



HAL
open science

Identification of rhythm guitar sections in symbolic tablatures

David Régnier, Nicolas Martin, Louis Bigo

► **To cite this version:**

David Régnier, Nicolas Martin, Louis Bigo. Identification of rhythm guitar sections in symbolic tablatures. International Society for Music Information Retrieval Conference (ISMIR 2021), 2021, Online, United States. hal-03335822

HAL Id: hal-03335822

<https://hal.science/hal-03335822>

Submitted on 13 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

IDENTIFICATION OF RHYTHM GUITAR SECTIONS IN SYMBOLIC TABLATURES

David Régnier¹

Nicolas Martin²

Louis Bigo¹

¹ Univ. Lille, CNRS, Centrale Lille
UMR 9189 CRISAL, F-59000 Lille, France

² Arobas Music, Lille, France

louis@algomus.fr

ABSTRACT

Sections of guitar parts in pop/rock songs are commonly described by functional terms including for example *rhythm guitar*, *lead guitar*, *solo* or *riff*. At a low level, these terms generally involve textural properties, for example whether the guitar tends to play chords or single notes. At a higher level, they indicate the function the guitar is playing relative to other instruments of the ensemble, for example whether the guitar is accompanying in *background*, or if it is intended to play a part in the *foreground*. Automatic labelling of instrumental function has various potential applications including the creation of consistent datasets dedicated to the training of generative models that focus on a particular function. In this paper, we propose a computational method to identify rhythm guitar sections in symbolic tablatures. We define rhythm guitar as sections that aim at making the listener perceive the chord progression that characterizes the harmony part of the song. A set of 31 high level features is proposed to predict if a bar in a tablature should be labeled as rhythm guitar or not. These features are used by an LSTM classifier which yields to a F_1 score of 0.95 on a dataset of 102 guitar tablatures with manual function annotations. Manual annotations and computed feature vectors are publicly released.

1. INTRODUCTION

1.1 Guitar tablatures

As many multi-stringed instruments, the guitar allows to play a same note in multiple locations on the neck. The location where the note is played, commonly designated by the term *position*, is specified by the combination of a string name and a fret number. For example, the pitch A3 can be played at fret 2 of the G string or at fret 7 of the D string. Guitar tablatures, as illustrated in Figure 1, aim at disambiguating these positions by indicating the string/fret combinations on which notes must be played. The choice

of the positions relates to playability and to some extent to the guitarist playing style [1].

1.2 Functions in guitar tablatures

Similarly to other instruments like the piano, the role of the guitar in a pop/rock ensemble can potentially be associated with different functions over a song. Most of the time, these functions can be gathered within two categories being accompaniment and melody, generally designated by the terms *rhythm guitar* and *lead guitar*. Although not central in this paper and less frequent in the context of a pop/rock ensemble, it is worth noting that the guitar, as the piano, can simultaneously perform accompaniment and melody. The piano will typically split the two functions into left hand and right hand while the guitar will generally use a specific playing technique called *finger picking*.

A more general way to describe the function of the guitar is to estimate if it is thought to be perceived in the *background* or in the *foreground* of the song. Accompaniment parts will generally fit the first category as they often aim at supporting a main musical part like a singing part or an instrumental solo. Although melodic parts are generally thought to be perceived in the foreground, it is not uncommon for a lead guitarist to play an accompanying melody, possibly improvised, during singing sections. Examples of this behavior include the verses of the song *What's Up* (4 Non Blondes) or the bridge of the song *Cryin'* (Aerosmith).

Rhythm guitar sections in the pop/rock repertoire mostly consist in (repetitively) realizing a chord sequence. Figure 1a illustrates two bars of rhythm guitar. In contrast, *lead guitar* appears to be less well-defined as it can be alternately associated with solo parts, as in Figure 1b, *riffs* and *licks*, or hybrid accompanying parts not directly related to the underlying chord sequence. In this work, we focus on the detection of rhythm guitar sections. Rhythm guitar sections are defined as guitar sections that aim at making the listener perceive the chord progression that characterizes the harmony part of the song. In pop/rock style, such chord progressions can often be indicated independently as chord symbols accompanying melodies and lyrics.

