Analyzing and Classifying Guitarists from Rock Guitar Solo Tablature

Orchisama Das
CCRMA,
Stanford University
orchi@ccrma.stanford.edu

Blair Kaneshiro
CCRMA,
Stanford University
blairbo@ccrma.stanford.edu

Tom Collins
Department of Psychology,
Lehigh University
tomthecollins@gmail.com

ABSTRACT

Guitar solos provide a way for guitarists to distinguish themselves. Many rock music enthusiasts would claim to be able to identify performers on the basis of guitar solos, but in the absence of veridical knowledge and/or acoustical (e.g., timbral) cues, the task of identifying transcribed solos is much harder. In this paper we develop methods for automatically classifying guitarists using (1) beat and MIDI note representations, and (2) beat, string, and fret information, enabling us to investigate whether there exist "fretboard choreographies" that are specific to certain artists. We analyze a curated collection of 80 transcribed guitar solos from Eric Clapton, David Gilmour, Jimi Hendrix, and Mark Knopfler. We model the solos as zero and first-order Markov chains, and do performer prediction based on the two representations mentioned above, for a total of four classification models. Our systems produce above-chance classification accuracies, with the first-order fretboard model giving best results. Misclassifications vary according to model but may implicate stylistic differences among the artists. The current results confirm that performers can be labeled to some extent from symbolic representations. Moreover, performance is improved by a model that takes into account fretboard choreographies.

1. INTRODUCTION

Avid listeners of rock music claim they can easily distinguish between a guitar solo by Jimi Hendrix versus Jimmy Page. This raises many questions about the types of features underlying such a task. For example, can artist identification of guitar solos be performed successfully from compositional features alone; or are other performance and timbral cues required?

Artist identification is an established research topic in Music Information Retrieval (MIR). Timbral features extracted from audio representations have been used for artist recognition [1–3] and for singer identification in popular music [4, 5].

Identification of artists/composers from symbolic representations (digital encodings of staff notation) has also been attempted [6–11]. Kaliakatsos-Papakostas et al. used a weighted Markov chain model trained on MIDI files for composer identification [8], as well as feedforward neural networks [12]. Markov models have been used to distinguish between Mozart and Haydn [9]. Existing work on feature extraction from symbolic music is extremely valuable for such a classification task. For example, Pienimaki et al. describe an automatic cluster analysis method for symbolic music analysis [13], while Collins et al. propose computational methods for generating music in the style of various composers [14, 15].

Past studies have modeled rhythm and lead content of guitar parts. Of particular relevance is work by McVicar et al. [16–18], in which models are trained to emulate playing styles of various guitarists such as Jimi Hendrix, Jimmy Page, and Keith Richards. The output is a stylistic generation of rhythm and lead guitar tablature based on string and fret rather than staff notation representations. It is unknown, however, whether this choice of representation confers any analytic or compositional advantage. A single MIDI note number (MNN) can be represented by several different (string, fret)-pairs on the fretboard, and it could be that such choices vary systematically from one artist to another. Methods for separating voices in lute tablature seemed to benefit from such a tablature-based representation [19].

In addition, Ferretti has modeled guitar solos as directed graphs and analyzed them with complex network theories to yield valuable information about playing styles of musicians [20]. Another study by Cherla et al. automatically generated guitar phrases by directly transcribing pitch and onset information from audio data and then using their symbolic representations for analysis [21].

To our knowledge, the task of identifying artists from guitar solos has not been attempted previously. Furthermore, McVicar et al.'s [18] work raises the question of whether fretboard representations are really more powerful than staff notation representations and associated numeric encodings (e.g., MIDI note numbers). In support of McVicar et al.'s [18] premise, research in musicology alludes to specific songs and artists having distinctive "fretboard choreographies" [22], but the current endeavor enables us to assess such premises and allusions quantitatively.

Widmer [23] is critical of the prevalence of Markov models in music-informatic applications, since such models lack incorporation of long-term temporal dependencies that most musicologists would highlight in a given piece. Collins et al. [15], however, show that embedding Markov chains in a system that incorporates such long-term dependencies is sufficient for generating material that is in some circumstances indistinguishable from human-composed excerpts. Whether the zero and first-order Markov models used in the present study are sufficient to identify the provenance of guitar solos is debatable; however, we consider them a reasonable starting point for the task at hand.

