

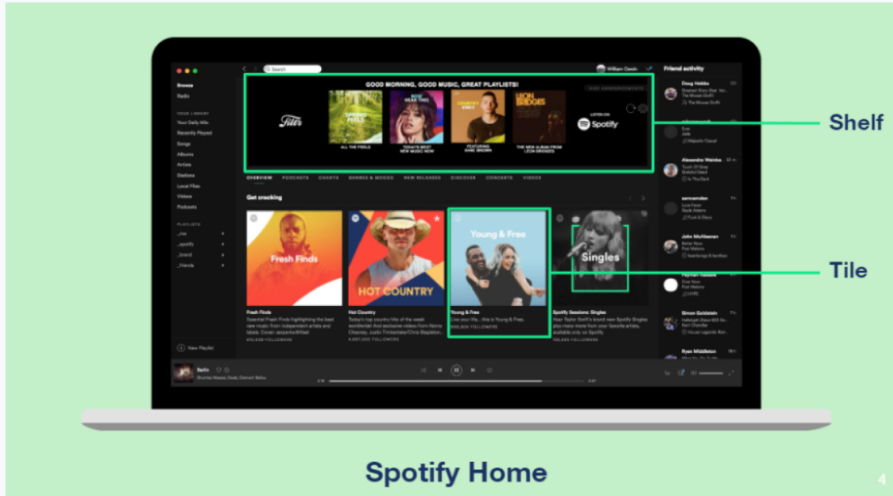


For Your Ears Only: Personalizing Spotify Home with Machine Learning



January 16, 2020

Published by Spotify Engineering



This article is based on the keynote given by Tony Jebara at TensorFlow World in Santa Clara, California, October 2019. You can watch the presentation [here](#).

Machine learning is at the heart of everything we do at Spotify. Especially on Spotify Home, where it enables us to personalize the user experience and provide billions of fans the opportunity to enjoy and be inspired by the artists on our platform. This is what makes Spotify unique.

Across our engineering community, we are working to unite autonomous teams and empower them to be more efficient by establishing best practices for tools and methods. Our recent adoption of a standardized machine learning infrastructure provides our engineers with the environment and tools that enable them to quickly create and iterate on models. We call it our 'Paved Road' approach, which includes utilizing services from TensorFlow, Kubeflow, and the Google Cloud Platform.

Before coming to Spotify, I worked on personalization algorithms and the home screen at Netflix. My previous experience is quite similar to the work I'm doing now at Spotify as the Vice President of Engineering and Head of Machine Learning. However, personalizing Spotify's Home screen presented a new set of challenges, which I'll share later.

Machine learning enables us to recommend artists, playlists, and podcasts so that users are active and more likely to subscribe in the long-term. Great in theory, but what enables us to do this better than our competitors? After all, we're not the only content streaming platform attempting to build a customized and engaging landing page that provides a unique value for our subscribers.

Power of personalization

Like Netflix, we organize our Home screen using a series of cards and shelves. A card is a square image that represents a playlist, podcast episode, artist page, album, and so on. Shelves are the rows we use to group a series of cards. Think of it as how a bookcase (Spotify Home), uses bookshelves (shelves) to hold and display books (cards).

Running with that analogy, each person's bookcase is uniquely curated by their interests and reading history of the books they collect over time. However, unlike a physical bookcase, Spotify uses machine learning to personalize the shelves and cards based on the content they previously enjoyed or might enjoy, and present it to millions of users.

We are able to create this level of personalization in real-time for our **248 million** monthly active users (MAU's), by combining the power of machine learning with their listening history, musical choices, duration of play, and willingness to act on recommendations. We like to say there is no 'one' true Spotify. Essentially, there are 248 million versions of the product, one for every user!