Although rhythm guitar looks more easily definable than lead guitar, it is common to find ambiguous guitar



(a) Extract of a rhythm guitar section from *Space Oddity* (David Bowie)

(b) Extract of a solo section from *Another Brick In The Wall* (Pink Floyd)

(c) An ambiguous extract from *Sultans of Swing* (Dire Straits)

Figure 1. Three guitar tablature extracts.

sections standing at the border of what rhythm guitar could be. Figure 1c illustrates this ambiguity with an extract of *Sultans of Swing* (Dire Straits). A rock song can typically begin with a guitar riff played as a foreground part, which is then repeated as a background accompaniment of a vocal verse. One example of this is the famous introducing riff of the song *Highway to Hell* (AC/DC) which switches from foreground to background as the vocal part begins. Ambiguities can also appear with punctual arrangement parts that are generally added during studio recording sessions.

The way ambiguous sections are labelled should be carefully considered if this labelling aims at separating a sub-corpus intended to train a rhythm guitar generative model. On one hand, a strict labelling would reach to a consistent sub-corpus with limited variety. On the other hand, a more flexible labelling would reach to a sparser sub-corpus but richer in variety. This aspect will be further discussed in Section 4.3.

1.3 Applications in MIR

Modeling instrumental function contributes to improve various applications in Music Information Retrieval including computational music analysis and generation. Identifying textural features that contribute to a function improves our knowledge in music theory and our understanding of musical style. Systematic studies bring our attention on unexpected and ambiguous cases which eventually encourage reconsiderations of common definitions.

Automatic function identification can also guide the division of large corpora into function-specific sub-corpora that will facilitate the effective training of machine learning models. For instance, a model trained exclusively on

rhythm guitar sections might be more performant in generating, analyzing, or transcribing such sections. In contrast, studying guitar playing techniques like bends, hammer-on/pull-off, and tapping will benefit from being done on a corpus of guitar solos as they appear predominantly in these sections.

2. RELATED WORK

2.1 Guitar tablature modelling

MIR research on guitar tablatures predominantly relates to automatic fingering, style analysis, and generation.

The automatic fingering task results from the fact that a same note can generally be played at multiple locations on the neck of the guitar. This task therefore consists in estimating a string/fret combination for each note of a score in order to optimize its playability. The fingering problem has been approached with various methods including HMM from audio signal [2] or symbolic scores [3], and visual detection [4].

Guitar tablature automatic analysis includes the detection in audio recordings of specific playing techniques (bends, hammer-on, pull-off, etc.) [5, 6]. Analysis of audio guitar recordings also include automatic transcription of tablatures [7] based on the training of CNNs on guitar recording spectrograms, that tackle both the pitch and fingering estimation. Automatic analysis of symbolic tablatures include guitarist style modeling with Markov models [1] or directed graphs [8], as well as the study of predominant fretboard positions [9].

Guitar tablature generation has been approached with various methods including HMMs to generate guitar arrangements from audio polyphonic music [10], integer programming to generate blues solos [11], and transformer neural networks to generate fingerpicking tablatures [12]. Of particular relevance in the context of this research, guitar tablature generation has also been limited rhythm guitar and lead guitar [13, 14] with probabilistic methods.

2.2 Musical function identification

The complementarity between *rhythm* and *lead* guitar sections in pop/rock tablatures can be generalized to the notion of musical function in musical scores. Identifying whether a section of a part corresponds to background accompaniment or to foreground melody relates to *texture modeling* [15, 16] which has been rarely addressed in symbolic scores so far. In audio recordings however, a number of works has been achieved to detect solo sections [17–19], which can employ similar techniques as vocal activity detection [20]. Solo detection contributes to the task of structure estimation for which a number of research has been done either on symbolic [21] and audio data [22].

Particularly related to this research, guitar playing modes (bass lines, chord comping and solo melody improvisations) can be detected in audio recordings with signal processing features [23] but to the best of our knowledge, there is no research detecting guitar-playing modes from symbolic tablatures.

3. HIGH LEVEL FEATURES

Rhythm guitar is considered in this work as a category of tablature sections that aim at making the listener perceive the chord progression underlying to the song. This section presents a set of 31 features that are designed and evaluated to detect such a behavior.