The rest of this paper is organized as follows. We describe the dataset and features, Markov models and maximum likelihood interpretations, and our classification procedure in Section 2. In Section 3 we visualize our data and report classification results. We conclude in Section 4 with discussion of results, insights into stylistic differences among the artists, potential issues, and avenues for future research.

2. METHOD

2.1 Dataset

We collated our own dataset for the present study, since no pre-existing dataset was available. First, we downloaded guitar tabs in GuitarPro format from UltimateGuitar. [1] The quality of tabs was assessed by us as well as the number of stars they received from UltimateGuitar users. Any tab with a rating below four stars was discarded. We then manually extracted the guitar solos from each song's score and converted them to MusicXML format with the free TuxGuitar software. [2] In total, our final dataset comprised 80 solos—20 each from Eric Clapton, David Gilmour, Jimi Hendrix, and Mark Knopfler. While the size of this dataset is in no way exhaustive, the number of songs curated was restricted by the availability of accurate tabs.

2.2 Representations

For parsing MusicXML data and symbolic feature extraction, we used a publicly available JavaScript library. [3] Using methods in this library, we wrote a script that returns ontime (symbolic onset time), MIDI note number (MNN), morphetic pitch number (MPN), note duration, string number, and fret number for each note in the solo. To obtain the beat of the measure on which each note begins, we took its ontime modulo the time signature of that particular solo. The tonic pitch of each song was identified from the key signature using an algorithm in the JavaScript library that finds the tonic MIDI note closest to the mean of all pitches in a song. We then subtracted this tonic MNN from each raw MNN to give a "centralized MNN", which accounted for solos being in different keys. When calculating pitch class, we took centralized MNN modulo 12 to limit values to the range [0, 11].

For guitar players, fingering positions on the fretboard are crucial. To account for variability in key along the fretboard, solos were transposed to the nearest C major/A minor fretboard position on the same string, making sure there were no negative frets. If a fret number was greater than or equal to 24 (the usual number of frets on an electric guitar), it was wrapped back around to the start of the fretboard by a modulo 24 operation, resulting in the fret range [0, 23]. The resulting dimensions of beat, MNN, pitch class, string and transposed fret were saved in JSON format for each song in the dataset. Finally, we generated two types of tuples on a per-note basis as our state spaces: the first state space comprises beat and centralized MNN, denoted (beat, MNN) hereafter; the second comprises beat, string, and transposed fret, denoted (beat, string, fret) hereafter. The quarter note is represented as a single beat. For example, an eighth note played on the fifth fret of the second string would be $(0.5, 64)$ in the 2D (beat, MNN) representation and $(0.5, 2, 5)$ in the 3D (beat, string, fret) representation.

2.3 Markov Model

A Markov model is a stochastic model of processes in which the future state depends only on the previous n states [24]. Musical notes can be modeled as random variables that vary over time, with their probability of occurrence depending on the previous n notes.

In the present classification paradigm, let x_i represent a state in a Markov chain at time instant i. In a first-order Markov model, there is a transition matrix $\mathbf{P}$ which gives the probability of transition from x_i to x_{i+1} for a set of all possible states. If $\{x_1, x_2, ...x_N\}$ is the set of all possible states, the transition matrix $\mathbf{P}$ has dimensions $N \times N$.

Given a new sequence of states $[x_1, x_2, ...x_T]$, we can represent it as a path with a probability of occurrence $P(x_1, ..., x_T)$. According to the product rule, this joint probability distribution can be written as:

$$P(x_1, ..., x_T) = P(x_T|x_{T-1}, ..., x_1)P(x_1, ..., x_{T-1}) \tag{1}$$

Since the conditional probability $P(x_T|x_{T-1}, ..., x_1)$ in a first-order Markov process reduces to $P(x_T|x_{T-1})$, we can write:

$$P(x_1, x_2, ...x_T) = P(x_T|x_{T-1})P(x_1, x_2, ..x_{T-1}) \tag{2}$$

Solving this recursively brings us to:

$$P(x_1, x_2, ...x_T) = P(x_1)\prod_{i=2}^{T} P(x_i|x_{i-1}) \tag{3}$$

Taking the log of $P(x_1, x_2, ...x_T)$ gives us the log likelihood, L_1 defined as:

$$L_1 = \log P(x_1) + \sum_{i=2}^{T} \log P(x_i|x_{i-1}) \tag{4}$$

Hence, the log likelihood can be calculated from the transition matrix $\mathbf{P}$ and initial distribution $P(x_1)$.