SHARE




CATEGORIES

Machine Learning

Related articles




It's All Just Wiggly Air:
Building Infrastructure to
Support Audio Research



Spotify Podcast Dataset
and TREC Challenge 2020

**Introducing the Spotify
Podcast Dataset and TREC
Challenge 2020**



Spotify Home
Shortcuts

**Reach for the Top: How Spotify
Built Shortcuts in Just Six
Months**

@SpotifyEng on Twitter

Tweets by @SpotifyEng



Spotify Engineering
@SpotifyEng


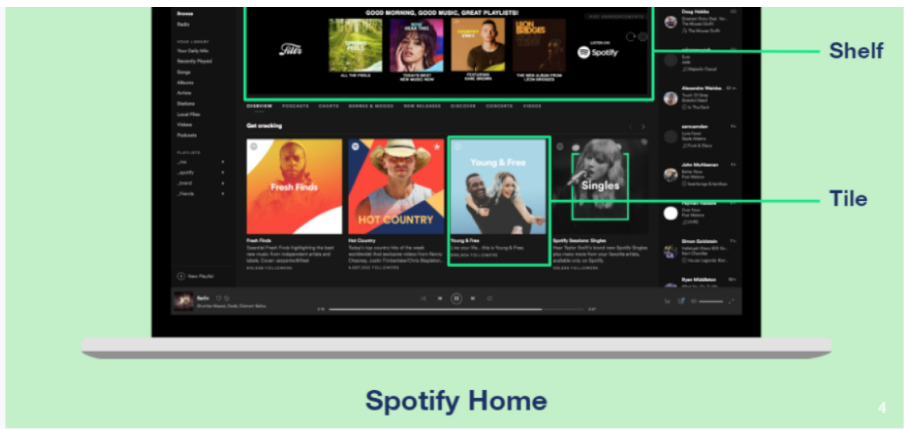
The recipe for efficient, scalable audio pipelines? #Python, @ApacheBeam and a cloud. Read more about [Klio.io](#), our open framework for processing media at scale, on the [@GCPcloud](#) blog featuring [@roguelynn](#): cloud.google.com/blog/products/... <https://twitter.com/GCPcloud/status/1334303207107989511>



Simplify creating data pipelines ...
Spotify open-sources Klio: scalabl...
cloud.google.com



16m



Engagement and research-based recommendations

From a machine learning perspective, we use a multi-armed bandit framework that balances **exploitation and exploration**. Beyond the engineering community, exploitation can have a negative connotation. However, in the recommender systems used to create personalized content experiences, exploitation means providing recommendations in the app that are based on previous music or podcasts selections. Unlike exploration, which is based on uncertain user engagement and is used more so as a research tool to learn more about how people interact with suggested content. This balanced framework helps us ensure our shelves and cards are personalized for new and existing users. We also employ counterfactual training and reasoning to evaluate our algorithms without requiring A/B testing or randomized experiments.

This enables us to deliver your favorite content including recently played albums, heavily rotated tracks, and subscribed podcasts. At the same time, we recommend new tracks and artists based on listening history. These tracks are packaged up in playlists such as 'More like', 'Recommended for you', and 'Made for you'. Although the Home screen looks organic and continuous, almost everything that you see is determined by the exploitation and exploration approach.

Of course, this only works well when you can experiment non-stop with huge volumes of data. To give you a sense of scale, let's compare Spotify Home to Netflix's home screens again. Both home screens have a similar setup and challenges, right? Not exactly, the most significant difference is the sheer volume of user and content data—as well as the extent to which we must scale our machine learning operation to match demand.

The numbers speak for themselves. Netflix has about **158m users**, and at Spotify, we have **248 million**. On top of that Netflix has around 5,800 movies and shows **across all territories**, while Spotify has more than 50 million music tracks and 500,000 podcast titles.

Remember too, that a single Netflix title ranges between a half-hour for sitcom-like content, up to a few hours for movies. Whereas, the typical length of a Spotify track is maybe three and a half minutes. When you think about the frequency and amount of content data we need to index, it results in an enormous difference in scale.

The result? We need a highly scalable environment where we can run high volume, non-stop, and real-time experiments that enable us to make better use of our data, to deliver an outstanding experience to listeners.

Spotify v Netflix: Different magnitude of data analytics requirements...

	Movies	Music
End users	158 million	230 million
Content volume	5,800 movies + shows	50 million tracks + 500,000 podcasts
Content duration	30 minutes - 2 hours	3 ½ minutes
Content replay	Low	High

...resulting in a non-stop need to scale and test at speed

From brittle to unbreakable

From brittle to unbreakable

In Spotify's early days, we wrote a lot of custom data libraries and APIs to drive the machine learning algorithms behind our personalization efforts. This presented several challenges for the machine learning team. Supporting multiple systems wasn't ideal for our engineers to maintain while trying to scale our machine learning practices, and it doesn't align with our current efforts to improve engineer productivity.

