




Music recommendation at Spotify

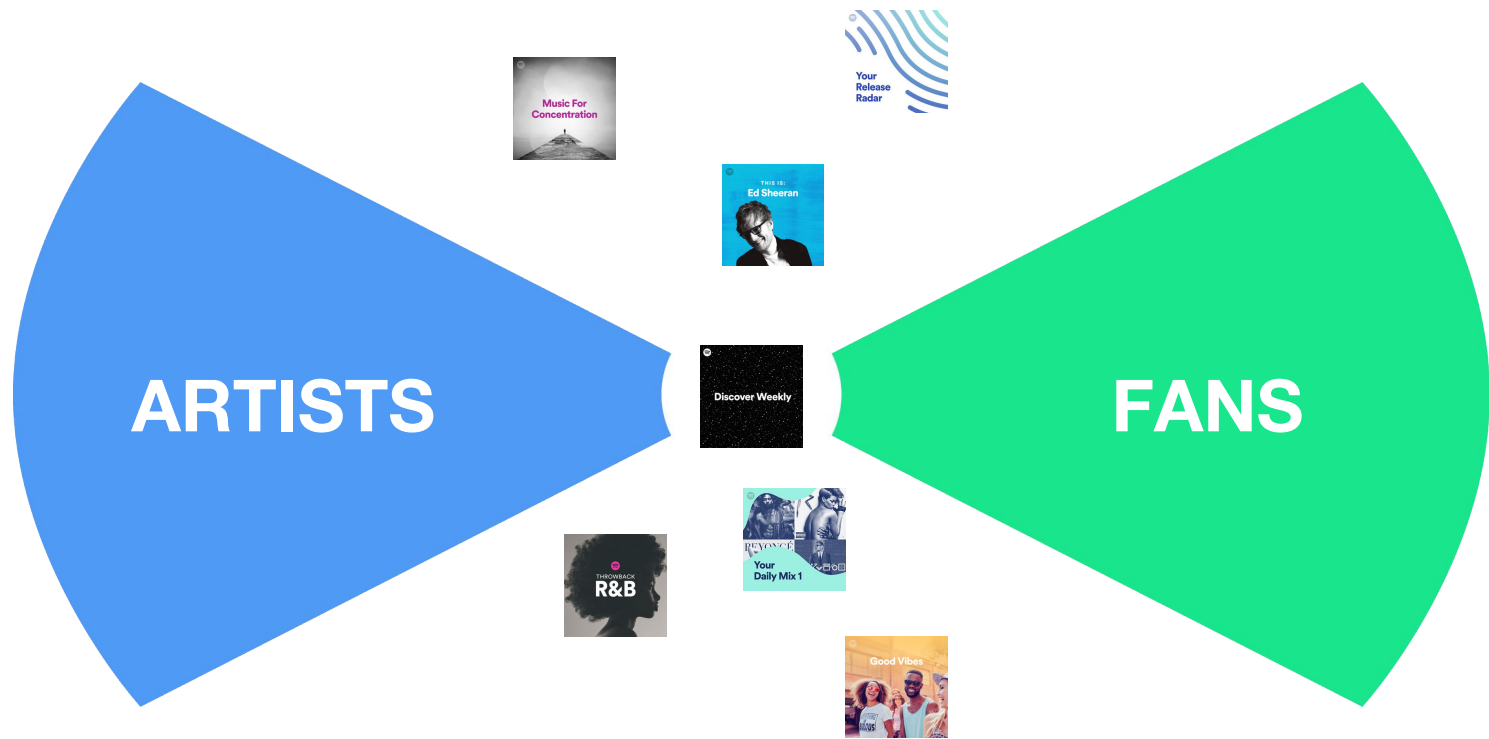
Ben Carterette

What we do

Spotify's mission is to unlock the potential of human creativity — by giving a million creative artists the opportunity to live off their art and billions of fans the opportunity to enjoy and be inspired by it.