For a zero order Markov model, the joint distribution is simply the product of the marginal distributions because the present state is independent of any of the past states:

$$P(x_1, x_2, ...x_T) = \prod_{i=1}^{T} P(x_i) \tag{5}$$

[1] https://www.ultimate-guitar.com/
[2] https://sourceforge.net/projects/tuxguitar/
[3] https://www.npmjs.com/package/maia-util

Therefore, the log likelihood is defined as

$$L_0 = \sum_{i=1}^{T} \log P(x_i) \tag{6}$$

2.4 Classification Procedure

For the present analysis, we performed classifications using a leave-one-out paradigm, i.e, we trained on all 79 songs except the song being classified, and repeated this for all 80 songs in our dataset.[4] We used two different state spaces in our analysis: the 2D state space comprising beat and MNN, and the 3D state space comprising beat, string, and transposed fret. Each state in a state space represents one note in a guitar solo. We then trained zero-order and first-order Markov models on these data, and used a maximum likelihood approach to classify each song.

For zero-order models, a probability list was constructed by obtaining the probability of occurrence of each unique state in the training data. This was done empirically by counting the number of times a unique state occurred, and dividing this value by the total number of occurrences of all states. A new set of states was obtained for an unseen song, and their probabilities were individually looked up in the probability list. If a state is previously unseen, we give the probability an arbitrarily small value (0.00005) for that state. The likelihood function L_0 was calculated for each artist, and the artist with maximum likelihood was selected.

For first-order models, the transition matrix **P** for a particular guitarist was found by training on all songs played by him except the song being classified. Once we computed **P** for each guitarist, we calculated the probability of finding the sequence of states observed in the unseen test song for each artist, and chose the artist whose transition matrix maximized the likelihood L_1, according to

$$\text{artist} = \arg\max_{a \in \mathcal{A}} L_{0,1}(a), \tag{7}$$

where $\mathcal{A} = \{\text{Clapton, Gilmour, Hendrix, Knopfler}\}$. Under a null hypothesis of chance-level classification, we assume a binomial distribution—having parameters $p = 1/\text{nClasses}$ ($= 0.25$), $k = \text{nSuccesses}$, and $n = \text{nTrials}$ ($= 80$)—for calculation of p-values. We correct for multiple comparisons using False Discovery Rate [25].

3. RESULTS

3.1 Data Visualization

3.1.1 Histograms

To better understand the choice of notes used by each guitarist, we plotted pitch class and note duration histograms. As shown in Figure 1, the pitch class distributions shed light on certain modes and scales that are used most frequently by the artists. For instance, the minor pentatonic scale (1, ♭3, 4, 5, ♭7) stands out prominently for all four artists. It is also interesting to note that Eric Clapton uses the 6th scale step more frequently than others. We performed chi-squared tests on all pitch class distributions to assess uniformity, with the null hypothesis being that the pitch class distributions are uniform and the alternate hypothesis being that they are not. The p-values were negligibly small, suggesting that none of the distributions are uniform, but we report the χ^2 statistic to see which distributions are *relatively* more uniform. Eric Clapton had the largest χ^2 value (5561.4), followed closely by Jimi Hendrix (4992.8), and then David Gilmour (3153.8) and Mark Knopfler (2149.7). A smaller chi-squared value indicates less distance from the uniform distribution, providing evidence that Knopfler is more exploratory in his playing style because he makes use of a wider variety of pitch classes. Note duration histograms indicate that all artists prefer the sixteenth note (0.125) and the eighth note (0.5) except for Knopfler, who appears to use more triplets ($0.333, 0.167$). Knopfler's exploratory use of pitch classes and triplets may be related. He may use more chromatic notes in the triplets to fill in the intervals between main beats. This could potentially set him apart from the other guitarists.

3.1.2 Self-similarity Matrices

To observe similarity between artists, we calculate a self-similarity matrix for each state space. To do so, we form vectors for each artist, denoted by $\mathbf{a}^{(1)}$, $\mathbf{a}^{(2)}$, $\mathbf{a}^{(3)}$, $\mathbf{a}^{(4)}$, with each element in the vector representing a (beat, MNN) or (beat, string, fret) state. To define similarity between artists, we use the Jaccard index [17]. Each element in the similarity matrix **S** is given as:

$$\mathbf{S}_{i,j} = \frac{|\mathbf{a}^{(i)} \cap \mathbf{a}^{(j)}|}{|\mathbf{a}^{(i)} \cup \mathbf{a}^{(j)}|} \tag{8}$$

where $|.|$ indicates set cardinality. The self-similarity matrices for both state spaces are shown in Figure 2. We observe that similarity among artists is slightly lower when we use a 3D state space with fretboard information. Among the artists, we observe that the similarity between Mark Knopfler and David Gilmour is highest, indicating the possibility for confusion between these artists.

