

Music Sequence Prediction with Mixture Hidden Markov Models

Tao Li, Minsoo Choi, Kaiming Fu, Lei Lin

Abstract—Recommendation systems that automatically generate personalized music playlists for users have attracted tremendous attention in recent years. Nowadays, most music recommendation systems rely on item-based or user-based collaborative filtering or content-based approaches. In this paper, we propose a novel mixture hidden Markov model (HMM) for music play sequence prediction. We compare the mixture model with state-of-the-art methods and evaluate the predictions quantitatively and qualitatively on a large-scale real-world dataset in a Kaggle competition. Results show that our model significantly outperforms traditional methods as well as other competitors. We conclude by envisioning a next-generation music recommendation system that integrates our model with recent advances in deep learning, computer vision, and speech techniques, and has promising potential in both academia and industry.

Index Terms—Recommendation System, Music Information Retrieval, Sequence Prediction, Mixture Hidden Markov Model.

1 INTRODUCTION

Traditional ways of listening to music, such as playing CDs or tapes, are being replaced by online music stores and streaming services [1], [2]. A report from *the International Federation of the Phonographic Industry* [3] shows that the revenue of streaming industry increased by 41.1% last year, and now the streaming industry is the most significant portion of the music industry [4], [5]. As opposed to traditional ways, the streaming services enable users to interact with music platforms dynamically. For example, one can select and delete a song from his/her playlist any time and replace it with an entirely different one. However, it is inconvenient for users to manually update the playlists every time they listen to a new song. To improve users' listening experience, various intelligent personalized music recommendation systems have been proposed [6]–[11].

Downie [12] introduced the discipline of music information retrieval (MIR), which is a multidisciplinary research endeavor that retrieves comprehensive data from music, including users' preferences as well as artworks, genres, beats, etc., in order to enhance the music-listening experience. MIR already has many real-world applications, powered by latest technologies ranging from computational musicology, audio analysis, and user preferences mining, to human-computer interaction [13]–[16].

Music sequence prediction is a particular type of recommendation system problem. Unlike usual recommendation system applications (e.g., Amazon and Netflix), predicting music sequences is more challenging and requires

application-specific tailoring. Considering how people listen to music in their daily lives, they usually play CDs or their playlists in Spotify repeatedly when they are in a car or at a gym [17]–[20]. However, it is unlikely that one would manually select songs or artists. Instead, users tend to repeat the same sequence of music, often in a pre-made playlist, unless one subscribed to a special recommendation service such as music station in Spotify. In this project, we verified these intuitions with some descriptive statistics analysis of training data and by comparing different models, discussed in the comparison section. The one important takeaway of this paper is that it doesn't deal with single factor such as music labels or sequential information separately but instead integrates all related factors together for the recommendation.

The rest of the paper is organized as follows: Section 2 provides a review of the dataset. Section 3 discusses state-of-the-art algorithms and our novel mixture model. Experiments are shown in Section 4. In Section 5, we outline limitations and provide future recommendations. Section 6 concludes the paper.

2 THE DATASET

We use the dataset provided by the *Kaggle Prediction Competition: Music Sequence Prediction* [21], which consists of artists played by different users on an online streaming service. In the dataset, each user has a sequence of 29 artists which they have played in order. Specifically, the dataset contains 972 rows and 29 columns where each row stands for each user and each column has an identification code of an artist. The order of columns corresponds to the order of music listening sequence of users. For example, if a user has artist id_1 , id_2 , and id_3 as values for column c_1 , c_2 , and c_3 respectively, it means that the user listens to music by the artist id_1 at first, and then id_2 and id_3 in his or her listening sequence. Note that the dataset only has information of artists but not of their songs. In the same way that one could have many

- Tao Li is with the Department of Computer Science, Purdue University, West Lafayette, IN 47907, USA.
E-mail: taoli@purdue.edu
- Minsoo Choi is with the School of Industrial Engineering, Purdue University, West Lafayette, IN 47907, USA.
- Kaiming Fu is with the School of Mechanical Engineering, Purdue University, West Lafayette, IN 47907, USA.
- Lei Lin is with the Goergen Institute for Data Science, University of Rochester, Rochester, NY 14627, USA.