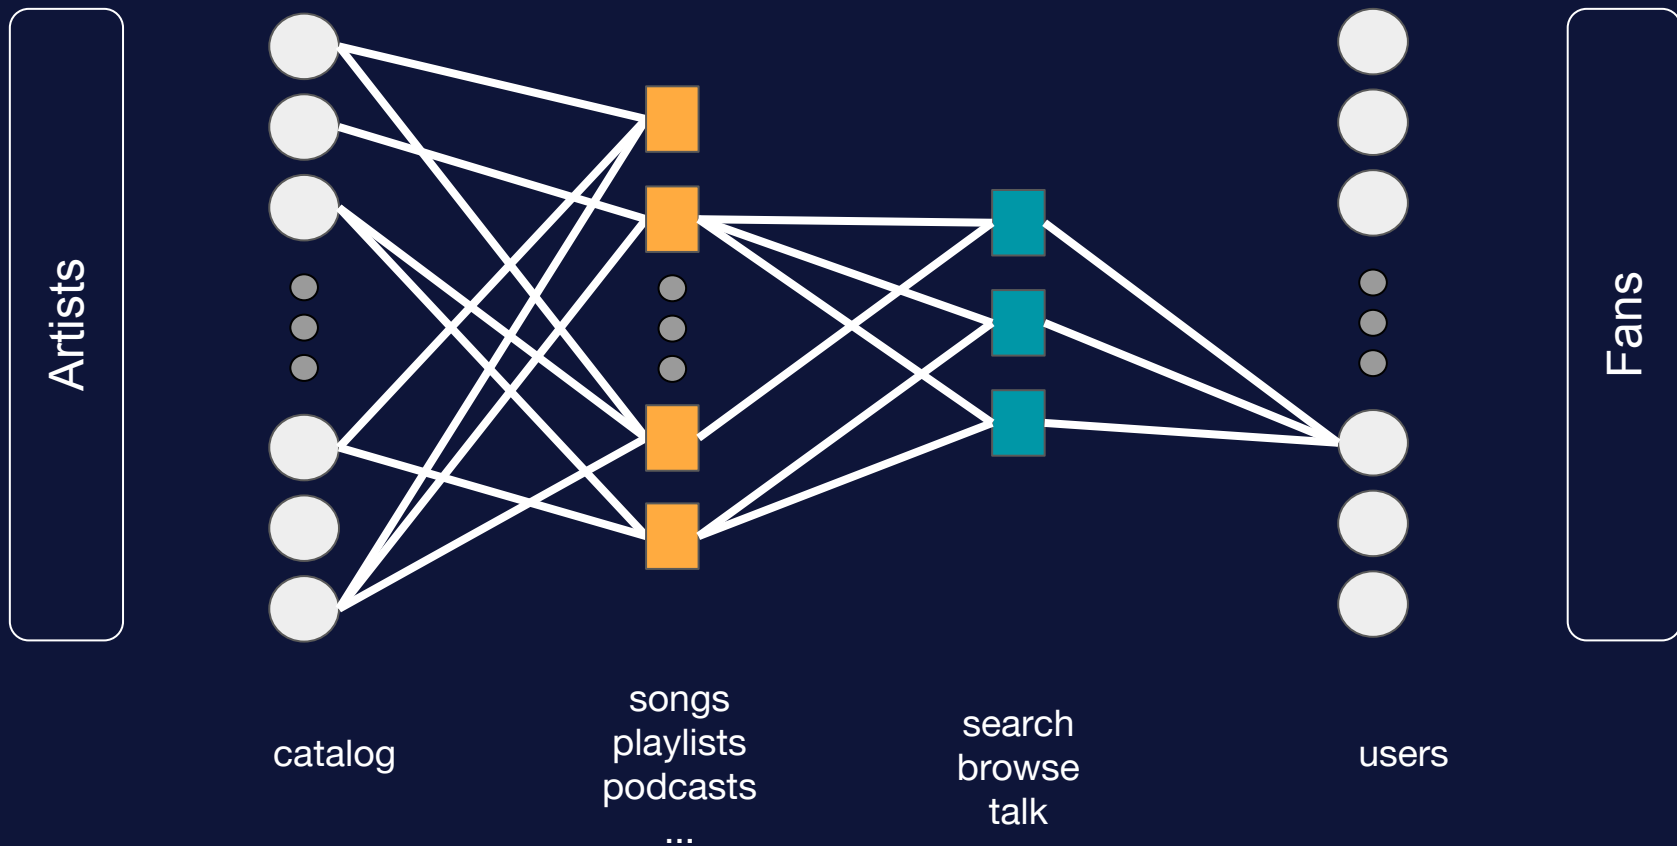
[illegible]



Our team mission:

Match fans and artists in a personal and relevant way.