3.1.3 Markov Chains

In Figure 3, we show the transition matrices for each artist as weighted graphs using R's *markovchain* package [26]. The vertices represent states in the transition matrix and edges represent transition probabilities. Sparsity is calculated as the ratio of number of zero entries to the total number of entries in the transition matrix. A combination of high sparsity with a large number of vertices (more unique states) indicates less repetition in the solos. While the details in these visualizations are not clear, some useful parameters of the graphs are given in Table 1. Although the difference in sparsity is relatively small between artists, it could play a significant role in classification. As expected, the transition matrices for the 2D state spaces are less sparse than their 3D counterparts. The 2D state space also contains fewer unique states, which makes intuitive

[4] The dataset consisting of JSON files, code for classification and analysis are all included in `https://github.com/orchidas/Guitar-Solo-Classification`

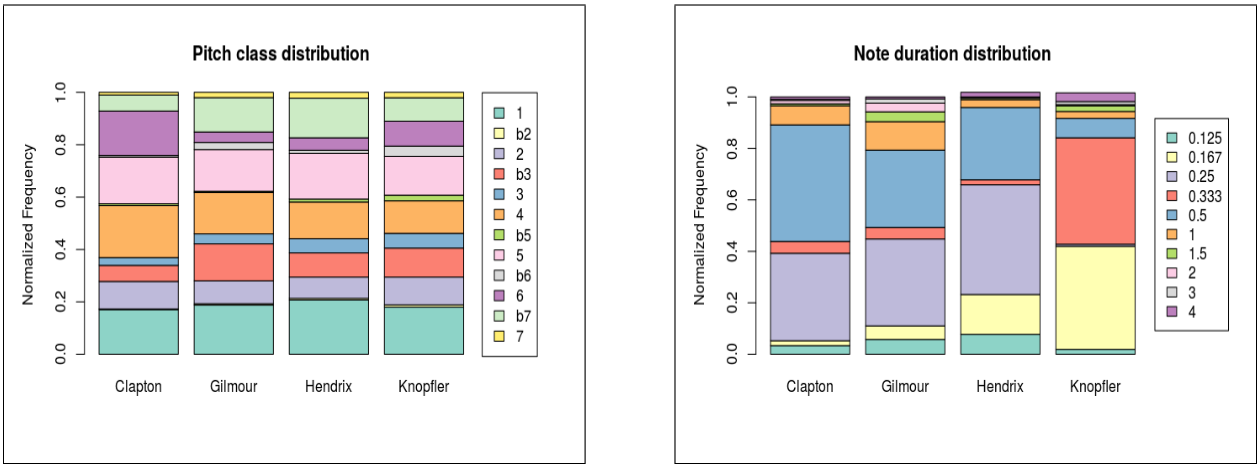

Figure 1. Left: Normalized pitch class histograms. Right: Note duration histograms.