Accepted to the 4th International Conference on Artificial Intelligence and Applications (AI 2018), October 27–28, 2018, Dubai, UAE.

songs of the same artist in their playlist, a sequence can have the same id more than once, and we do not differentiate the values, artist's id, no matter where they appear; that is, there is no difference between $artist_1$ appearing on the first column, c_1 on a playlist and $artist_1$ appearing on the tenth column c_{10} on the playlist. Lastly, the total number of artists in the pool is 3924.

The goal of the project is to predict the 30th artist that users will play, given the previous 29 playlists of their music listening sequence. We will give ten candidates of artists for the 30th, where the order of candidates has meaning regarding prediction accuracy. We assume that the earlier an artist appears, the more probable it is that the user plays the artist's song. It is reasonable to give a sequence of artists as the predictor instead of one target artist regarding that our ultimate objective is on recommendation not on exact prediction itself. For recommending purposes, figuring out preference tendencies of users would be more meaningful than just acquiring one prediction for the 30th artist they play. Accordingly, we used an evaluation metric which is devised to evaluate this sequential information of the predictors.

3 METHODS

The problem of music sequence prediction is intriguing because its a combination of techniques in both recommendation system and sequence prediction, which have been studied extensively. In this section, we first discuss those state-of-the-art methods and then introduce our approach to attack the new issue.

3.1 User-based Collaborative Filtering

Collaborative filtering (CF) is a machine learning algorithm for a user-specific recommendation system, which does not need extraneous information about users or items. Instead, CF utilizes internal user-item information which represents preferences for items by users. [22] Depending on the system, the user-item information can vary from ratings to purchase history. CF uses the information to find neighborhood and utilize it to predict a specific users preference or future behavior. Although there are many variants, CF-based algorithms fall into categories: user-based and item-based [23].

User-based CF assumes that, if some users whose past rating behaviors on other items are similar to user u , u tends to behave similarly to an item i . User-based CF, therefore, uses the similarity function s with respect to user dimension. $s : U \times U \rightarrow \mathbb{R}$. There are various kinds of similarity functions such as Pearson correlation and cosine similarity. Using this similarity between the user u and others, user-based CF computes a neighborhood which consists of the most similar k users. Then it computes the prediction of the user u 's preference to item i combining the ratings of the neighborhood. How many neighbors to select also depends on the problem.

$$P_{u,i} = \bar{r}_u + \frac{\sum_{u' \in N} s(u, u')(r_{u',i} - \bar{r}_{u'})}{\sum_{u' \in N} |s(u, u')|} \quad (1)$$

where \bar{r}_u is the average rating of the user u , $r_{u,i}$ is a rating of item i by user u , and $s(I, j)$ is the similarity score between user i and user j .

3.2 Item-based Collaborative Filtering

While the user-based CF algorithm is intuitive and effective, it is computationally expensive and has problems with handling sparsity. The most important reason is because the algorithm does and has to calculate similarity scores at the time when predictions or recommendations are needed; each operation takes $O(U)$. In a system which has large number of users and fast changing such as e-commerce, this cost could be prohibiting. Item-based CF solves the problem allowing the system to pre-compute and store the similarity scores and use the pre-stored data at the time when it predicts or recommends items to users. Item-based CF makes pre-computing possible using similarities between the rating patterns of items instead of users. The methods for computing item similarity are the same as user-based CF, except that the item-based CF computes with respect to items. The most popular similarity metrics is cosine similarity.

$$s(u, v) = \frac{\mathbf{r}_u \cdot \mathbf{r}_v}{\|\mathbf{r}_u\|_2 \cdot \|\mathbf{r}_v\|_2} \quad (2)$$

After collecting a similarity scores of items similar to i , item-based CF finds the most similar k neighbors as user-based CF. Then, it predicts user u 's rating with respect to the neighborhood as follows:

$$P_{u,i} = \bar{r}_u + \frac{\sum_{u' \in N} s(u, u')(r_{u',i} - \bar{r}_{u'})}{\sum_{u' \in N} |s(u, u')|} \quad (3)$$

For non-real-values ratings scales, such as purchase history data which can be represented as binary, we can compute pseudo-predictions with some simple tactics such as summation of all k similarities.