What does it mean to match fans and artists in a personal and relevant way?



What does it
mean to match
fans and artists
in a personal
and relevant
way?



Personalization



Research @ Personalization

Areas of research expertise

Machine learning

Information retrieval

Evaluation

Language technologies

Content analysis

Algorithmic bias

Human computer
interaction

Recommender systems

User modeling

5 labs

... Boston, London, New York & San Francisco

hai: we research the interactions between the rich diversity of people and personalized audio experiences that matter to them.

LiLT: we research how Spotify users and creators communicate using written and spoken language, and how machine-learning models using this knowledge can improve user satisfaction.

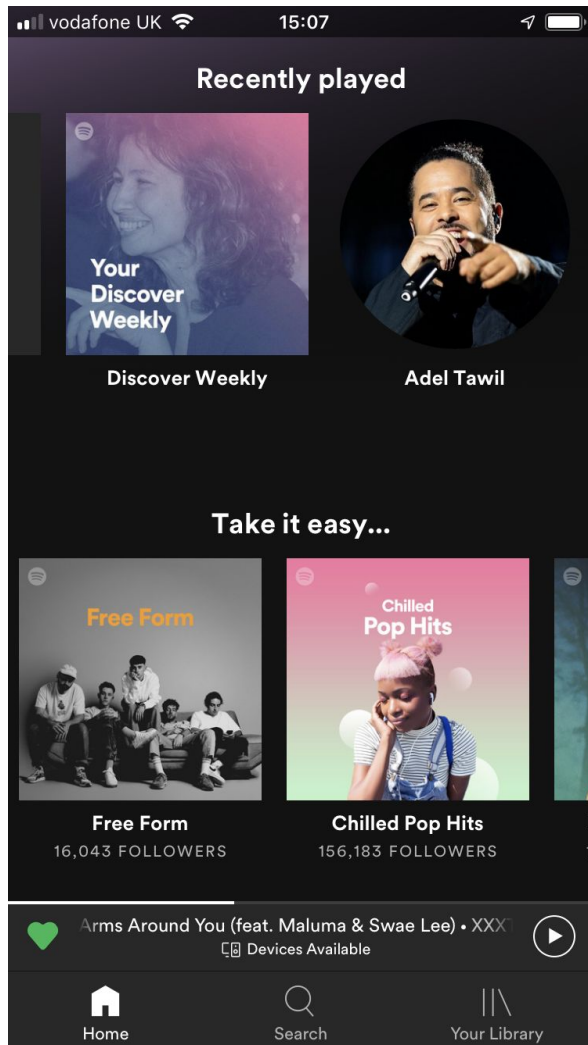
preamp: we research how to match audience to artists using machine learning, search & recommendation, and rigorous experimentation.

SIA: we develop machine learning based solutions to understand, interpret and influence interactions and consumption signals.

algo-bias: we empower Spotify teams to assess & address algorithmic bias and better serve underserved audiences & creators.