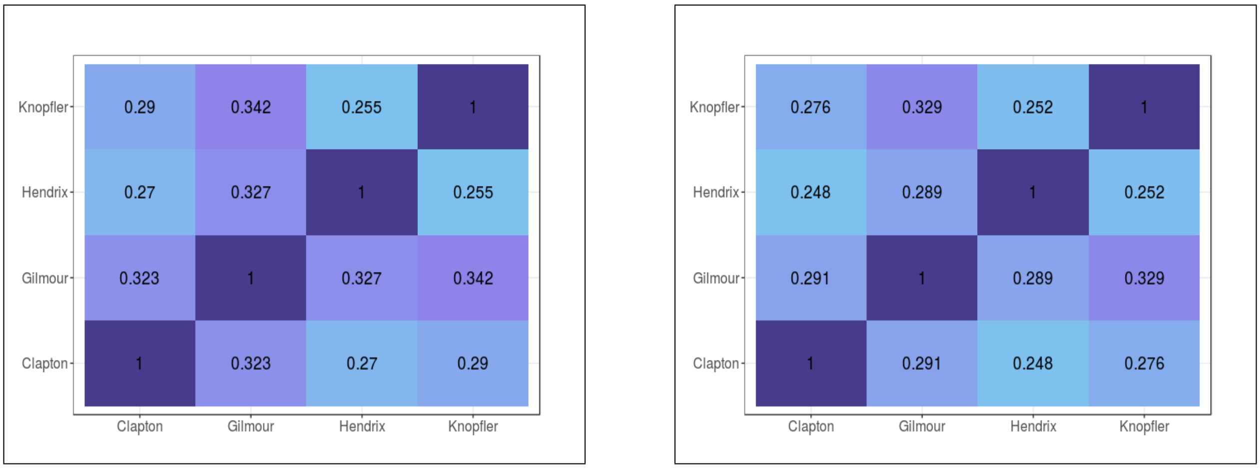

Figure 2. Artist self-similarity matrices. Left: State space of (beat, MNN). Right: State space of (beat, string, fret).

sense because a single MNN can be represented by several different sets of (string, fret) tuples on the fretboard. We observe that Hendrix has the largest number of unique states and his transition matrix is sparsest, indicating that he is least repetitive among the artists. Knopfler, on the other hand, has fewer unique states and more transitions, which means he repeats similar patterns of notes in his solos. It is curious how this analysis complements our interpretation of the histograms—even though Knopfler was shown to employ relatively unusual pitch and rhythmic material in Section 3.1.1, the current analysis shows he does so in a more repetitive manner, while Hendrix is vice versa.

	Clapton	Gilmour	Hendrix	Knopfler
Vertices	921	913	1751	889
Edges	3437	3616	5704	3564
Sparseness	0.9959	0.9956	0.9981	0.9955

	Clapton	Gilmour	Hendrix	Knopfler
Vertices	1443	1536	2679	1414
Edges	3817	4044	6241	4028
Sparseness	0.9981	0.9982	0.9991	0.9979

Table 1. Properties of transition matrices for (beat, MNN) space (top) and (beat, string, fret) space (bottom).

3.2 Classification Results

We performed classification tests on four models:

- Zero-order model, state space of (beat, MNN);
- Zero-order model, state space of (beat, string, transposed fret);
- First-order model, state space of (beat, MNN);
- First-order model, state space of (beat, string, transposed fret).

The classification results and overall accuracy are given in the four confusion matrices of Table 2. All classification accuracies are significantly above chance at the .05 level after correction for multiple comparisons. The first-order Markov model with 3D state space (beat, string, transposed fret) performs best, with a classification accuracy of 50%. This confirms that a model comprising temporal information and preserving fretboard information is the best choice