3.1 Bar-level labels

Although the role of a guitarist in a pop/rock song can strictly be limited to *rhythm* or *lead* guitar, it is common to see guitar tablatures switching between *rhythm* parts and *lead* parts over a same song. A number of bands have a single guitarist who alternates during a song between accompanying *rhythm* parts, and *riff/solo lead* parts. A global labelling of guitar tablatures as *rhythm* or *lead* might therefore lead to approximations and wrong interpretations. In contrast, trying to characterize the role of the guitar at the beat level would require unnecessary complexity as these functional labels tend to span over much larger time frames. In this work, we propose to assign rhythm guitar labels to *bars* of a tablature.

3.2 High level features

The 31 high level features described in this section and summarized in Table 1 are intended to be computed at each bar from raw tablature informations. These informations include note pitches, onsets, durations, string and fret indications, as well as occurrences of some technical playing modes specific to the guitar. Note that some features may derive from combinations of others. For example, the pitch of a note can be deduced from its string and fret value.

3.2.1 Note-related features

Note related features include the number of notes in the bar, as well as the presence of single notes (*i.e.*, not played simultaneously to any other). Pitch-related features include mean/min/max pitch, pitch ambitus and pitch variety (*i.e.*, number of distinct pitches). Pitch interval related features include min/max interval found between 2 successive single notes and interval variety. Finally we added the variety of note durations found in the bar.

3.2.2 Chord-related features

A chord is considered here as a set of at least two notes that are plucked simultaneously. Note that arpeggiated chords are generally notated in guitar tablatures as successive single notes labeled with a *let-ring* indication. Arpeggios are therefore not included in this definition of chords.

Chord related features include the presence of chords, the number of distinct chords and more specifically the number of n -note chords with n in [2..6]. Two additional features indicate whether a triad (either minor or major) or a fifth interval can be formed with the whole set of notes in the bar.

note features		chord features		tab features	
# notes	(7e+2)	chords*	(2e+3)	min fret	(2e+3)
single notes*	(1e+3)	# 2-chords	(1e+1)	max fret	(2e+3)
min pitch	(3e+3)	# 3-chords	(3e+2)	mean fret	(2e+3)
max pitch	(8e+2)	# 4-chords	(5e+2)	min string	(3e+3)
mean pitch	(2e+3)	# 5-chords	(2e+2)	max string	(4e0)
pitch ambitus	(1e+3)	# 6-chords	(9e+1)	mean string	(7e+2)
pitch variety	(2e+3)	chord variety	(9e+2)	<i>l-r(s)</i> *	(1e+2)
min interval	(3e+1)	m/M triad*	(5e+2)	<i>l-r (100%)*</i>	(1e+2)
max interval	(1e-1)	fifth interval*	(1e+2)	<i>w.b(s)</i> *	(6e0)
interval var	(2e+2)			<i>bend(s)</i> *	(2e+3)
duration var	(1e+2)			<i>l-h vibr(s)</i> *	(8e+2)

Table 1. Features describing tablature bars for the rhythm guitar detection task. Binary features are indicated with a *. The importance of each feature in the dataset is indicated by its ANOVA F -value.

3.2.3 Guitar tablature specific features

For each bar, the min/max/mean values of both frets and string are computed. Playing technique features respectively include the presence of at least one *let-ring* (*l-r*), *vibrato*, *whammy bar* (*w.b*) and *bend* indication. A feature indicating whether the whole bar is covered by a *let-ring* indication is added.

4. EXPERIMENTS

The detection of rhythm guitar bars is formulated as a binary classification problem with two classes being *rhythm-guitar* and *other*. Each bar is described by the set of features presented above. A classifier is then trained to predict the label of a bar from its feature values.

