HUMAN VS. AUTOMATED JUDGEMENTS OF CROSS-CULTURAL MUSICAL SIMILARITY

Hideo Daikoku, Marino Kinoshita, Shinya Fujii, Patrick E. Savage

Keio University, Shonan Fujisawa Campus
{t16559hd, t14513mk, sfujii, psavage}@sfc.keio.ac.jp

ABSTRACT

Although MIR has demonstrated great success in automatic analysis of Western music, no study has tested automatic algorithms against perceptual ground-truth data for a global musical sample. It thus remains unknown whether MIR algorithms can be meaningfully applied to automatically compare diverse music from around the world. In this pilot study, we aim to establish ground truth perceptual data on similarity between diverse musical recordings from across the world and use this perceptual data to test the accuracy of existing audio similarity algorithms. Preliminary results (two participants, ten recordings) suggest that perceptual ratings of musical similarity are significantly correlated between participants, but these similarities are only weakly correlated with similarities measured by an automatic algorithm. While this is consistent with more pessimistic assessments of MIR's current ability to accommodate non-Western music, we hypothesize that collecting more perceptual data, comparing against more algorithms, and creating new algorithms based on more universal musical theories will enable meaningful automatic analysis of all the world's music within the near future. This would have important implications for our understanding of cross-cultural music diversity, including applications to domains such as music recommendation and cultural heritage preservation.

1. INTRODUCTION

The growth of Music Information Retrieval (MIR) has made possible rapid, objective, and relatively accurate automatic analysis of Western classical and popular music, which constitutes the bulk of the music industry [1]. MIR has achieved this success in part through developing algorithms that incorporate aspects of Western music theory (e.g., 12-tone equal tempered scales, 4/4 meters). In contrast, there have been very few automatic studies of non-Western musics, for which such theories do not necessarily apply [2, 4, 5, 7]. Crucially, no MIR studies have

© Hideo Daikoku, Marino Kinoshita, Shinya Fujii, Patrick E. Savage. Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). **Attribution:** Hideo Daikoku, Marino Kinoshita, Shinya Fujii, Patrick E. Savage "Human vs. Automated Judgements of Cross-Cultural Musical Similarity", 19th International Society for Music Information Retrieval Conference, Paris, France, 2018.

successfully validated automatic analysis of a global musical sample against real perceptual data. Thus, it remains

unknown whether current MIR technology can meaningfully be used to automatically compare diverse music from throughout the world.

Therefore, our principal task is to collect perceptual similarity measurements to establish a reliable ground-truth data set for a diverse global musical sample and use it to evaluate both existing algorithms and new ones.

2. METHODS

2.1. Perceptual similarity

In our pilot study we used a subset of ten music recordings out of a set of 20 diverse recording from around the world analyzed in [6] (Fig. 1), which includes diverse genres such as folk, popular, and art music (e.g., BaAka hunting song, Brazilian samba, Mozart piano concerto).

Figure 1. Map of the musical recording sample.

To enable experiments comparing all possible pairs, each recording was shortened to a 10s excerpt and RMS-normalized to a consistent volume. In the pilot experiments, participants (N=2) listened to all 45 possible pairs among the 10 recordings and rated each pair on their musical similarity on a scale from 0 ("very different") to 100 ("very similar").

2.2. Automatic similarity

We measured automatic musical similarity using the untrained algorithm published in [7]. For each recording, this algorithm extracts features relevant to timbre (MFCCs), melody (pitch bi-histograms), harmony (chromatograms) and rhythm (onset patterns). A window size

of 40ms and a hop size of 5ms was used for all recordings. Normalized Mahalanobis distances were then measured in each pair in order to create a multivariate similarity matrix.

3. RESULTS

We found a strong and significant correlation between the perceptual ratings from the two participants ($r = .61$, $p < .001$) as shown in Fig. 2a. On the other hand, there was only a weak correlation between the perceptual similarity score and the algorithmic similarity value ($r = 0.08$, $p < .05$) as shown in Fig. 2b.