3.3 Hidden Markov Models

The Hidden Markov Model (HMM) is an extended version of the Markov Chain Model which is a stochastic process following Markovian property. Markovian property is an assumption on which the future state depends only on a current state which can be written as [24]

$$P(q_t | q_{t-1}, q_{t-2}, \dots, q_0) = P(q_t | q_{t-1}) \quad (4)$$

However, the Markov Model has a limitation in that it assumes each state corresponds to an observable event. To overcome the problem, HMM introduces a separation between observations and states of the system. That is, HMM assumes there is a latent and unobservable stochastic process within the system itself, but we can only observe another set of a stochastic processes as observations. With them, we could guess the hidden model of the system [25]. To be specific, we get the hidden model with the following process; decide model topology of the system such as the number of states in the model and the number of distinct observations; set parameters for the transition probability ($A=a_{ij}$), the observation symbol probability ($B=b_j(K)$), and initial state distribution $\pi = \pi_i$ which can be written as the following equations:

$$a_{ij} = P[q_{t+1} = S_i | q_t = S_j], 1 \leq i, j \leq N \quad (5)$$

$$b_j(k) = P[v_k \text{ at } t | q_t = S_j], 1 \leq j \leq N, 1 \leq k \leq M \quad (6)$$

$$\pi_i = P[q_1 = S_i], 1 \leq i \leq N \quad (7)$$

Next, with the initial parameters $\lambda = (A, B, \pi)$, we are going to keep re-estimating until we have no or little improvement regarding the probability of the observations given the parameters λ : $P[O|\lambda]$. The HMM has been applied to various kinds of sequential analysis from natural language processing to genetic sequence alignment [25], [26].

3.4 Mixture Models

The proposed mixture hidden Markov model (MHMM) is given as follow:

$$\text{MHMM}(n : \mathbb{D}) = \text{HMM}(n_1 : \mathbb{D}) \odot \text{CF}(n_2 : \mathbb{D}) \quad (8)$$

where \mathbb{D} are observed data, n_1 and n_2 are hyperparameters subject to $n = n_1 + n_2$, and \odot is a special operator we defined to concatenate two sequences. In practice, there are lots of tied rankings given by HMM and CF. A trick here is that we add weights to each artist, which are learned from the entire artists population - the more frequent a artist in the population, the higher weight he or she has. This trick fixes tied situations and is critical to overall performance which will be shown in section 4.

4 EXPERIMENTS

4.1 Evaluation Method

We use mean average precision at K (MAP@K) [27] to evaluate the prediction power of models. MAP@K is the average value of precision at K (AP@K), where AP@K is a score which gives higher value for a sequence which has actual value at smaller index [28]. For example, if actual value is 1 and two prediction sequences are $\{1, 2, 3, 4, 5\}$ and $\{2, 1, 3, 4, 5\}$, the first prediction has higher AP@K score than the second prediction since 1 comes earlier in the first prediction. The formula of MAP@K for this problem is

$$\text{MAP@K} = \frac{1}{N} \sum_{i=1}^N \frac{1}{f(y_i, \mathbb{D}_i)} \quad (9)$$

where N is the total number of users, y_i is the actual artist that user i listens to, \mathbb{D}_i is the predicted sequence of size K , and

$$f(y_i, \mathbb{D}_i) = \begin{cases} \text{index of } y_i \text{ in } \mathbb{D}_i, & y_i \in \mathbb{D}_i \\ 0, & y_i \notin \mathbb{D}_i \end{cases}$$

In our experiments, N is 972, K is 10, and y_i is the 30th artist that user i listens to.

4.2 Comparison

A comparison of different methods can be found in Table 1, where HF_{corpus} and $HF_{current}$ mean predicting with the highest frequent artists from the whole corpus and from the current list respectively, CF_{user} and CF_{item} represent user-based collaborative filtering and item-based collaborative filtering respectively, HMM is the hidden Markov model, and MHMM is the mixture model we propose. Results show our method significantly outperform state-of-the-art models. Figure 1 is the final leaderboard of the competition.