For example, we needed to investigate different learning models underneath the multi-armed bandit framework including logistic regression, boosted trees, and more complex models such as the latest deep neural network architecture. We were constantly rewriting code to hold things together. When the experiments were done, we were left with a potentially brittle system that scaled poorly and required us to support multiple frameworks in a fast-paced production setting. It's hard to iterate and innovate in that environment.

This is why our move to a standardized machine learning infrastructure is so valuable and timely. With the help of TensorFlow Estimators (TFE) and Tensorflow Data Validation (TFDV) we're able to minimize a lot of the custom work previously required. With TFE we can train and evaluate models quicker than before, which creates a significantly faster iterative process.

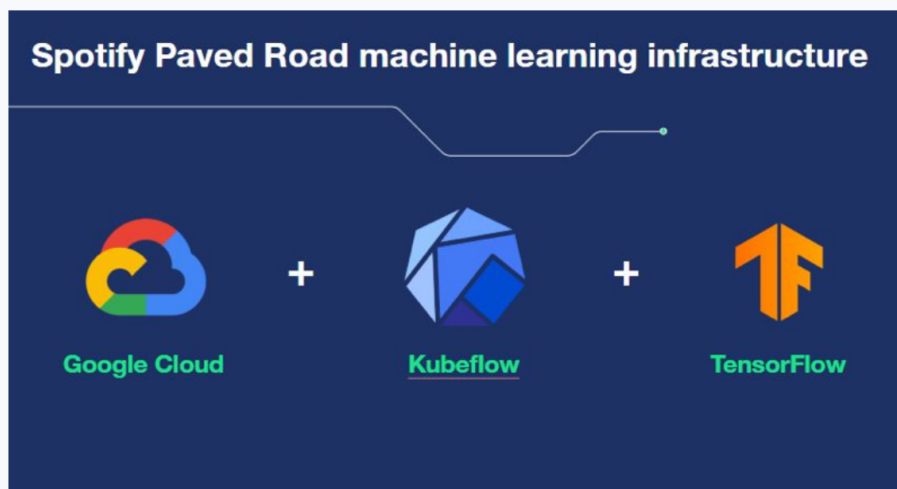
Additionally, migrating to Kubeflow is extremely valuable because it helps us better manage workloads, and accelerate the pace of experimentations and roll out. Now, it is significantly faster to automatically train and speed up our machine learning training algorithms.

With TFDV we can find bugs in data pipelines and machine pipelines while we are developing them, as well as during evaluation and roll out. Now, we can quickly see if there's any missing data or inconsistencies in our pipelines using a dashboard that plots the distribution of features and counts across different data sets.

In one case, we saw premium tier samples were missing from the training data and free shuffle tier data samples were missing from the evaluation pipeline. From an ML perspective, it's horrible to exclude these valuable data sets, but we were able to catch it quickly using TFDV. Even better, the dashboard can be configured to trigger alerts for certain thresholds so our engineers don't have to worry about data pipelines feeding into their systems.

Stepping up the pace with the Spotify Paved Road

Creating a personalized Home experience doesn't come without a cost. Previously, our engineers spent a considerable amount of time maintaining data and backend systems. Standardizing our ML infrastructure with the tools mentioned earlier (TensorFlow Extended, Kubeflow, and the Google Cloud Platform ecosystem) taught us several things about our engineering practices and productivity. A notable benefit of our newly established best practice is that it helped us achieve significant lifts in user satisfaction over popularity-based baselines in a very short amount of time.



That said, we are only just scratching the surface. We are determined to further harness the power of machine learning and AI to deliver the best possible personalized experience to our users.

Investing in engineers

Our investment in these technologies is also an investment in our machine learning engineers and their productivity. We want our engineers to focus on being innovative and pushing the bounds of what is possible with machine learning at Spotify, instead of participating in time-consuming infrastructure maintenance work. Our engineering culture revolves around being as productive and efficient as possible to help us continue scaling our platform while also creating a great user experience for creators and consumers.

Customizing Spotify Home is just the tip of the personalization iceberg at Spotify. While Machine Learning is an exciting and innovative space, we are continuously working to address challenges regarding creating personalized experiences. If you're interested in helping us address our ML engineering challenges, take a look at our [available positions](#) to join the band.

SHARE THIS:



Sign up for engineering updates

By clicking sign up you'll receive occasional emails from Spotify. You always have the choice to adjust your interest settings or unsubscribe.

[Spotify Jobs](#)
[Spotify.com](#)
[Newsroom](#)

[Spotify R&D Research](#)
[Spotify R&D Design](#)