

A DATA-DRIVEN METHODOLOGY FOR CONSIDERING FEASIBILITY AND PAIRWISE LIKELIHOOD IN DEEP LEARNING BASED GUITAR TABLATURE TRANSCRIPTION SYSTEMS

Frank Cwitkowitz[†]
University of Rochester
fcwitkow@ur.rochester.edu

Jonathan Driedger
Chordify
jonathan@chordify.net

Zhiyao Duan
University of Rochester
zhiyao.duan@rochester.edu

ABSTRACT

Guitar tablature transcription is an important but understudied problem within the field of music information retrieval. Traditional signal processing approaches offer only limited performance on the task, and there is little acoustic data with transcription labels for training machine learning models. However, guitar transcription labels alone are more widely available in the form of tablature, which is commonly shared among guitarists online. In this work, a collection of symbolic tablature is leveraged to estimate the pairwise likelihood of notes on the guitar. The output layer of a baseline tablature transcription model is reformulated, such that an inhibition loss can be incorporated to discourage the co-activation of unlikely note pairs. This naturally enforces playability constraints for guitar, and yields tablature which is more consistent with the symbolic data used to estimate pairwise likelihoods. With this methodology, we show that symbolic tablature can be used to shape the distribution of a tablature transcription model’s predictions, even when little acoustic data is available.

1. INTRODUCTION

Automatic Music Transcription (AMT) is a well-known task within the Music Information Retrieval (MIR) community dealing with the estimation of note content within a music signal [1]. Guitar tablature transcription refers to the specific problem of identifying all of the notes within a solo guitar recording and assigning the string that was used to play them. The task represents the determination of not only what was played, but also how it was played on the instrument. This information is necessary to realize tablature, a type of prescriptive notation for stringed instruments where fret numbers are superimposed atop lines representing each string. The fret numbers correspond to the notes that are to be played for a specific piece.

The guitar is a very popular musical instrument with users spanning all skill levels. The value of knowing how a piece

was played on guitar is immeasurable for the vast community of guitarists learning to play the instrument. Guitar has a relatively low barrier to entry w.r.t. music theory knowledge, and many guitarists use guitar tablature instead of standard staff notation. Even more experienced players often prefer tablature for storing and communicating guitar-specific music ideas due to the intuitiveness and simplicity. Tablature is also widely shared across the internet through primarily user-curated websites such as *Ultimate-Guitar*¹.

Despite the popularity of guitar, the instrument has received considerably less attention when it comes to music transcription. The main obstacles stem from a lack of audio recordings with transcription labels, or acoustic data, capturing the exceeding variability of the instrument. The guitar has many expressive dimensions such as the plucking style, plucking location, the use of embellishments², etc. These and many other factors can affect the audio. Standard guitars do not have a digital interface, and expensive manual processing is required to obtain qualitative note labels. Without large datasets capturing the breadth of the instrument, it is very difficult to train reliable models and avoid over-fitting.

Since a standard guitar has six independent strings, it is a polyphonic instrument. This means that guitar transcription carries all of the intrinsic challenges of polyphonic note transcription. One further challenge is the estimation of the string on which each note was played. The pitch ranges of adjacent strings have a significant overlap, with their lowest pitches being only 4-5 semitones apart. The timbral cues of different strings are very subtle, and it is difficult to learn these without a lot of data.

There have been several attempts to realize systems which can transcribe solo guitar audio into tablature. Often, the approaches consist of two-stage systems, which first estimate pitch salience and then map pitch estimates to the guitar. Many works employ a basic signal processing pipeline, e.g. [2–4], whereby some form of spectral analysis or peak-picking is carried out. Some approaches attempt to estimate the string of detected pitches based off of inharmonicity measurements [5], which vary across strings. Others employ graph-search algorithms to find the optimal path through a list of fingerings for the observed pitches [6–8]. Often, to determine if a fingering or a transition is optimal or even feasible, these approaches have relied on rule-based procedures with hand-crafted weightings.

[†] Main work completed as a research intern at Chordify.

Copyright: © 2022 Frank Cwitkowitz et al. This is an open-access article distributed under the terms of the [Creative Commons Attribution 3.0 Unported License](https://creativecommons.org/licenses/by/3.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

¹ <https://www.ultimate-guitar.com/>

² Examples include bends, slides, hammer-ons/offs, vibrato, etc.

