preferences. Specifically, the predicted interest *pi* of the target user u in music piece m is the cosine similarity between u's music preference $\mathbf{p}_u$ and m's distributed representation $\mathbf{v}_m$, which is defined as follows:

$$pi(m \mid u, \mathbf{p}_u) = \cos(\mathbf{p}_u, \mathbf{v}_m) \tag{4}$$

Therefore, the ranking of music pieces $>_{u,\mathbf{p}_u}$ in our approach is defined as

$$m_i >_{u,\mathbf{p}_u} m_i' :\Leftrightarrow pi(m_i \mid u, \mathbf{p}^u) > pi(m_i' \mid u, \mathbf{p}^u) \tag{5}$$

We then can recommend the music pieces with high ranking scores (similar to user's musical preference) to the target user.

3. EXPERIMENTS

The experiments consist of two parts: evaluation of *music2vec* and the comparison of the proposed approach with baselines.

Firstly, we illustrate the effect of *music2vec* model by visualizing similarity among music pieces given in Table 1. As shown in Figure 1, music pieces with similar styles, such as singers, tags, and genres, have similar distributed representations. For example, "Summer" and "Moonlit Sea of Clouds", which have the same genre and player, do lie nearby in the real-valued distributed representation space. Besides, neither of these two music pieces has similar distributed representations with the other music pieces in Table 1. Therefore, the learned distributed representations with *music2vec* capture useful features effectively and depict music pieces well.

Table 1. Basic information of music examples

No	Name-Singer	Tags
1	Hero-Mariah Carey	pop, female vocalists, 90s, ballad
2	Without You-Mariah Carey	pop, female vocalists, soul, love
3	Drowning-Backstreet Boys	pop, boy bands, ballad
4	My Love-Westlife	pop, boy bands, Irish
5	Don't Cry-Guns N' Roses	classic rock, hard rock, ballad
6	Hotel California-Eagles	classic rock, rock, 70s
7	Fall Again-Kenny G	smooth jazz, R&B, Soul
8	Heart and Soul-Kenny G	smooth jazz, Rhythm and blues
9	Summer-Joe Hisaishi	sound track, Japanese, anime, instrumental, classical
10	Moonlit Sea of Clouds-Joe Hisaishi	sound track, Japanese, anime, instrumental, classical

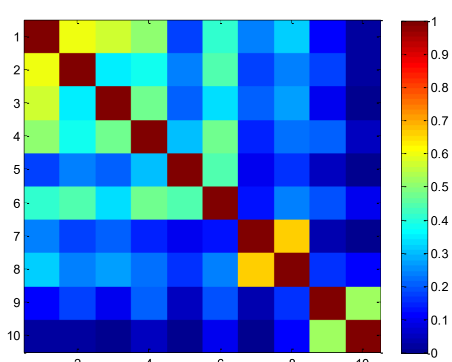

Figure 1. Similarity visualization of music examples with distributed representations

Then, we compare the proposed method with three state-of-the-art recommendation algorithms, including Bayesian Personalized Ranking (BPR) [2], FISMauc (FISM) [3], and user based collaborative filtering method (UserKNN) [4] on a real world dataset collected from Xiami Music (http://www.xiami.com/). From the comparison results in Table 2, we can see that our approach outperforms baselines in terms of F1 score and hitrate. Taking the F1 score as an example, when compared with BPR, FISM, and UserKNN with the recommending number being 10, the relative performance improvement achieved by the proposed approach is around 33.3%, 20.7%, and 42.2%, respectively. The improvements indicate that our approach is more effective than baselines in acquiring users' preference and assisting music recommendation.

In conclusion, the proposed approach can effectively recommend music pieces appropriate for target users and satisfy their preferences well.

Table 2. Comparisons with baselines

Methods	F1, %		Hitrate, %	
	@10	@20	@10	@20
Our approach	**7.88**	**9.23**	**36.65**	**43.14**
BPR	5.91	7.13	26.96	30.21
FISM	6.53	7.50	29.68	33.52
UserKNN	5.54	6.34	22.72	25.61

4. CONCLUSIONS AND FUTURE WORK

We present a music recommendation approach, which learns the distributed representations of music pieces from users' historical listening records, and utilizes these distributed representations to acquire users' music preferences and recommend appropriate music pieces. Experimental results show the effectiveness of the proposed approach. There are two possible future directions. Firstly, we plan to combine the distributed representation with more advanced recommendation techniques [5, 6], to further improve the performance. Secondly, we will try to evaluate our approach by online experiments.

5. REFERENCES

[1] T. Mikolov, I. Sutskever, K. Chen, G.S. Corrado, and J. Dean, Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*, 3111-3119, 2013.

[2] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme, BPR: Bayesian personalized ranking from implicit feedback. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*, 452-461, 2009.

[3] S. Kabbur, X. Ning, and G. Karypis, Fism: factored item similarity models for top-n recommender systems. In *Proceedings of the 19th ACM international conference on Knowledge discovery and data mining*, 659-667, 2013.

[4] P. Resnick, N. Iacovou, M. Suchak, P. Bergstrom, and J. Riedl, GroupLens: an open architecture for collaborative filtering of netnews. In *Proceedings of the ACM conference on Computer supported cooperative work*, 175-186, 1994.

[5] S. Deng, L. Huang, and G. Xu. Social network-based service recommendation with trust enhancement. *Expert Systems with Applications*, 41, 18 (12/15/), 8075-8084, 2014.

[6] A. Karatzoglou, X. Amatriain, L. Baltrunas, and N. Oliver, Multiverse recommendation: n-dimensional tensor factorization for context-aware collaborative filtering. In *Proceedings of the fourth ACM conference on Recommender systems*, 79-86, 2010.