4.1 Annotated dataset

For this work, 102 guitar tablatures in the *Guitar Pro* format from the *mySongBook* corpus¹ were analyzed, annotated and checked by two musicians experts in the pop/rock style. Selected tablatures are mostly in the pop/rock style with a few exceptions in swing/jazz. Only tablatures of six strings with standard tuning (E³ A³ D⁴ G⁴ B⁴ E⁵) were included in the annotated dataset. Among the 7487 non-empty bars (60% of the whole dataset), 6051 (82%) were labeled as *RhythmGuitar* (the other 1368 ones were complementarily labeled as *other*). Different functions were identified within this complementary class including solos, licks, riffs and studio arrangements. No finger-style tablatures were included as this playing style generally mixes both accompaniment and lead melody, making its annotation ambiguous.

Raw tablatures are not available due to legal constraints. However, computed features and manual annotations are released² in an open licence.

4.2 Feature analysis

File parsing and feature computation were performed with the *music21* python library [24] using a dedicated

¹ <https://www.mysongbook.com/>

² <http://algomus.fr/data/>

parser [25]. Figure 2 shows the value distribution of a selection of features extracted from bars of both classes in the annotated dataset. To facilitate the comparison of the two classes, the histograms indicate the proportion of feature values in each class rather than the actual number of bars. As expected, rhythm guitar and non-rhythm guitar bars appear to be respectively correlated with the presence of chords (80% of rhythm guitar bars) and the presence of single notes (92% of non-rhythm guitar bars). Non-rhythm guitar can also be distinguished by a lower number of notes and distinct chords. Rhythm guitar bars can finally be distinguished by a lower register that appears in pitch, fret, and string related features. An ANOVA Fischer test is performed for each feature as an indication of its correlation with the two classes. The results are displayed on Table 1.

4.3 Rhythm guitar prediction

4.3.1 Evaluation measure

The choice of an evaluation measure of the performance of classifier that predicts whether a guitar tab bar is rhythm guitar or not varies depending on the way the result of the classification is intended to be used.

On one hand, maximizing *precision* penalizes false positives and potentially leads to a consistent rhythm guitar sub-corpus although possibly small and uniform. Such a corpus would facilitate the training of a model that is expected to produce *typical*, but not necessary surprising, rhythm guitar tablatures. On the other hand, maximizing *recall* penalizes false negatives and potentially leads to a larger sub-corpus with more diversity although more sparse and including more debatable rhythm guitar examples. Such a corpus would be appropriate for the training of a model that aims at generating creative rhythm guitar tablatures, at the expense of outputs that possibly diverge from the common definition of rhythm guitar. Note that for a classifier that outputs a probability (like neural networks) moving the decision threshold, that is generally set by default to 0.5, could also be a way to balance between consistency and variety.

From an analysis point of view, improving our comprehension of what makes a rhythm guitar bar requires to take into account both false negatives and false positives, which could be achieved by using accuracy. As the dataset is unbalanced, we propose to evaluate the F_1 score which is defined by the harmonic mean of precision and recall.

4.3.2 Leave-one-piece-out evaluation

Training a machine learning model is often performed by splitting the dataset into a training set and a validation set. As bars can highly repeat, in particular in rhythm guitar sections, all bars belonging to the same piece should belong to the same subset to avoid overfitting. The small size of our dataset lets us adopt a *leave-one-piece-out* validation process: given the dataset of n pieces, the model is trained on $n - 1$ pieces and then evaluated on the remaining one. The process is repeated for the n pieces and the evaluation is therefore performed on the whole set of pieces of the dataset. The *leave-one-piece-out* method allows to

	<i>r.g</i> precision	<i>r.g</i> recall	F_1 score
chords/single notes presence	0.86	0.88	0.87
note + chord features	0.95	0.94	0.94
tab features	0.95	0.93	0.94
all features	0.96	0.94	0.95

Table 2. Precision, recall and F_1 score obtained for the detection of rhythm guitar (*r.g*) with a LSTM trained on different set of features.

maximize the quantity of training datas and evaluate the model on the whole dataset.

Different classifiers were tested including logistic regression, SVM, decision tree and random forest thanks to the scikit-learn framework [26]. A Long Short-Term Memory model (LSTM) implemented with the Keras framework [27] happened to provide the best results. The LSTM has 2 hidden layers of 75 and 10 units. An early stopping process was used to identify the optimal number of 12 epochs. A batch size of 32 was used and bars were presented to the model by subsequences of 5. It is not surprising to see a recurrent model outperforming standard classifiers given that bars of the same label are likely to occur consecutively in the piece as outlined in section 5. The code is publicly provided³.