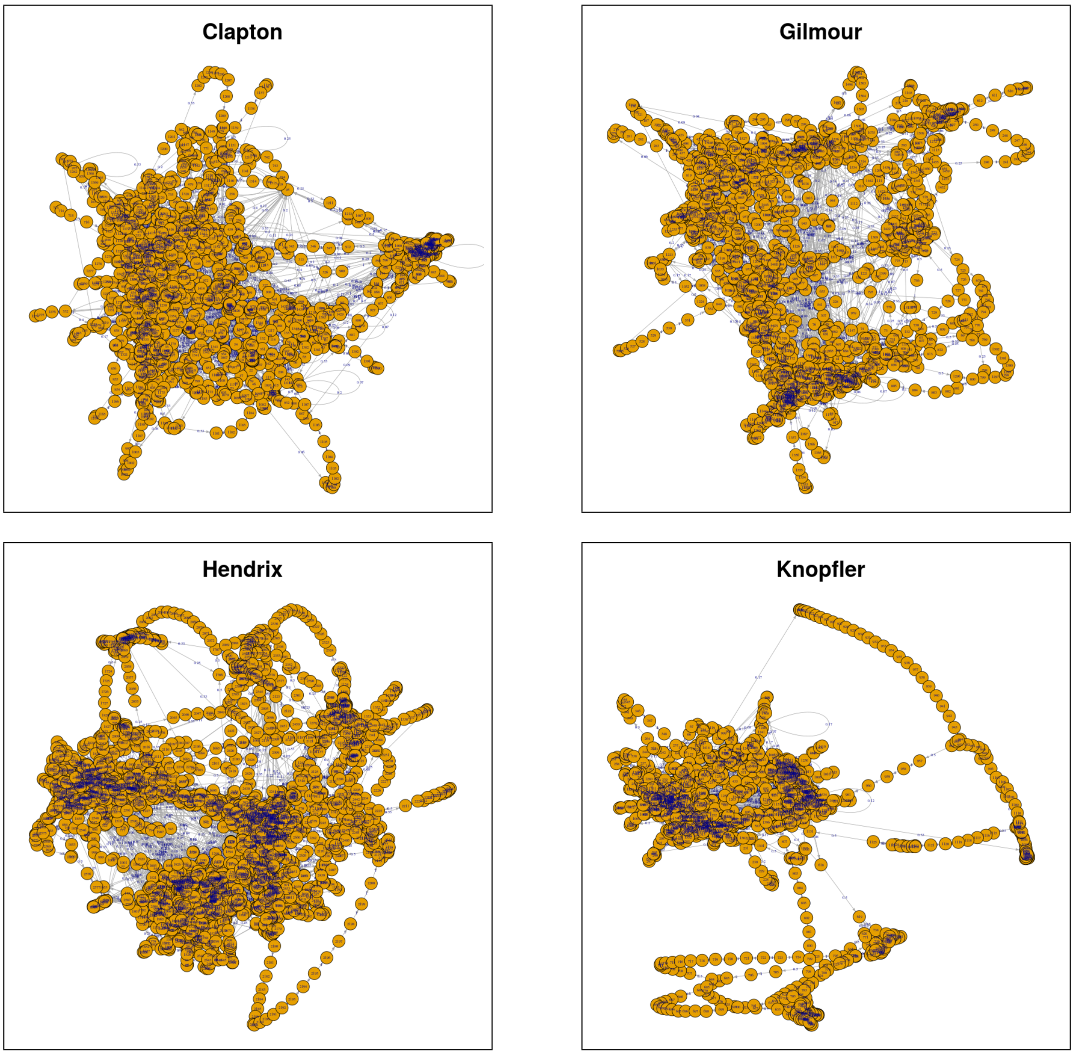

Figure 3. Markov chain visualizations of guitar solos with a 3D state space of (beat, string, fret).

for this classification problem.

4. DISCUSSION

4.1 Interpreting Results

In this study we have proposed a novel classification problem and used simple yet intuitive Markov models to achieve automatic classification of guitarists. Motivated by claims in the literature about (1) guitarists' "fretboard choreographies" [22], and (2) fretboard information leading to stronger models than ones built on pitch-based representations alone [18, 19], we considered two different state spaces—a 2D state space comprising beat and transposed MIDI note, and a 3D state space comprising beat, string, and transposed fret information. Pitch class and duration histograms reveal basic stylistic differences between artists. Classification was significantly above chance for all models, with performance strongest for the first-order Markov model built on beat and fretboard information, substantiating the claims about the efficacy of tablature space. Results also highlight some useful inferences about the playing styles of the guitarists.

We observe from the classifier output that Clapton always has the lowest number on the diagonal of the confusion matrices, which means he is particularly hard to classify correctly. This may be because each of his solos is distinctly different from the others, and it is hard to detect previously seen states in them. Although his transition matrix is not the sparsest, his transition probabilities themselves are rather small.

The models that include MNN as a state space are able to classify Gilmour more accurately. In fact, the classification for such models is biased toward Gilmour, i.e, more artists are misclassified as Gilmour. There may exist a number of (beat, MNN) states where Gilmour's transition probabilities dominate. If any of these states are detected in an unseen song, the likelihood function will be heavily biased toward Gilmour even if the true artist is different. This is also due to the fact that we lose information by representing pitches merely as MIDI notes. For example, the note C4 represented by MNN 48 can be represented in (string, fret) format as $(2, 1)$, $(3, 5)$, $(4, 10)$, $(5, 15)$ or $(6, 20)$. In a state-space model that includes string and fret information, these would be five unique states, but in the MNN model

(a) Zero order, (beat, MNN)
Overall accuracy–35.00%

P \ R	C	G	H	K
C	2	0	2	2
G	6	14	5	9
H	6	3	5	2
K	6	3	8	7

(b) Zero order, (beat,string,fret)
Overall accuracy–32.50%

P \ R	C	G	H	K
C	3	0	3	0
G	7	9	6	9
H	4	2	5	2
K	6	9	6	9

(c) First order, (beat, MNN)
Overall accuracy–43.75%

P \ R	C	G	H	K
C	4	0	3	0
G	4	10	1	7
H	7	4	12	4
K	5	6	4	9

(d) First order, (beat,string,fret)
Overall accuracy–50.00%

P \ R	C	G	H	K
C	6	0	4	0
G	2	8	1	4
H	6	3	12	2
K	6	9	3	14

Table 2. Confusion matrices for all classification models, where "P" stands for prediction, "R" for reference, and "C", "G", "H", and "K" our four artists Clapton, Gilmour, Hendrix, and Knopfler, respectively.

they would all denote the same state.