More recently, machine learning has become a popular strategy for guitar tablature transcription. Several works train Hidden Markov Models (HMMs) to model the transition between fingerings and chords [9, 10]. Other works perform classification to estimate, among other parameters, the string associated with each note using a collection of hand-crafted features [4], or the use of a Deep Belief Network (DBN) to produce pitch estimates [11]. Bayesian classification has also been proposed [12–14] to estimate the string and fret of notes.

Convolutional Neural Network (CNN) based approaches have been proposed [15, 16] to perform the task of pitch estimation and tablature arrangement jointly. These models benefit from being able to learn features directly from acoustic data, and tend to generalize much more effectively to real-world data. Our main contributions stem from a simple observation: the output layer formulation of these CNNs is prone to falsely producing tablature with duplicated pitches. This is because the output layer formulation represents six independent classification problems, *i.e.*, one per string. While the output neurons of each softmax group implicitly share information from previous layers of the network, they do not explicitly communicate when determining which class to choose. In contrast, the fretting of one string is highly correlated to other string frettings. This is related to what pitch intervals are likely to be played at different locations as well as bio-mechanical feasibility.

In order to incorporate the knowledge of likely fingerings and feasibility, we propose a new output layer formulation for guitar tablature transcription models³. In the new formulation, a novel inhibition objective is applied during training to discourage the concurrent activation of unlikely or infeasible fingerings. The inhibition weights are derived from the likelihood of co-occurrence for each pair of notes on the guitar, which is estimated using DadaGP [17], a large dataset of guitar tablature. The proposed output layer formulation essentially learns a language model without requiring acoustic data. We directly compare the new formulation to the previous formulation, and show that its predictions more closely match the distribution of DadaGP.

2. PROPOSED METHOD

Given a collection of symbolic tablature data, the pairwise likelihood of each string and fret (S/F) combination can be estimated. The new output layer formulation is amenable to training with the pairwise likelihoods through a pairwise inhibition loss. As a result, the predictions of the new output layer more closely match the distribution of the tablature within the collection. Assuming a preponderance of the tablature is playable, the generated tablature will naturally be more feasible to play. In the following sections, we introduce the baseline CNN architecture used in this work, and discuss the new output layer formulation in more detail, the process of estimating the pairwise likelihood of each S/F combination, and the proposed inhibition loss.

³ All code is available at <https://github.com/cwitkowitz/guitar-transcription-with-inhibition>.

2.1 Baseline Model

We employ TabCNN [16] as our baseline model for guitar tablature transcription, leaving most of the original design choices largely unchanged. TabCNN is a simple CNN which processes Constant-Q Transform (CQT) frames and produces sets of fret class predictions. It was designed to be compatible with real-time processing, so it is relatively lightweight and operates on multiple CQT frames in order to make one set of predictions. The model comprised three 2D convolutional layers with ReLU activations, followed by a max pooling layer, a fully-connected layer with ReLU activation, and finally the output layer discussed in the following section. During training, dropout is applied after the max pooling layer and directly before the output layer. The network is trained using AdaDelta optimizer with an initial learning rate of 1.0. We implemented the model from scratch in PyTorch, using all of the same hyperparameters as in the original paper for feature extraction and the model architecture. We also insert a uni-directional long short-term memory (LSTM) [18] layer before the output layer. This is a simple modification which results in a relatively significant improvement (see Sec. 4), without disrupting the real-time processing capacity of TabCNN.

2.2 Output Layer Formulation

The output layer of TabCNN is a fully-connected layer with one softmax activation for each string. The output neurons represent all combinations of string $s \in \{1, \dots, 6\}$ and fret class $f \in \{-1, 0, 1, \dots, F\}$, where $f = -1$ represents a class for silence, $f = 0$ represents the open string, and F is the total number of frets supported. In the standard model, the softmax operations are applied independently to the neurons associated with each string. This results in probability activations $z_{s,f,n}$ for each frame n where $z_{s,f,n} > 0 \forall s, f, n$ and $\sum_{f=-1}^F z_{s,f,n} = 1, \forall s, n$. The loss for a track is then computed by summing the categorical cross entropy for each string group and averaging across the N total frames. This can be expressed as