5. RESULTS AND DISCUSSIONS

Different sets of features among those presented in Section 3.2 were tested to evaluate the model. We first consider a baseline model that only looks at the presence of chords and single notes in each bar. We then evaluate score based features (first two columns of Table 1). We then evaluate tablature based features only (third column of Table 1). Finally, we evaluate a model taking into account the whole set of features. In addition to F_1 score, Table 2 displays the *precision* and the *recall* on rhythm guitar label predictions.

The LSTM baseline model achieves a F_1 score of 0.87. The LSTM model combining the whole set of features reaches a F_1 score of 0.95, which outperforms other tested models including logistic regression ($F_1=0.93$), decision tree ($F_1=0.91$) and random forest ($F_1=0.93$). Although disjoint, the score feature set and the tab feature set interestingly achieve similar performance. This can partly be explained by the fact that pitch informations in score features can be derived from string+fret combinations in tab features. It is interesting to observe that string/fret and playing technics indications seem to counterbalance the absence of chord related informations, although presumably crucial for rhythm guitar detection. It also worth to note that both these two feature sets almost yield to the score obtained with the whole set of features which means that none of them much improves the other. In the following, we present wrong predictions obtained with the model trained with the whole set of features.

³ <https://gitlab.com/lbigo/rhythm-guitar-detection>

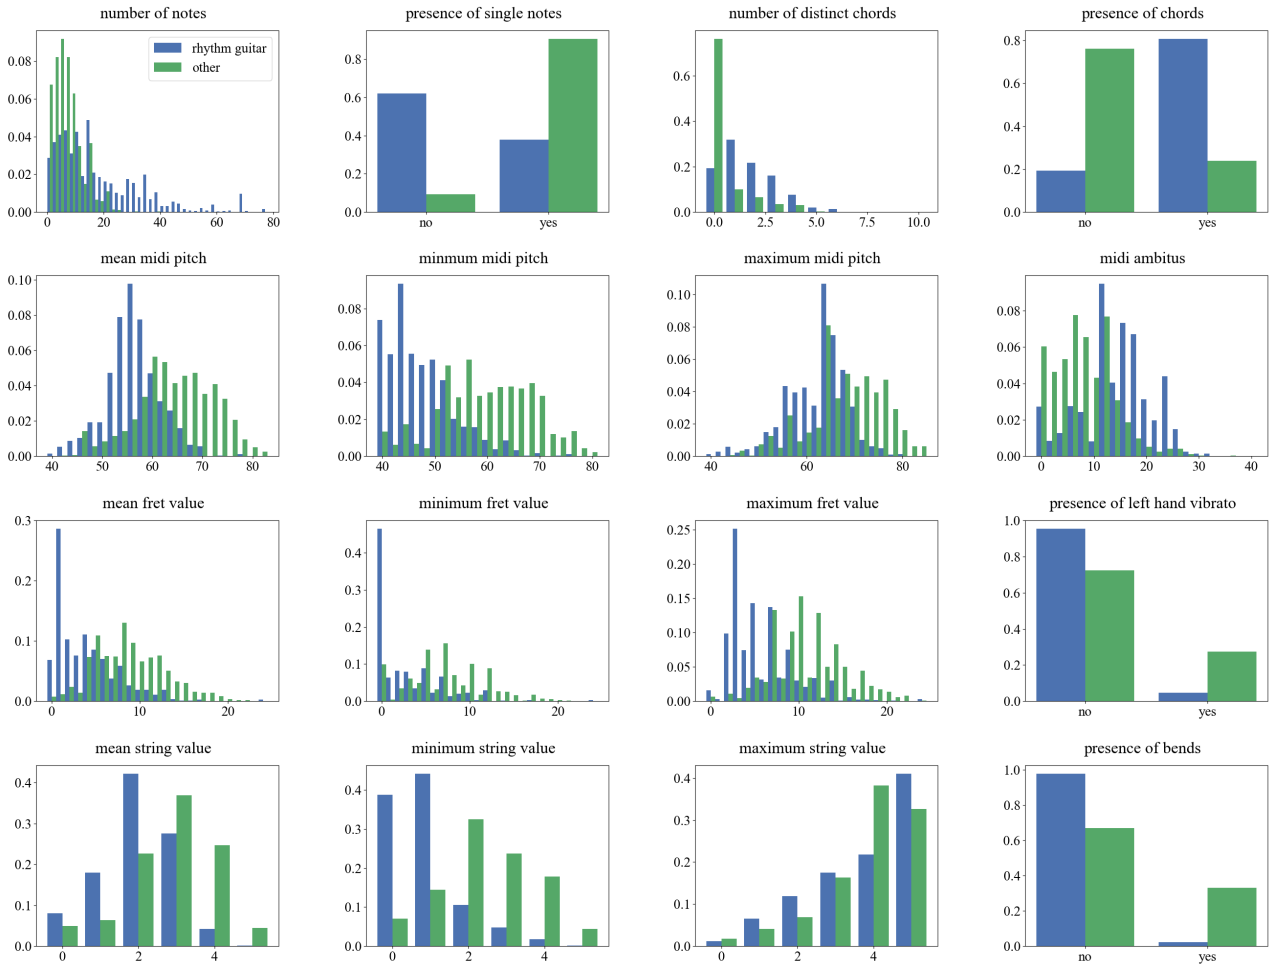


Figure 2. Disitribution of some features on bars annotated with label *Rhythm guitar* (blue) and *other* (green).

Figure 3 displays a comparison between reference annotations (top line) and predictions (bottom line) for a selection of tablatures of the corpus. Although the model succeeds in identifying large scale sections, it can still predict unlikely short sections, sometimes for one unique bar. For example, the model wrongly predicts unlikely short rhythm guitar sections in the song *Sultans of Swing* (Dire Straits). Similarly, it wrongly predicts unlikely short non-rhythm guitar sections in the songs *Stairway To Heaven* (Led Zeppelin) and *You Only Live Once* (The Strokes) as discussed below.