Examples

Home



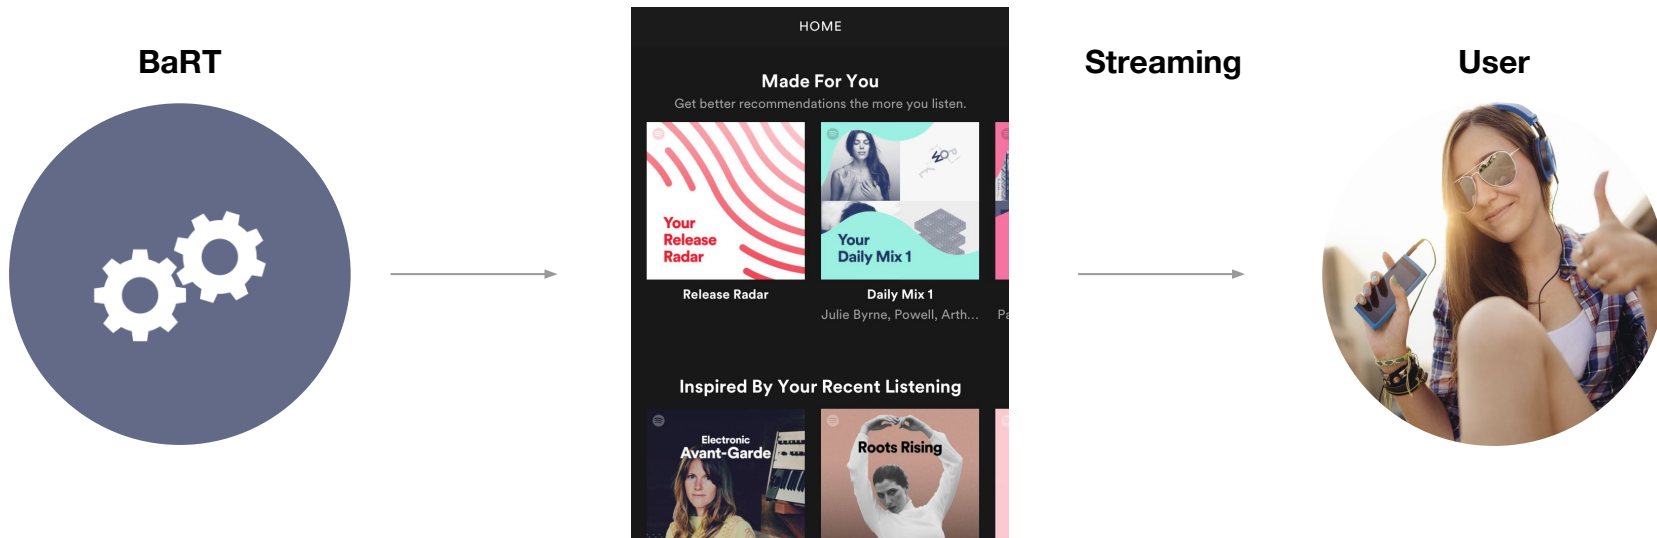
Home

Home is the **default screen** of the mobile app for all our users worldwide.

It surfaces the **best of what Spotify has to offer**, including music and podcasts for every situation, personalized playlists, new releases, old favorites, and undiscovered gems.



Value to the user here means helping them find something they're going to enjoy listening to, **quickly**.

BaRT: Machine learning algorithm for Spotify Home



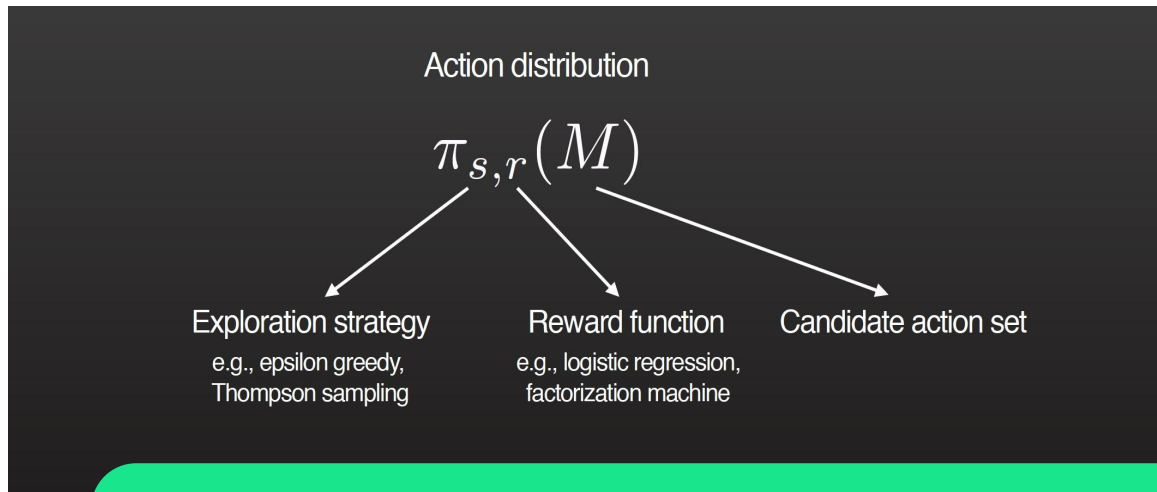
Explore, Exploit, Explain: Personalizing Explainable Recommendations with Bandits, J McInerney, B Lacker, S Hansen, K Higley, H.Bouchard, A Gruson, R Mehrotra, ACM RecSys 2018.

