

Digital audio streaming is on the rise, with a reported 61% of U.S. listeners streaming audio in 2021 ([MRC data](#)), and audio song streaming increasing globally by 25.6% between 2021 and 2022 ([Luminate 2022](#)). Playlists play an increasingly important role in curating and organizing digital listening experiences, with a 2017 survey by [Nielsen Music](#) finding that 75% of listeners who stream music online use playlists and over half create their own. As playlist creation and curation is often a time consuming process, users frequently turn to track recommendation systems to help them in constructing new playlists ([Inman et al., 2020](#)). Most existing track recommendation systems rely on listeners to provide seed tracks, and then utilize a variety of different approaches to recommend additional tracks in either a playlist-like listening session or as sequential track recommendations based on user feedback.

Our recommendation engine takes a different approach, allowing listeners to generate playlists based on a semantic string, such as the title of desired playlist, specific mood (happy, relaxed), atmosphere (tropical vibe), or function (party music, focus). Using a publicly available dataset of [existing playlists](#), we combine a semantic similarity vector model with a matrix factorization model to allow users to quickly and easily generate playlists to fit any occasion.

Our engine progresses in two steps:

- 1.) The user prompt is run through a semantic similarity vectorized space, identifying the closest nearest neighbor playlists with similar semantic content, and returns the top 10 frequently occurring tracks within those playlists.
- 2.) Once approved, the 10 tracks are then run through a matrix-factorization model which was trained using gradient descent on a subset of the million playlists dataset. This generates a large list of recommended tracks that are then further filtered using cosine similarity to produce 40 more tracks to complete the playlist.

The recommendation engine simplifies the process of creating novel playlists from scratch, allowing users to bypass the more cognitively taxing process of identifying initial tracks and artists and allowing them to focus on curating and fine-tuning the desired playlist. By starting with a limited number of recommendations based on user input, users can more readily evaluate if suggested tracks are appropriate for the desired playlist before generating a longer list of recommended tracks.

With our recommendation engine, users can generate playlists independent of a specific streaming platform, nor do they need to provide a listening history or user profile to generate suggestions. This also means our engine can be readily integrated into new or existing streaming platforms as means of enhancing user experience and assisting in playlist generation.

Finally, there is a high potential for further refining recommendations, both through the incorporation of additional layers, such as [music sentiment](#) to improve recommendations for specific moods or emotions, and through user feedback and interaction with the engine that can inform improvements to both the model and the user experience.