

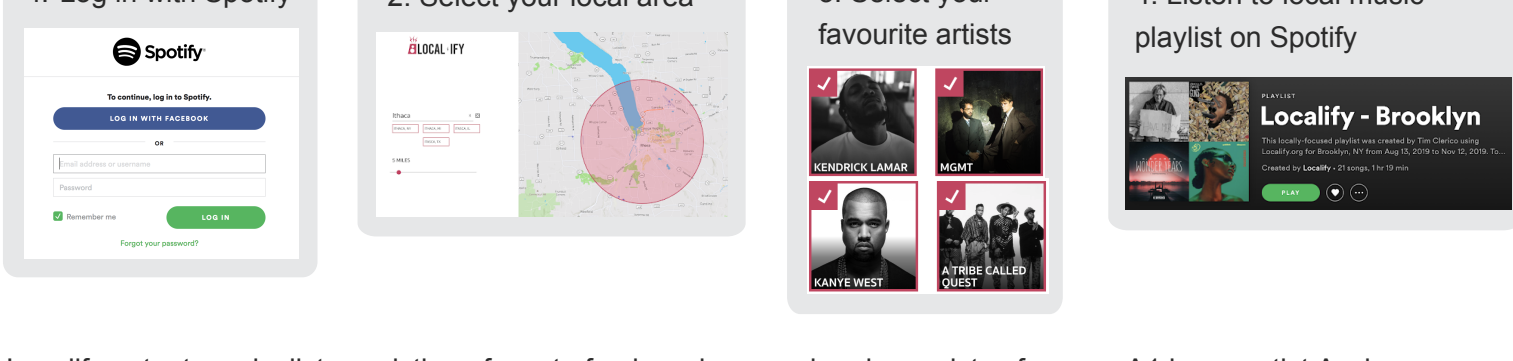
Exploring Acoustic Similarity for Novel Music Recommendation

Derek Cheng, Thorsten Joachims, Doug Turnbull

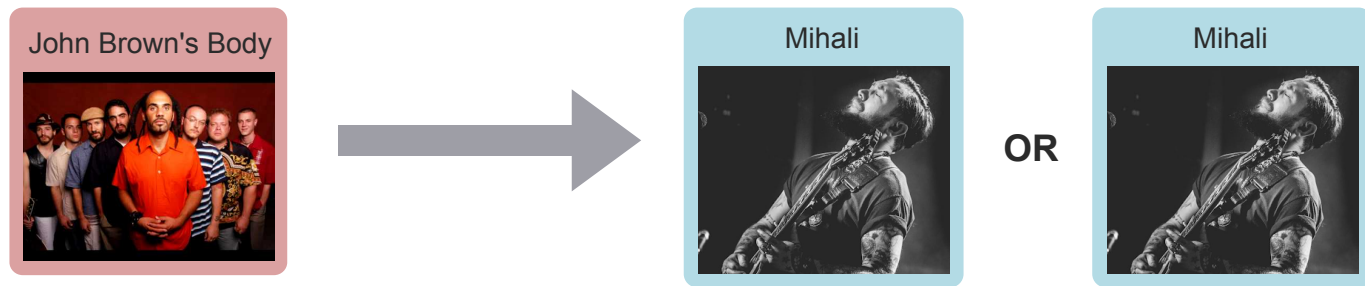
CORNELL UNIVERSITY · ITHACA COLLEGE

Introduction: Localify

Localify is a music recommender that uses your Spotify data to help you discover local artists that match your tastes.



Localify outputs a playlist consisting of a set of pairs, where each pair consists of a song A1 by an artist A who you already like (a seed song), followed by some song by a novel local artist B.



However, given song A1, which of artist B's songs would be most suitable to use in the playlist after A1?

Research Questions

Intuitively, it seemed that the most acoustically similar song would be the most suitable, because it would "flow" most naturally in the playlist. We formulated the following research questions to investigate this idea:

RQ1: Are human judgments about acoustic similarity consistent across users?

RQ2: Is acoustic similarity a good proxy for how an individual might

- construct a **playlist**,
- recommend** music to a friend, or
- prefer** one song over another?

RQ3: If so, what are some of the measurable acoustic properties that correlate with how humans judge acoustic similarity?

Methodology

Song Tuples

We asked participants to listen to and answer questions about song tuples we collected.

The song tuples were formed as shown on the right, closely emulating the original problem encountered for Localify.



3
Songs per tuple

4
Music genres represented

12
Song tuples collected

36
Songs used

Song Tuple Questions

- If you were creating a playlist with Song A1 and either B1 or B2, which one would you pick?
- If you had a friend who likes Song A1 by Artist A and you wanted to introduce them to Artist B, which would you recommend to them first?
- Which song is most acoustically similar to Song A1?
- Which song do you prefer?

Open Ended Questions

- When deciding to pick specific songs for a playlist, what do you consider to be most important?
- When deciding songs to recommend to a friend, what do you consider to be most important?
- When comparing songs in terms of acoustic similarity, what do you consider to be most important?

