on our experiences working with some of the world's largest companies.

[Watch the series →](#)



ARTICLE

ARE YOU OVERTHINKING DATA INTELLIGENCE?

Where should business leaders start when it comes to AI/ML? Danny Estavillo shares a few simple steps to help you transcend your understanding of data intelligence.

[Read the insights →](#)



GUIDE

APPLYING AGILE TO SPEED UP YOUR DATA PROJECTS

Whether your data project involves building a data model, a data visualization, applying machine learning, or data integration, taking an agile approach will get you more successful results, faster.

[Learn more →](#)





REVUNIT

RevUnit is a data technology studio that helps retail, transportation, and logistics organizations create change with their data — faster. Through disruptive thinking, RevUnit has helped Fortune 500 companies like Walmart, Tyson, and J.B. Hunt reimagine their data systems to improve operations.

revunit.com →

Get in Touch

Have a question, comment, or something we may be able to help with?

[Drop us a line](#) →