TABLE 1
Comparison of Different Models

Model	MAP@K
HF_{corpus}	0.00721
$HF_{current}$	0.10050
CF_{user}	0.01143
CF_{item}	0.01233
HMM	0.12838
MHMM	0.13958

#	Team Name	Score
1	Tao Li	0.14312
2	IG 88	0.13835
3	Darth Krayt	0.13817
4	Boba Fett	0.13702
5	Padme Amidala	0.11926
6	4-LOM	0.10882
7	Captain Rex	0.08109
8	Brakiss	0.01970

Fig. 1. Final leaderboard of the competition

5 LIMITATIONS AND FUTURE WORKS

Despite the success of our models in the competition and on real-world datasets, there are still many opportunities for further enhancements, including:

Similarity Metrics. For collateral filtering-based models, various cosine similarity-based measurements [29] have been proposed and the one used in our models are not necessarily optimal. *Resnik* [30] proposed an information-based measurement, semantic similarity, which might achieve more reasonable similarity results by using customized weights for different users based on information from other sources.

Markov Property. Although the mixture hidden Markov models have convincing performance in the dataset, the markov property might be too strong to assume for general cases. It is natural to come out that using information from previous n steps instead of only one step before to achieve better prediction. *Lafferty et al.* [31] introduce conditional random fields (CRFs) which takes context into consideration. *Sutton et al.* [32] provide a diagram of the relationship between HMMs and general CRFs as well as naive Bayes, logistic regression, linear-chain CRFs, and generative models.

Explore/Exploit Tradeoff. A central concern of adaptive learning system is the relation between the exploration of new possibilities and the exploitation of old certainties [33]. The explore/exploit (EE) tradeoff, since first introduced [34], it has been discussed extensively in the RecSys community. Although the models we used outperform others in the competition in terms of prediction power, the recommended items are restricted to a limited pool and fail to further explore new possibilities of user preference. As recommendation systems are usually built for commercial purposes, user experience is a major concern. *Gaver et al.* [35] argued the importance of users' non-instrumental needs, including diversion and surprise, which are needed to be addressed in software systems.

Future Techniques. With the advent and successes of recent deep learning techniques (e.g., CNNs and GANs [36]–[39]), great improvements have been achieved in computer vision and speech. Advances in these techniques provide new opportunities for music recommendation systems as a system can not only rely on offline information (e.g. music genre, user’s playing history) but is also possible to track users’ visual and auditory feedbacks in real-time. Facial expression detection, tracking, and classification of facial expressions have been widely studied [40]–[43] and even micro-expression detection [44], [45] for digging hidden emotions is becoming possible. Many efforts have been made from the music synthesis community [46]–[48] and a novel autoencoder framework was recently proposed [49]. We can expect customized real-time music recommendation systems that learn from human, compose for human, and even know human better than human themselves in the near future.

6 CONCLUSION

In this paper, we compare various state-of-the-art methods for music sequence predictions and propose a novel mixture hidden Markov model which shows promising results and significantly outperforms others in a Kaggle competition. We discuss limitations of similarity metrics, Markov-based models and explore/exploit trade-offs. Furthermore, we point out future directions for next-generation music recommendation systems.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Bruno Ribeiro for interesting lectures and well-designed projects in *Purdue CS 573, Spring 2018* [50]. This work was supported by a grant from the Goodata Foundation (GDF grant b5d-8d-564-f0).