Figure 4 illustrates three examples of false negatives, *i.e.* rhythm guitar bars predicted as being non-rhythm guitar bars. Examples 4a and 4b are extracted from songs *You Only Live Once* (The Strokes) and *Stairway To Heaven* (Led Zeppelin). In these two examples, only the middle bar is wrongly estimated as non-rhythm guitar. Both these bars have the particularity to be the final bar of a musical phrase, leading to a new phrase beginning on the next bar. In these cases, the rhythm guitar punctually plays a short melodic lick often referred as a *fill*, which is not identified as rhythm guitar by the model. This kind of wrong predictions could probably be avoided by improving the faculty of the model to capture the tendency of adjacent bars to

have the same label and avoid the prediction of isolated labels, for example using a bidirectional LSTM. Example 4c is extracted from the song *When The Sun Goes Down* (Arctic Monkeys). In this example, the guitar starts to play bass single notes and produces a melodic line which is wrongly estimated by the model as non-rhythm guitar. This behavior could arguably be qualified as being at the edge of the common definition of rhythm guitar and it would be difficult to avoid this kind of wrong predictions without looking at the other tracks of the song (in particular the singing part), which is out of the scope of this work.

Figure 5 illustrates three examples of false positives, *i.e.* non-rhythm guitar bars predicted as rhythm guitar bars. Example 5a illustrates an extract of a solo part of the song *Hotel California* (Eagles) where the guitar repetitively plays arpeggios of the underlying chord sequence. Although the played notes belong to a rather high register, the model is probably misled by the repetitiveness, low variety and the presence of perfect triad as these features are often correlated with rhythm guitar sections. Example 5b consists in a short interlude between a solo section and the bridge of the song *La Grange* (ZZ Top). In this case, the function of the guitar seems to consist in doing a transition between two sections and could hardly be unam-



Figure 3. Comparison of manual annotations (top lines) and predictions (bottom lines) of a some tablatures of the dataset. Sections labelled as rhythm guitar are displayed in blue. Other sections are displayed in green. Empty bars are left in gray.