$$L_{CCE} = -\frac{1}{N} \sum_{n=1}^N \sum_{s=1}^6 \log(z_{s,f',n}), \quad (1)$$

where $z_{s,f',n}$ is the activation corresponding to the ground truth fret f' for string s during frame n . Inference then consists of choosing the frets with the highest activation, resulting in six predictions for each frame. Using a binary representation the predictions can be written as

$$y_{s,f,n} = \mathcal{I}(\operatorname{argmax}_f \{z_{s,f,n}\} = f), \quad (2)$$

where $\mathcal{I}(\cdot)$ is the indicator function. We refer to this formulation as the 6D Softmax formulation. One advantage of this formulation is that the model is not capable of generating invalid predictions, *i.e.*, only one note can be chosen at maximum for each string. While this property is highly desirable, the six softmax activations act independent of one another, causing the model to treat transcription as six somewhat independent classification tasks. In practice, the ground truth S/F combinations making up a fingering arrangement are highly correlated at all times. This

is due to both the physical limitations of what a human hand can play as well as musical motivations that make it unlikely to play certain pitches at the same time.

In contrast, we propose to formulate the output layer as representing a binary classification problem for each fret of each string. In this way, the likelihood of each S/F combination being active is independently computed using a sigmoid activation. As such, we refer to the new formulation as the Logistic formulation. This formulation allows us to expand the loss function to include an inhibition objective to discourage the co-activation of certain pairs of S/F combinations. This additional objective would otherwise conflict with the 6D Softmax formulation, due to normalization of the activations in the softmax function.

The new loss is computed by summing the binary cross entropy across each fret class for each string and again averaging across all frames. For convenience, let us now fold string and fret into a single variable $c \in \{1, \dots, C\}$ representing each distinct S/F combination, where $C = 6 \times (F + 2)$. The loss can then be expressed as

$$L_{BCE} = -\frac{1}{N} \sum_{n=1}^N \sum_{c=1}^C t_{c,n} \log(z_{c,n}) + (1 - t_{c,n}) \log(1 - z_{c,n}), \quad (3)$$

where $t_{c,n}$ is a binary number indicating whether there is a positive class label in the ground-truth at the corresponding S/F combination c at frame n . This new output layer formulation is similar to that of standard piano transcription models [19]. However, we still use Equation (2) to obtain the final predictions, rather than considering each activation above a certain threshold a positive prediction. This is because, ultimately, we can still only choose one fret class per string. Note that the inhibition objective is applied during training and is unaffected by inference.

2.3 Estimating Pairwise Likelihood

Certain S/F combinations on the guitar have a very low chance of being played at the same time. These can include S/F combinations located on the same string, S/F combinations with the same pitch, S/F combinations which are far apart, or simply musically uncommon S/F combinations. In order to incorporate these considerations into the Logistic formulation, we estimate the likelihood of co-occurrence for all S/F combination pairs, to inform a novel training objective for pairwise S/F combination inhibition.

The pairwise likelihood of two S/F combinations c_i and c_j can be estimated using an arbitrary collection of symbolic tablature data. Here we define symbolic tablature data as one-hot encoded annotations for each string of the guitar at the frame-level. Given the symbolic tablature for a single track, we compute the intersection over union (*IoU*) of frame-level occurrences for all pairs of S/F combinations that co-occur in at least one frame. Mathematically, the intersection of a pair is defined as the number of frames where both combinations occur concurrently:

$$inter(i, j) = \sum_{n=1}^N t_{c_i, n} \wedge t_{c_j, n}. \quad (4)$$

The union is defined as the number of frames where either of the combinations occur.

$$union(i, j) = \sum_{n=1}^N t_{c_i, n} \vee t_{c_j, n}. \quad (5)$$

Let $\mathcal{T}'(i, j)$ be the set of tracks where c_i and c_j , independently, each occur in at least one frame. Then, the *IoU* of the pair is averaged across these valid tracks:

$$IoU(i, j) = \frac{1}{|\mathcal{T}'(i, j)|} \sum_{t \in \mathcal{T}'(i, j)} \frac{inter(i, j)_t}{union(i, j)_t}, \quad (6)$$

where $|\mathcal{T}'(i, j)|$ is the cardinality of $\mathcal{T}'(i, j)$. Note that this is only valid for pairs where $|\mathcal{T}'(i, j)| > 0$. All pairs where $|\mathcal{T}'(i, j)| = 0$ receive $IoU(i, j) = 0$.