REFERENCES

- [1] M. A. 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] A. Van den Oord, S. Dieleman, and B. Schrauwen, “Deep content-based music recommendation,” in *Advances in neural information processing systems*, 2013, pp. 2643–2651.
- [3] “Global music report 2018,” 2018. [Online]. Available: <http://www.ifpi.org/downloads/GMR2018.pdf>
- [4] T. Li, L. Lin, M. Choi, K. Fu, S. Gong, and J. Wang, “Youtube av 50k: an annotated corpus for comments in autonomous vehicles,” *arXiv preprint arXiv:1807.11227*, 2018.
- [5] T. Li, M. Choi, and L. Lin, “Youtubestat.py: a python module to download youtube statistics and rankings,” *GitHub Repository*, pp. 1–3, June 2017. [Online]. Available: <https://github.com/Eroica-cpp/YouTube-Statistics>
- [6] H.-C. Chen and A. L. Chen, “A music recommendation system based on music data grouping and user interests,” in *Proceedings of the tenth international conference on Information and knowledge management*. ACM, 2001, pp. 231–238.
- [7] W. Hicken, F. Holm, J. Clune, and M. Campbell, “Music recommendation system and method,” Feb. 17 2005, uS Patent App. 10/917,865.
- [8] H.-S. Park, J.-O. Yoo, and S.-B. Cho, “A context-aware music recommendation system using fuzzy bayesian networks with utility theory,” in *International conference on Fuzzy systems and knowledge discovery*. Springer, 2006, pp. 970–979.
- [9] V. G. Alcalde, C. M. L. Ulloa, A. T. Bonet, A. T. Llopis, J. S. Marcos, D. C. Ysern, and D. Arkwright, “Method and system for music recommendation,” Jul. 25 2006, uS Patent 7,081,579.
- [10] Y. Song, S. Dixon, and M. Pearce, “A survey of music recommendation systems and future perspectives,” in *9th International Symposium on Computer Music Modeling and Retrieval*, vol. 4, 2012.
- [11] P. Hoffmann, A. Kaczmarek, P. Spaleniak, and B. Kostek, “Music recommendation system,” *Journal of telecommunications and information technology*, 2014.
- [12] J. S. Downie, “Music information retrieval,” *Annual review of information science and technology*, vol. 37, no. 1, pp. 295–340, 2003.
- [13] B. Bel and B. Vecchione, “Computational musicology,” *Computers and the Humanities*, vol. 27, no. 1, pp. 1–5, 1993.
- [14] S. Pfeiffer, S. Fischer, and W. Effelsberg, “Automatic audio content analysis,” in *Proceedings of the fourth ACM international conference on Multimedia*. ACM, 1997, pp. 21–30.
- [15] S. Holland, M. Ester, and W. Kießling, “Preference mining: A novel approach on mining user preferences for personalized applications,” in *European Conference on Principles of Data Mining and Knowledge Discovery*. Springer, 2003, pp. 204–216.
- [16] A. Dix, “Human-computer interaction,” in *Encyclopedia of database systems*. Springer, 2009, pp. 1327–1331.
- [17] L. Baltrunas, M. Kaminskas, B. Ludwig, O. Moling, F. Ricci, A. Aydin, K.-H. Lüke, and R. Schwaiger, “Incarmusic: Context-aware music recommendations in a car,” in *International Conference on Electronic Commerce and Web Technologies*. Springer, 2011, pp. 89–100.
- [18] C. Wang, S. Gong, A. Zhou, T. Li, and S. Peeta, “Cooperative adaptive cruise control for connected autonomous vehicles by factoring communication-related constraints,” in *The 23rd International Symposium on Transportation and Traffic Theory*, 2019.
- [19] T. Li, “Modeling uncertainty in vehicle trajectory prediction in a mixed connected and autonomous vehicle environment using deep learning and kernel density estimation,” in *The Fourth Annual Symposium on Transportation Informatics*, 2018.
- [20] S. Gong, A. Zhou, J. Wang, T. Li, and S. Peeta, “Cooperative adaptive cruise control for a platoon of connected and autonomous vehicles considering dynamic information flow topology,” in *The 21st IEEE International Conference on Intelligent Transportation Systems*, 2018.
- [21] “Kaggle prediction competition: Music sequence prediction,” 2018. [Online]. Available: <https://www.kaggle.com/c/music-sequence-prediction>
- [22] Y. Koren and R. Bell, “Advances in collaborative filtering,” in *Recommender systems handbook*. Springer, 2015, pp. 77–118.
- [23] M. D. Ekstrand, J. T. Riedl, J. A. Konstan *et al.*, “Collaborative filtering recommender systems,” *Foundations and Trends® in Human-Computer Interaction*, vol. 4, no. 2, pp. 81–173, 2011.
- [24] P. A. Gagniuc, *Markov Chains: From Theory to Implementation and Experimentation*. John Wiley & Sons, 2017.
- [25] L. R. Rabiner, “A tutorial on hidden markov models and selected applications in speech recognition,” *Proceedings of the IEEE*, vol. 77, no. 2, pp. 257–286, 1989.
- [26] R. Hughley and A. Krogh, “Hidden markov models for sequence analysis: extension and analysis of the basic method,” *Bioinformatics*, vol. 12, no. 2, pp. 95–107, 1996.
- [27] W. Kan, “Map@k demo,” 2016. [Online]. Available: <https://www.kaggle.com/wendykan/map-k-demo>
- [28] C. D. Manning, P. Raghavan, and H. Schütze, “Chapter 8: Evaluation in information retrieval,” *Introduction to information retrieval*, 2008.
- [29] M. Li, X. Chen, X. Li, B. Ma, and P. M. Vitányi, “The similarity metric,” *IEEE transactions on Information Theory*, vol. 50, no. 12, pp. 3250–3264, 2004.
- [30] P. Resnik, “Semantic similarity in a taxonomy: An information-based measure and its application to problems of ambiguity in natural language,” *Journal of artificial intelligence research*, vol. 11, pp. 95–130, 1999.
- [31] J. Lafferty, A. McCallum, and F. C. Pereira, “Conditional random fields: Probabilistic models for segmenting and labeling sequence data,” 2001.
- [32] C. Sutton and A. McCallum, *An introduction to conditional random fields for relational learning*. MIT Press, 2006, vol. 2.
- [33] J. G. March, “Exploration and exploitation in organizational learning,” *Organization science*, vol. 2, no. 1, pp. 71–87, 1991.
- [34] J. A. Schumpeter, “The theory of economic development,” 1934.