	<i>r.g</i> measures	<i>r.g</i> sections	mean <i>r.g</i> section length	isolated <i>r.g</i> measures
reference	6051	101	77	0
prediction	5923	223	34	44

Table 3. Comparison of consecutiveness of annotated and predicted rhythm guitar (*r.g*) bars.

biguously described as rhythm guitar or not. Example 5c is extracted from a solo section of the song *Minor Swing* (Django Reinhardt). The model is clearly misled by the sudden occurrence of chords here. As it is often the case gypsy jazz, the guitar punctually includes series of chords within a solo, that do not necessarily precisely feet the underlying chord sequence. This behavior typically lasts a few bars before the guitar goes back to melody.

Table 3 illustrates the difficulty of the model to reconstruct continuous rhythm guitar sections. Although the proportion of rhythm guitar bars predicted by the model is close to the one of the reference, these bars are grouped in smaller and more numerous sections. The model particularly tends to detect isolated rhythm guitar bars although the reference annotation do not include any of them.



(a) *You Only Live Once* (The Strokes)

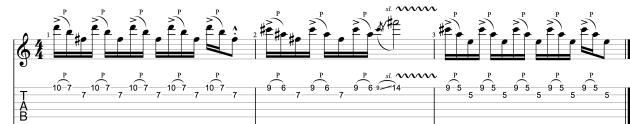


(b) *Stairway To Heaven* (Led Zeppelin)

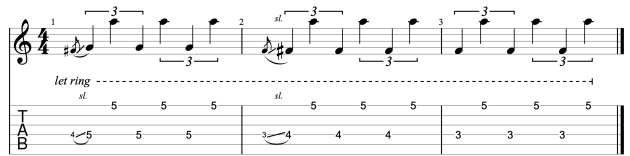


(c) *When The Sun Goes Down* (Arctic Monkeys)

Figure 4. Examples of false negatives. The second bar is wrongly predicted as non-rhythm guitar on each example.



(a) *Hotel California* (Eagles)



(b) *La Grange* (ZZ Top)



(c) *Minor Swing* (Django Reinhardt)

Figure 5. Examples of false positives. Second and third bars are wrongly predicted as rhythm guitar.

6. CONCLUSIONS

This study improved our understanding of which features contribute to a rhythm guitar section. We believe that this approach can be used to separate a corpus of pop/rock guitar tablatures into consistent sub-corpora dedicated to tablature generation limited to a specific function.

The method presented here could benefit from several improvements. A finer tuning of the LSTM, or the use of a bidirectional LSTM, would probably better capture the tendency of adjacent bars to have the same label and therefore to limit isolated predictions which appear to be very unlikely across the corpus. Futur works also include adding features that look at more structural aspects of the song like bar location and activity of other tacks, in particular singing tracks as rhythm guitar if often intended to accompany singing.

Acknowledgements. The authors are grateful to anonymous reviewers and all the Algomus team for fruitful comments on the manuscript. This work is partially funded by French CPER MAuVE (ERDF, Région Hauts-de-France).