The confusion matrices also show notable confusion between Knopfler and Gilmour. This makes sense based upon the artist self-similarity matrices. The two first-order models classify Jimi Hendrix and Mark Knopfler well. One may expect Knopfler to have a high classification accuracy since he was shown to have the most repetitive pattern of states. Hendrix's high classification accuracy is more surprising, but can be explained. Although Hendrix has a large number of unique states in his graph, his transition probabilities between those states are relatively high. If these states with high transition probabilities are found in a new solo, Hendrix's log likelihood will have a high value for that solo.

Ultimately, we conclude that the first-order Markov model with 3D state space performs best because it captures fretboard choreographies that are somewhat unique to guitarists. Musically, this makes sense because guitarists are likely to repeat certain "licks" and riffs. These licks consist of notes dependent on each other and linked in a sequential manner, which a first-order Markov model can capture. It is to be noted that a model built on a 3D state space will *not* necessarily outperform one built on a 2D state space, as we can see from the overall accuracy of the zero-order models. Higher-dimensional spaces are sparser, and more zero probabilities can lead to inaccurate classifications. The relationship between dimensionality and sparsity is more complex and calls for detailed study which is beyond the scope of this paper.

4.2 Limitations

As ever with work on symbolic music representations, the dataset could be enlarged and could contain a greater number of lesser-known songs. Second, the tabs have been transcribed by different people and may not be exactly accurate, so there exists some noise in the data, especially as we have chosen songs on the basis of average review rating (e.g., not by the quantity of reviews). Moreover, the transcribers may have included their own stylistic gestures during transcription, which were unrelated to the guitarist being analyzed but could have affected or even confounded classification performance [27]. This was one of the reasons we chose not to incorporate techniques such as bends, hammer-ons and pull-offs into our representations and classification systems, although in its most accurate form such information could lead to superior results. While there are more than 20 songs by each artist, we found overall that increasing the number of songs per artist resulted in lower-rated tabs.

As note dependencies extend beyond closest neighbors, a first-order Markov chain may be too simple to model music [15, 23]. Using musical phrases as states in an auxiliary Markov model could yield better results, although it would be challenging to define when two phrases should be considered "the same" for the purposes of defining states.

One might argue that we claim fretboard representations are an important feature in distinguishing guitarists, yet we transpose all solos to bring them to the same key. Transposition is justified because fretboard preferences among guitarists are largely independent of key. Without transposition, the probability of finding similar states in different solos is very low, yielding even sparser transition matrices, which makes classification an even more difficult task.

4.3 Future Work

While the present findings are only an initial step for this topic, they point to several directions for future research. Identifying and distinguishing guitarists according to their musical choices is useful because deeper understanding of their style can aid in algorithmic composition and shed light on what constitutes a stylistically successful guitar solo. Music recommendation is another area where the present research could have applications—for example, someone who likes Eric Clapton may also enjoy B.B. King, since both guitarists often exhibit slow, melodic playing styles; but is not as likely to enjoy Eddie Van Halen, who is known mostly for being a shredder.

In the future, we would like to extend the current approach to more data, and more complex and hybrid machine learning methods. Finally, the present analyses considered only symbolic features. As remarked above, it is likely that acoustic features also contribute to the successful human classification of guitar solos. It will therefore be interesting to compare the present results to performance of a classification system based on acoustic features, or combined acoustic and symbolic features. Assessing human identification of guitar solos based on symbolic features alone (e.g., using MIDI performances of the solos) could further help to disentangle the contributions of compositional and acoustical cues.

5. REFERENCES

[1] H. Eghbal-Zadeh, M. Schedl, and G. Widmer, "Timbral modeling for music artist recognition using i-vectors," in *Signal Processing Conference (EU-*

SIPCO), 2015 23rd European. IEEE, 2015, pp. 1286–1290.

[2] B. Whitman, G. Flake, and S. Lawrence, "Artist detection in music with Minnowmatch," in *Neural Networks for Signal Processing XI, 2001. Proceedings of the 2001 IEEE Signal Processing Society Workshop*. IEEE, 2001, pp. 559–568.

[3] M. I. Mandel and D. P. W. Ellis, "Song-level features and support vector machines for music classification," in *ISMIR*, vol. 2005, 2005, pp. 594–599.