- [35] B. Gaver and H. Martin, "Alternatives: exploring information appliances through conceptual design proposals," in *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. ACM, 2000, pp. 209–216.
- [36] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, p. 436, 2015.
- [37] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [38] L. Lin, S. Gong, and T. Li, "Deep learning-based human-driven vehicle trajectory prediction and its application for platoon control of connected and autonomous vehicles," in *The Autonomous Vehicles Symposium*, vol. 2018, 2018.
- [39] T. Li, K. Fu, M. Choi, X. Liu, and Y. Chen, "Toward robust and efficient training of generative adversarial networks with bayesian approximation," in *the Approximation Theory and Machine Learning Conference*, 2018.
- [40] P. Rodriguez, G. Cucurull, J. Gonzalez, J. M. Gonfaus, K. Nasrollahi, T. B. Moeslund, and F. X. Roca, "Deep pain: Exploiting long short-term memory networks for facial expression classification," *IEEE transactions on cybernetics*, no. 99, pp. 1–11, 2017.
- [41] A. T. Lopes, E. de Aguiar, A. F. De Souza, and T. Oliveira-Santos, "Facial expression recognition with convolutional neural networks: coping with few data and the training sample order," *Pattern Recognition*, vol. 61, pp. 610–628, 2017.
- [42] N. Zeng, H. Zhang, B. Song, W. Liu, Y. Li, and A. M. Dobaie, "Facial expression recognition via learning deep sparse autoencoders," *Neurocomputing*, vol. 273, pp. 643–649, 2018.
- [43] T. Li, "Single image super-resolution: A historical review," *ObEN Research Seminar*, August 2018.
- [44] X. Li, H. Xiaopeng, A. Moilanen, X. Huang, T. Pfister, G. Zhao, and M. Pietikainen, "Towards reading hidden emotions: A comparative study of spontaneous micro-expression spotting and recognition methods," *IEEE Transactions on Affective Computing*, 2017.
- [45] M. Takalkar, M. Xu, Q. Wu, and Z. Chaczko, "A survey: facial micro-expression recognition," *Multimedia Tools and Applications*, vol. 77, no. 15, pp. 19 301–19 325, 2018.
- [46] C. Dodge and T. A. Jerse, *Computer music: synthesis, composition and performance*. Macmillan Library Reference, 1997.
- [47] C.-i. Wang and S. Dubnov, "Guided music synthesis with variable markov oracle," in *Tenth Artificial Intelligence and Interactive Digital Entertainment Conference*, 2014.
- [48] M. O'Sullivan, B. Srbinovski, A. Temko, E. Popovici, and H. McCarthy, "V2hz: Music composition from wind turbine energy using a finite-state machine," in *Signals and Systems Conference (ISSC), 2017 28th Irish*. IEEE, 2017, pp. 1–6.
- [49] J. Colonel, C. Curro, and S. Keene, "Improving neural net auto encoders for music synthesis," in *Audio Engineering Society Convention 143*. Audio Engineering Society, 2017.
- [50] B. Ribeiro, "Cs57300: Graduate data mining," 2018. [Online]. Available: <https://www.cs.purdue.edu/homes/ribeiro/courses/Spring2018/>