Factoring in What Gets Listened To

Performance of Factorisation Machine using Musical Features for Children

Konrad Barbers | k.f.barbers@student.tudelft.nl | Faculty EEMCS, TU Delft

Supervisors: Robin Unruh, Sole Pera

Background

- Many recommender systems are well researched for "normal" users[1]
- Children and other non-normative listeners often have traditional recommenders perform worse[2]
- Children's listening behaviour has been studied, but no recommender has been evaluated[3]

Research Question

How well does a music recommender system using matrix factorisation leveraging various audio features perform for child users?

Resources

Dataset of users' listening behaviour: LFM-2b

- filtered for children, at least 10 listens

Songs' feature values, extracted from Spotify:

- Tempo
- Average Loudness
- Musical mode
- Time signature

Experiment

- Use 0.1% of user's listening events for training, 99.9% for heuristics (nDCG, MRR) (hardware limitations)
- Factorisation Machine without features as control, evaluate extension with all features and all combinations of 3 features

Results

Features used	nDCG	MRR
Control	0.13077	0.31161
All features	0.15900	0.32225
No loudness	0.15974	0.37828
No mode	0.15369	0.35678
No tempo	0.18303	0.49598
No time signature	0.17544	0.41342

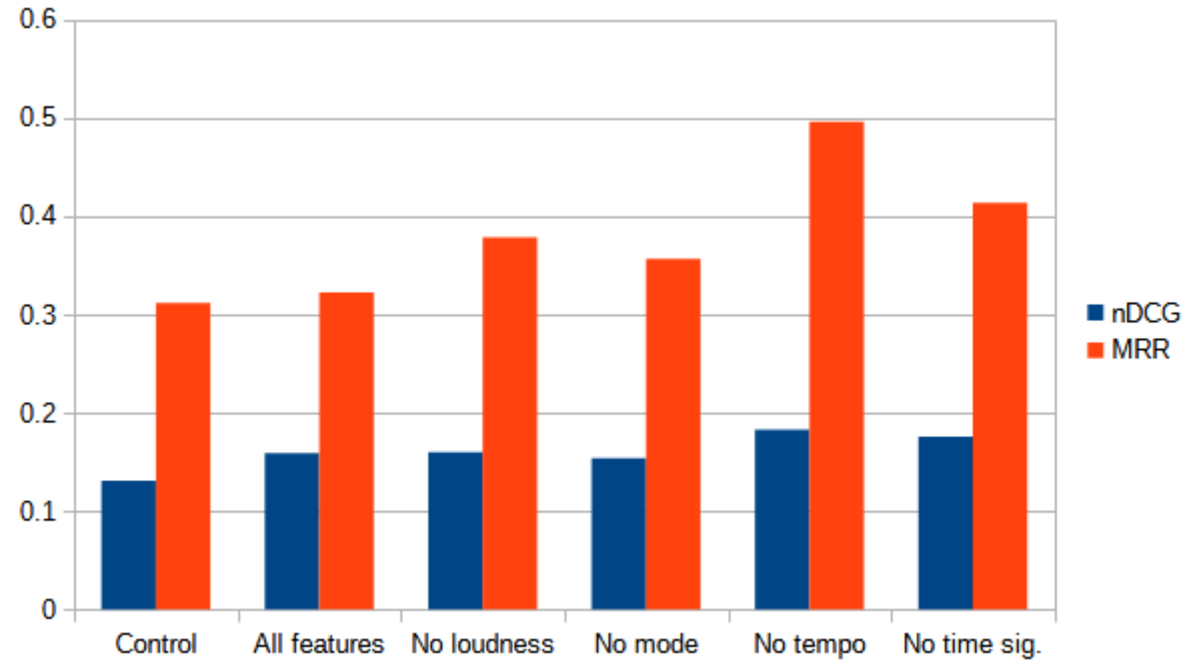

- All recommenders with added features performed better than the control, but fewer features tend to cause better scores
- Removing mode has little to no positive effect
- Using no tempo values significantly improves the scores

Discussion

- Quality over quantity for selecting features, some features may oppose each other
- Very small training set due to demanding performance of FM
- Because of the low split, comparison of the results to similar research is not advisable

Conclusion

- Do features improve the recommender? Yes, but...
- For better results, this experiment should be repeated with a larger training set (80%)
- What other features can be used to improve the recommender?
- How do features extracted through other means (e.g. signal processing) perform?
- How good is this recommender for recommending music to adults?

1. Schedl et al.. Music Recommender Systems, pages 453–492. Springer US, Boston, MA, 2015.
2. Kowald et al.. Support the underground: characteristics of beyond-mainstream music listeners. EPJ Data Science, 10, 2021.
3. Spear et al.. Baby shark to barracuda: Analyzing children's music listening behavior. In Proceedingsof the 15th ACM Conference on Recommender Systems, RecSys '21, page 639–644, New York, NY, USA, 2021. Association for Computing Machinery.