BaRT (Bandits for Recommendations as Treatments)



How to rank playlists (cards) in each shelf first, and then how to rank the shelves?

Multi-armed bandit algorithms




Explore vs Exploit


Flip a coin with given probability of tail

If head, pick best card in M according to predicted reward $r \rightarrow$ EXPLOIT

If tail, pick card from M at random \rightarrow EXPLORE



<https://hackernoon.com/reinforcement-learning-part-2-152fb510cc54>



Discover Weekly

Richer understanding of user satisfaction

Unambiguously
positive signals for
Discover Weekly



- Album view duration
- Artist view duration
- Downstream msPlayed
- Ds completed plays
- Album views count
- Artist views count
- Collection saves count
- Playlist adds count

Four main goals emerged; behaviors differ by goal

Play new background music

- No skipping
- ↑ Saves or adds
- ↑ Listening time
- ↑ Sessions per week

Listen to new music now and later

- ↑ Saves or adds
- ↑ % tracks heard
- ↑ Streams over half the song
- ↑ Downstream listening

Find new music for later

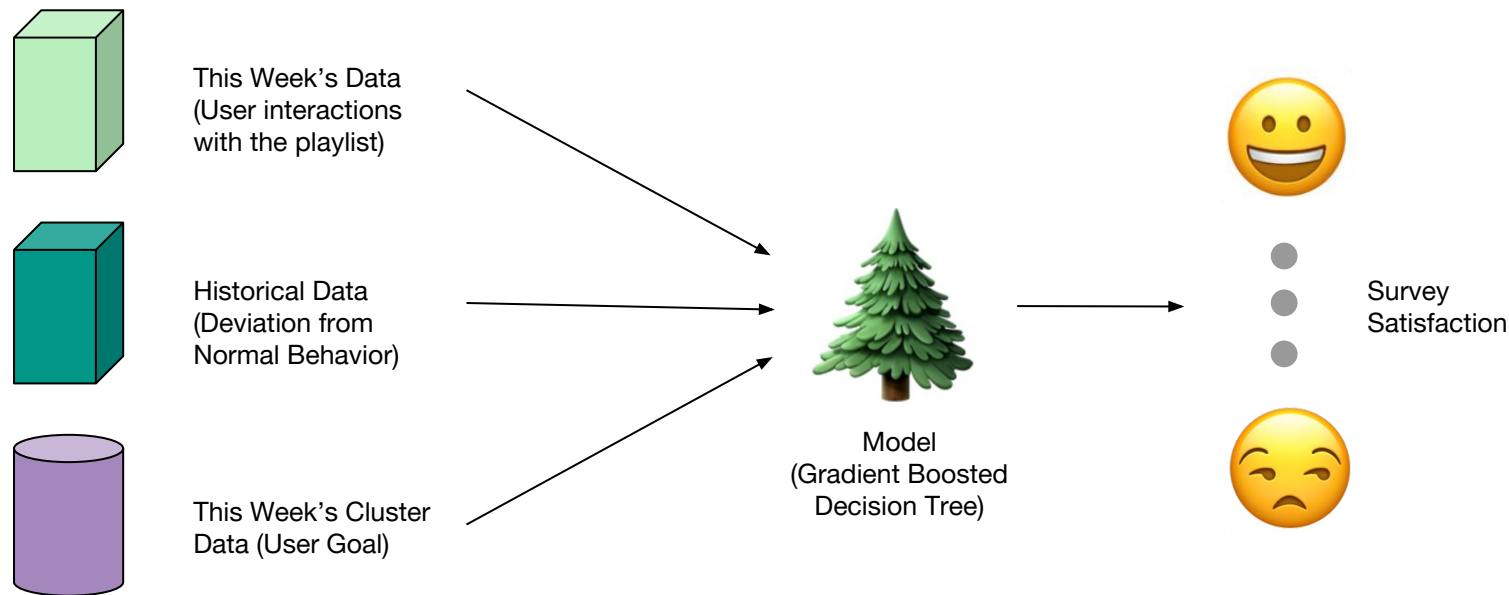
- ↑ Saves or adds
- ↑ Streams
- ↑ Downstream listening