[4] W. H. Tsai and H. M. Wang, "Automatic singer recognition of popular music recordings via estimation and modeling of solo vocal signals," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 1, pp. 330–341, 2006.

[5] A. Mesaros, T. Virtanen, and A. Klapuri, "Singer identification in polyphonic music using vocal separation and pattern recognition methods," in *ISMIR*, 2007, pp. 375–378.

[6] T. Hedges, P. Roy, and F. Pachet, "Predicting the composer and style of jazz chord progressions," *Journal of New Music Research*, vol. 43, no. 3, pp. 276–290, 2014.

[7] R. Hillewaere, B. Manderick, and D. Conklin, "String quartet classification with monophonic models," in *ISMIR*, 2010, pp. 537–542.

[8] M. A. Kaliakatsos-Papakostas, M. G. Epitropakis, and M. N. Vrahatis, "Weighted Markov chain model for musical composer identification," in *European Conference on the Applications of Evolutionary Computation*. Springer, 2011, pp. 334–343.

[9] Y. W. Liu and E. Selfridge-Field, *Modeling music as Markov chains: Composer identification*. Citeseer.

[10] D. Meredith, "Using point-set compression to classify folk songs," in *Fourth International Workshop on Folk Music Analysis*, 2014, p. 7 pages.

[11] D. Conklin, "Multiple viewpoint systems for music classification," *Journal of New Music Research*, vol. 42, no. 1, pp. 19–26, 2013.

[12] M. A. Kaliakatsos-Papakostas, M. G. Epitropakis, and M. N. Vrahatis, "Musical composer identification through probabilistic and feedforward neural networks," in *European Conference on the Applications of Evolutionary Computation*. Springer, 2010, pp. 411–420.

[13] A. Pienimäki and K. Lemström, "Clustering symbolic music using paradigmatic and surface level analyses," in *ISMIR*, 2004.

[14] T. Collins, R. Laney, A. Willis, and P. H. Garthwaite, "Developing and evaluating computational models of musical style," *AI EDAM*, vol. 30, no. 1, pp. 16–43, 2016.

[15] T. Collins and R. Laney, "Computer–generated stylistic compositions with long–term repetitive and phrasal structure," *Journal of Creative Music Systems*, vol. 1, no. 2, 2017.

[16] M. McVicar, S. Fukayama, and M. Goto, "AutoGuitarTab: Computer-aided composition of rhythm and lead guitar parts in the tablature space," *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, vol. 23, no. 7, pp. 1105–1117, 2015.

[17] ——, "AutoRhythmGuitar: Computer-aided composition for rhythm guitar in the tab space," in *ICMC*, 2014.

[18] ——, "Autoleadguitar: Automatic generation of guitar solo phrases in the tablature space," in *Signal Processing (ICSP), 2014 12th International Conference on*. IEEE, 2014, pp. 599–604.

[19] R. de Valk, T. Weyde, and E. Benetos, "A machine learning approach to voice separation in lute tablature," in *ISMIR*, 2010, pp. 537–542.

[20] S. Ferretti, "Guitar solos as networks," in *Multimedia and Expo (ICME), 2016 IEEE International Conference on*. IEEE, 2016, pp. 1–6.

[21] S. Cherla, H. Purwins, and M. Marchini, "Automatic phrase continuation from guitar and bass guitar melodies," *Computer Music Journal*, vol. 37, no. 3, pp. 68–81, 2013.

[22] T. Collins, "Constructing maskanda," *SAMUS: South African Music Studies*, vol. 26, no. 1, pp. 1–26, 2006.

[23] G. Widmer, "Getting closer to the essence of music: the con espressione manifesto," *ACM Transactions on Intelligent Systems and Technology, Special Issue on Intelligent Music Systems and Applications*, vol. 8, no. 2, p. 13 pages, 2016.

[24] E. Fosler-Lussier, *Markov models and hidden Markov Models: A brief tutorial*. International Computer Science Institute, 1998.

[25] Y. Benjamini and D. Yekutieli, "The control of the false discovery rate in multiple testing under dependency," *The Annals of Statistics*, vol. 29, no. 4, pp. 1165–1188, 2001.

[26] G. A. Spedicato, *markovchain: Easy Handling of Discrete Time Markov Chains*, 03 2017, r package version 0.6.8.1.

[27] B. Sturm, "A simple method to determine if a music information retrieval system is a "horse"," *IEEE Transactions on Multimedia*, vol. 16, no. 6, pp. 1636–1644, 2014.