113 participants

58% male, 42% female



Most popular streaming service

Results



21-25
Median age group

≤ 17
Youngest group

61-70
Oldest age group

Qualitative Feedback

Most common factors considered important for:

Acoustic similarity

- Instrumentation
- Tempo
- Beat / rhythm
- Vibe / tone
- Vocal tone

Playlist selection

- Vibe / tone*
- Mood / expression*
- Acoustic similarity*
- Genre / style
- Tempo*

Recommendation

- Friend will like
- Acoustic similarity*
- Personal preference
- Genre / style
- Vibe / tone*

For the tasks of selecting a song for a playlist and recommending a song to a friend, many participants considered factors related to acoustic similarity (italicized) to be important to their decision!

Quantitative Feedback

The following table summarizes the results obtained for each song tuple. The first line for acoustic similarity, playlist selection, recommendation, and personal preference displays the number of participants respectively that selected the first place B1/B2 song / indicated that they were the same / selected second place song based on acoustic similarity. The second line represents the p-value for a binomial hypothesis test in which the null hypothesis assumes that B1/B2 songs are equally likely to be selected by a participant. Bold font indicates statistically significant differences at the $\alpha < 0.05$ level. Italics indicate that participants generally preferred the less acoustically similar song.

Genre	Artists	Acoustic Similarity	Playlist Selection	Recommendation	Personal Preference
Rock	The Beatles Aviator Stash	45 / 12 / 5 0.000	45 / 15 / 5 0.000	45 / 15 / 5 0.000	31 / 18 / 13 0.006
Hip hop	Nicki Minaj Mulatto	19 / 0 / 1 0.000	14 / 3 / 3 0.010	14 / 5 / 1 0.001	5 / 6 / 9 <i>0.244</i>
Rock	Paramore Tonight Alive	14 / 2 / 1 0.001	13 / 3 / 1 0.002	12 / 1 / 4 0.056	8 / 7 / 2 0.108
Pop	Post Malone Lil Xan	22 / 8 / 5 0.001	17 / 8 / 10 0.126	15 / 11 / 9 0.156	10 / 10 / 15 0.195
Pop	Billie Eilish Gabbie Hanna	29 / 5 / 10 0.002	29 / 6 / 9 0.001	28 / 5 / 11 0.006	20 / 13 / 11 0.079
Hip hop	Cardi B Kash Doll	13 / 2 / 3 0.017	12 / 2 / 4 0.056	8 / 4 / 6 0.367	6 / 5 / 7 <i>0.419</i>
Rock	Imagine Dragons The Score	13 / 5 / 3 0.017	8 / 10 / 3 0.161	10 / 8 / 3 0.070	3 / 9 / 9 <i>0.107</i>
Hip hop	Drake Kahiem Rivera	14 / 2 / 4 0.023	13 / 3 / 4 0.036	11 / 6 / 3 0.044	7 / 9 / 4 0.322
Pop	The Weeknd Myer Clarity	21 / 5 / 10 0.041	13 / 8 / 15 0.279	18 / 6 / 12 0.161	13 / 6 / 17 <i>0.223</i>
R&B	Beyoncé Keri Hilson	9 / 2 / 3 0.107	12 / 0 / 2 0.011	11 / 0 / 3 0.044	7 / 4 / 3 0.234
R&B	Camila Cabello Ally Brooke	9 / 2 / 4 0.175	9 / 2 / 4 0.175	10 / 0 / 5 0.183	10 / 1 / 4 0.122
R&B	Frank Ocean Syd	9 / 1 / 5 0.244	10 / 2 / 3 0.070	9 / 3 / 3 0.107	10 / 2 / 3 0.070

In 9 of the 12 tuples, the similarity judgment was significantly ($\alpha < 0.05$) pointing in one direction, suggesting that there was a winner between the B1/B2 songs. Overall, the results suggest that listeners are somewhat consistent in their judgment of acoustic similarity even when comparing songs by similar artists.

This table shows the correlation coefficients for our song tuple results when comparing pairs of acoustic similarity, playlist selection, song recommendation, and personal preference. Here, acoustic similarity is significantly correlated with playlist selection and recommendation!

	Playlist Selection	Recommendation	Personal Preference
Acoustic Similarity	0.716	0.595	0.387
Playlist Selection		0.596	0.386
Recommendation			0.116

Conclusion

RQ1: Are human judgments about acoustic similarity consistent across users?

Yes, there is a degree of consistency among human judgments of acoustic similarity.

RQ2: Is acoustic similarity a good proxy for how an individual might construct a playlist, recommend music to a friend, or prefer one song over another?

Yes, acoustic similarity seems to be highly correlated with playlist selection and recommendation, *but not* personal preference.

RQ3: If so, what are some of the measurable acoustic properties that correlate with how humans judge acoustic similarity?

While certain acoustic features seem to be related to acoustic similarity, further investigation is necessary to determine the extent to which these features are indeed related to acoustic similarity.

Contact us

Derek Cheng
dsc252@cornell.edu

Thorsten Joachims
tj@cs.cornell.edu

Doug Turnbull
dturnbull@ithaca.edu

This project is supported by NSF grant IIS-1615706/1615679 and IIS-1901168/1901330.