Engage with new music


- ↑ Artist page views
- ↑ Album page views
- ↑ Downstream listening

Trained model to predict satisfaction for each track

Features were informed by hypotheses from user interviews



Current work: Modeled metric as an optimization target



Modeled metric
(user-track scores)

Learn to
Rank

Broken Heart

Spiritualized • Ladies & Gentlemen We Are Floati... ..

You Always Hurt The One You Love

The Mills Brothers • The Anthology: 1931 - 1968

Our Mutual Friend

The Divine Comedy • Absent Friends

Autumn Serenade

John Coltrane, Johnny Hartman • John Coltrane... ..

In the Kingdom

Mazzy Star • Seasons of Your Day

Artifact #1

Conor Oberst • Upside Down Mountain

While You Were Sleeping

Elvis Perkins • Ash Wednesday

**What we are working
on now ... some examples**

Multiple objective functions

$$a_t = \arg \min_{a_t} \pi$$

Metric 1

VS


$$a_t = \arg \min_{a_t} f(\pi_1, \pi_2, \pi_3 \dots)$$

Metric 2

Metric 3

Home

Optimising for fairness and satisfaction at the same time



Search

Large catalog

40M+ songs, 3B+ playlists

2K+ microgenres

Many languages

78 countries

Different modalities

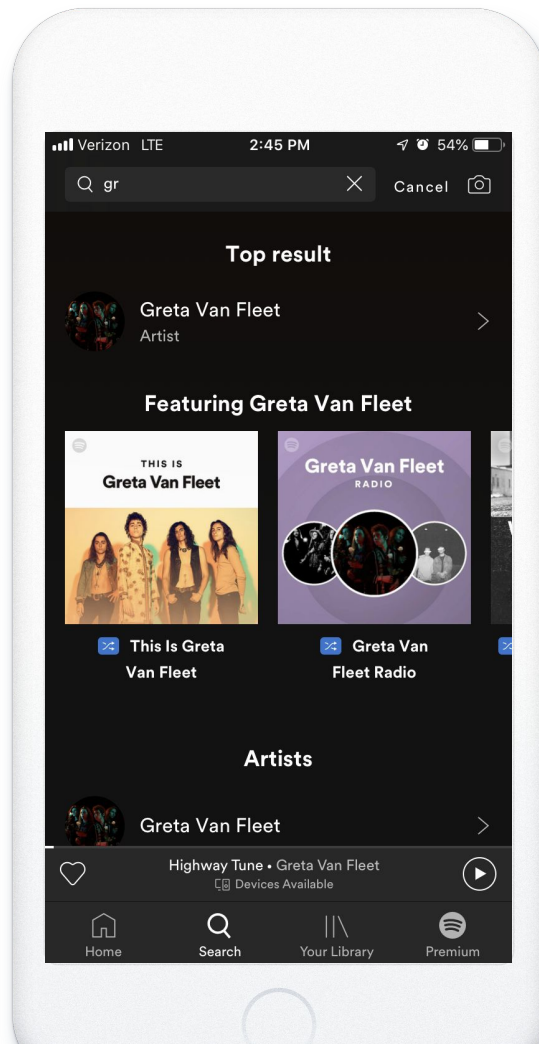
Typed, voice

Various granularities

Song, artist, playlist

Various goals

Focus, discover, lean-back, mood



Search

FOCUSED

One specific thing in mind

- Find it or not
- Quickest/easiest path to results is important

OPEN

A seed of an idea in mind

- From nothing good enough, good enough to better than good enough
- Willing to try things out
- But still want to fulfil their intent