7. REFERENCES

- [1] O. Das, B. Kaneshiro, and T. Collins, “Analyzing and classifying guitarists from rock guitar solo tablature,” in *Proceedings of the Sound and Music Computing Conference, Limassol, Chypre*, 2018.
- [2] A. M. Barbancho, A. Klapuri, L. J. Tardón, and I. Barbancho, “Automatic transcription of guitar chords and fingering from audio,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no. 3, pp. 915–921, 2011.
- [3] G. Hori and S. Sagayama, “Minimax viterbi algorithm for hmm-based guitar fingering decision.” in *ISMIR*, 2016, pp. 448–453.
- [4] A.-M. Burns and M. M. Wanderley, “Visual methods for the retrieval of guitarist fingering,” in *Proceedings of the 2006 conference on New Interfaces for Musical Expression*. Citeseer, 2006, pp. 196–199.
- [5] L. Reboursière, O. Lähdeoja, T. Drugman, S. Dupont, C. Picard-Limpens, and N. Riche, “Left and right-hand guitar playing techniques detection.” in *NIME*, 2012.
- [6] Y.-P. Chen, L. Su, Y.-H. Yang *et al.*, “Electric guitar playing technique detection in real-world recording based on f0 sequence pattern recognition.” in *ISMIR*, 2015, pp. 708–714.
- [7] A. Wiggins and Y. Kim, “Guitar tablature estimation with a convolutional neural network.” in *ISMIR*, 2019, pp. 284–291.
- [8] S. Ferretti, “Guitar solos as networks,” in *2016 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, 2016, pp. 1–6.
- [9] J. Cournut, M. Giraud, L. Bigo, N. Martin, and D. Régnier, “What are the most used guitar positions?” in *International Conference on Digital Libraries for Musicology (DLfM 2021)*, Online, United Kingdom, 2021. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-03279863>
- [10] S. Ariga, S. Fukayama, and M. Goto, “Song2guitar: A difficulty-aware arrangement system for generating guitar solo covers from polyphonic audio of popular music.” in *ISMIR*, 2017, pp. 568–574.
- [11] N. d. S. Cunha, A. Subramanian, and D. Herremans, “Generating guitar solos by integer programming,” *Journal of the Operational Research Society*, vol. 69, no. 6, pp. 971–985, 2018.
- [12] Y.-H. Chen, Y.-H. Huang, W.-Y. Hsiao, and Y.-H. Yang, “Automatic composition of guitar tabs by transformers and groove modeling,” *arXiv preprint arXiv:2008.01431*, 2020.
- [13] M. McVicar, S. Fukayama, and M. Goto, “Autorhythmguitar: Computer-aided composition for rhythm guitar in the tab space,” in *ICMC*, 2014.
- [14] ———, “Autoleadguitar: Automatic generation of guitar solo phrases in the tablature space,” in *2014 12th International Conference on Signal Processing (ICSP)*. IEEE, 2014, pp. 599–604.
- [15] B. Duane, “Texture in eighteenth-and early nineteenth-century string-quartet expositions,” Ph.D. dissertation, Northwestern University, 2012.
- [16] M. Giraud, F. Levé, F. Mercier, M. Rigaudière, and D. Thorez, “Towards modeling texture in symbolic data,” in *ISMIR*, 2014, pp. 59–64.
- [17] K. A. Pati and A. Lerch, “A dataset and method for guitar solo detection in rock music,” in *Audio Engineering Society Conference: 2017 AES International Conference on Semantic Audio*. Audio Engineering Society, 2017.
- [18] G. Peterschmitt, E. Gomez, and P. Herrera, “Pitch-based solo location,” in *Proc. of MOSART Workshop on Current Research Directions in Computer Music*, 2001.
- [19] F. Fuhrmann, P. Herrera, and X. Serra, “Detecting solo phrases in music using spectral and pitch-related descriptors,” *Journal of New Music Research*, vol. 38, no. 4, pp. 343–356, 2009.
- [20] M. Mauch, H. Fujihara, K. Yoshii, and M. Goto, “Timbre and melody features for the recognition of vocal activity and instrumental solos in polyphonic music.” in *ISMIR*, 2011, pp. 233–238.
- [21] P. Allegraud, L. Bigo, L. Feisthauer, M. Giraud, R. Groult, E. Leguy, and F. Levé, “Learning sonata form structure on mozart’s string quartets,” *Transactions of the International Society for Music Information Retrieval (TISMIR)*, vol. 2, no. 1, pp. 82–96, 2019.
- [22] K. Ullrich, J. Schlüter, and T. Grill, “Boundary detection in music structure analysis using convolutional neural networks.” in *ISMIR*, 2014, pp. 417–422.
- [23] R. Foulon, P. Roy, and F. Pachet, “Automatic classification of guitar playing modes,” in *International Symposium on Computer Music Multidisciplinary Research*. Springer, 2013, pp. 58–71.
- [24] M. S. Cuthbert and C. Ariza, “music21: A toolkit for computer-aided musicology and symbolic music data,” 2010.

- [25] J. Cournut, L. Bigo, M. Giraud, and N. Martin, “Encodages de tablatures pour l’analyse de musique pour guitare,” in *Journées d’Informatique Musicale (JIM 2020)*, Strasbourg, France, 2020. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-02934382>
- [26] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [27] F. Chollet *et al.* (2015) Keras. [Online]. Available: <https://github.com/fchollet/keras>