Figure 2. Correlations between musical similarity matrices for **a)** perceptual ratings from two participants ($r = .61$, $p < .001$), and **b)** the average perceptual ratings vs. an automatic algorithm ($r = .08$, $p < .05$).

4. DISCUSSION

Our preliminary results showed that it was possible to measure perceptual similarity in a diverse sample of the world’s music that is significantly correlated across participants ($r = .61$), despite substantial variation in perception between participants. In contrast, our comparison of perceptual data against automatic similarity showed only a weak correlation ($r = .08$), suggesting that the automated algorithm is not capturing musical similarity in a very perceptually meaningful way. This is consistent with criticisms that MIR algorithms may not yet be useful for cross-cultural automatic comparison due to a historical emphasis only on a narrow range of Western popular and classical music [4, 7]. However, this conclusion remains tentative until we have collected more perceptual data and compared it against a broader range of leading similarity algorithms, including the full trained version of the algorithm in [2].

5. CONCLUSION AND FUTURE WORK

We should have more conclusive results after we have conducted our experiments on a larger scale. We plan to eventually test many participants from around the world using the full sample of 20 recordings (we estimate that the full experiment comparing all 190 pairs among the 20 recordings will take approximately 3 hours per participant). We also plan to test our results not only on the trained and untrained versions of the algorithm from [2] but other algorithms that have been most successful in the annual MIREX audio similarity evaluation. We are also conducting feature-wise comparisons across songs.

We hypothesize that our algorithmic results will be more consistent with perceptual data once data from

many participants is averaged and the algorithm is trained on a cross-cultural musical sample as in [2]. However, we still suspect that correlations with perceptual ratings may remain weak until we develop new algorithms and features based on more universal musical theories rather than ones designed primarily for Western music [3, 4].

If we can succeed in developing and validating methods for automated analysis that work with all the world’s music, we hope to expand the scale of analysis so they can be used to investigate broader questions such as whether automated analysis can predict cross-cultural aesthetic preferences. We hope to apply such findings to basic research in fields such as information science, music cognition, anthropology, and to broader social and economic applications including music recommendation services, music copyright law, music education, cultural heritage preservation, and cross-cultural understanding through music.

6. ACKNOWLEDGMENTS

This work was funded by a Keio Research Institute at SFC Startup Grant and a Keio Gijuku Academic Development Fund Individual Grant to P.E.S. We thank Rei Konno and Shota Shimazu for assisting with experimental design.

7. REFERENCES

- [1] M. Casey, R. Veltkamp, M. Goto, M. Leman, C. Rhodes, and M. Slaney, “Content-Based Music Information Retrieval: Current Directions and Future Challenges,” *Proceedings of the IEEE*, vol. 96, no. 4, pp. 668–696, 2008.
- [2] M. Panteli, E. Benetos, and S. Dixon, “A Computational Study on Outliers in World Music,” *PLoS One*, vol. 12, no. 12, p. e0189399, 2017.
- [3] P. E. Savage, “Alan Lomax’s Cantometrics Project: A Comprehensive Review,” *Music Sci.*, vol. 1, pp. 1–19, 2018.
- [4] P. E. Savage, “The Need for Global Studies,” in *Oxford Handbook of Music and Corpus Studies*, D. Shanahan, A. Burgoyne, and I. Quinn, Eds. New York: Oxford University Press. Preprint doi: 10.31235/osf.io/nxtbg
- [5] X. Serra, “The Computational Study of a Musical Culture Through its Digital Traces,” *Acta Musicol.*, vol. 89, no. 1, pp. 24–44, 2017.
- [6] M. Tenzer, “A Cross-Cultural Topology of Musical Time: Afterword to the Present Book and to Analytical Studies in World Music (2006),” in *Analytical and cross-cultural studies in world music*, M. Tenzer and J. Roeder, Eds. New York: Oxford University Press, 2011, pp. 415–439.
- [7] G. Tzanetakis, A. Kapur, W. A. Schloss, and M. Wright, “Computational Ethnomusicology,” *Journal Interdisciplinary Music Studies*, Vol. 1, no. 2, pp. 1–24, 2007.