How the user thinks about results

EXPLORATORY

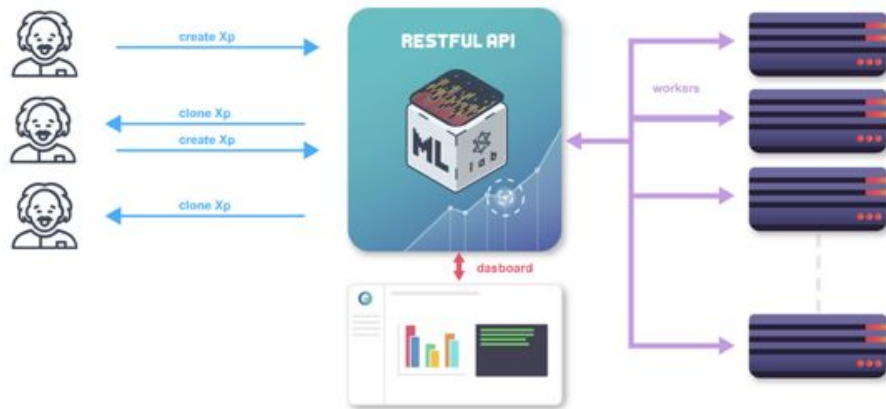
A path to explore

- Difficult for users to assess how it went
- May be able to answer in relative terms
- Users expect to be active when in an exploratory mindset
- Effort is expected

Evaluation

ML Lab

An offline evaluation framework to launch, evaluate and archive machine learning studies, ensuring reproducibility and allowing sharing across teams.



Other things we are doing

RecSys Challenge 2018



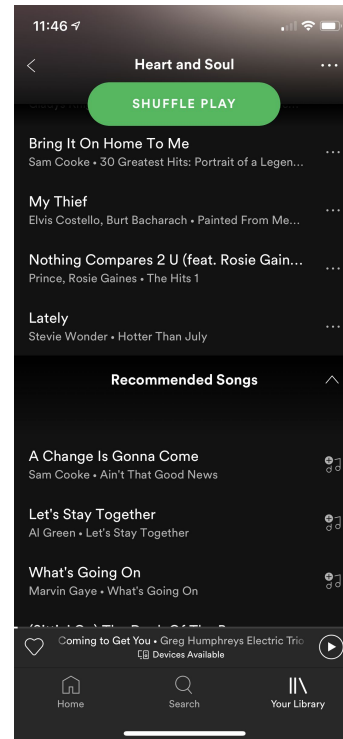
Earlier in 2018 we hosted the RecSys Challenge on Automatic Playlist Continuation, together with researchers from JKU Linz and UMass Amherst.

The dataset was 1 million user-created playlists from Spotify.

The challenge was to predict tracks that would complete a given playlist. This is similar to the Recommended Songs feature on Spotify.

Participation

791 participants from over 20 countries & 410 teams with 1497 submissions.



WSDM Cup 2019

We are currently running the WSDM Cup 2019 challenge on Sequential Skip Prediction.

The dataset is 130 million listening sessions on Spotify, along with associated interactions.

The challenge is to predict which tracks in a session will be skipped.



crowdAI

[Challenges](#) [Knowledge Base](#) [Job Board](#) [Sign up](#) [Log in](#)



Spotify Sequential Skip Prediction Challenge

Predict if users will skip or listen to the music they're streamed

By Spotify

24 days left

10634 319 129
Views Participants Submissions

♥ 53

FOLLOW

bit.ly/spotify-wsdm-cup-2019

Thank you