The final pairwise likelihoods are stored in a symmetric matrix, ordered by string and fret on both axes. As a result of Equation (6), the likelihood of a pair co-occurring with itself is always 1, and the likelihood of pairs within the same string co-occurring is always 0. This is convenient for the next step, as self-occurrence will never be inhibited, whereas same-string-occurrence will be maximally inhibited. The pairwise likelihood for the rest of the S/F combinations with differing strings will fall somewhere in between 0 and 1 (inclusive). See Fig. 1 for an example of a pairwise likelihood matrix.

2.4 Inhibition Loss

In order to apply the estimated pairwise likelihoods to the problem of guitar tablature transcription, we introduce a new loss term for inhibiting the co-activation of unlikely pairs. We refer to this as the inhibition loss L_{inh} . The inhibition loss requires that each pair of S/F combinations receive an inhibition weight $w(c_i, c_j)$ between 0 and 1, indicating how much to penalize the model for producing high activations for the combinations in the pair in a single frame. The inhibition weights here are chosen to be the complement of the pairwise likelihood (*IoU*) estimated in the previous step:

$$w(c_i, c_j) = (1 - IoU(i, j))^b, \quad (7)$$

where b is a parameter which boosts the effective pairwise likelihood by pushing the mass of the inhibition weights closer to 0. Many S/F combinations which tend to co-occur do not necessarily have a high likelihood of occurring together, relative to the amount of times they occur separately. This can make the estimated pairwise likelihood for many combinations small, leading to a high inhibition weight. Since we do not wish to discourage the activations of pairs which commonly co-occur, it can be helpful to boost the pairwise likelihood in this way. Empirically, we found that $b = 2^7$ produced weights with nice contrast between common pairs and uncommon pairs. Boosting can also be thought of as computing the joint probability of not observing c_i and c_j across b total frames, ≈ 3 seconds here. Given the inhibition weights, the inhibition loss for a sequence of N frame is computed as

$$L_{inh} = \frac{1}{2N} \sum_{n=1}^N \sum_{i=1}^C \sum_{j=1}^C z_{c_i, n} z_{c_j, n} w(c_i, c_j). \quad (8)$$




Figure 1. Visualization of estimated pairwise likelihood for all S/F combinations, ordered row- and column-wise by string s and fret f , computed using DadaGP [17] with $b = 1$ (left) and $b = 2^7$ (right). Grid lines are overlaid atop the string boundaries. The likelihoods range from 0 (dark purple) to 1 (bright yellow), where 1 represents pairs which always co-occur (e.g., identical pairs on diagonal) and 0 represents pairs which never co-occur (e.g., non-identical same-string pairs).

Here, the product for every combination of activations produced by the model is taken and scaled by the appropriate inhibition weight, and the result is subsequently summed over all combinations. Since both permutations of each S/F combination are included in the summation, we divide by two to remove redundancy. The total loss then becomes

$$L_{total} = L_{BCE} + \lambda L_{inh}, \quad (9)$$

where λ is a scaling term for balancing the two terms.

3. EXPERIMENTAL SETUP

In order to evaluate the efficacy of the proposed output layer formulation, we use it within TabCNN [16]. We compare it to the 6D Softmax layer, and experiment with several variations to study the effect of inhibition.

3.1 Datasets

We use two datasets for our experiments. GuitarSet [20] is used to train, validate, and test our models, and DadaGP [17] is used to compute weights for the inhibition loss applied during training and employed as an evaluation metric. We briefly introduce the datasets in the following sections.

3.1.1 GuitarSet

GuitarSet [20] is a guitar transcription dataset comprising roughly three hours of acoustic guitar audio. It contains various types of annotations, including pitch and note annotations with string labels. The provided labels were obtained by employing monophonic pitch tracking on the independent audio of each string, recorded with a hexaphonic pickup mounted to the guitar. The dataset features

six guitarists playing two unique interpretations over 30 different chord progressions from various keys, resulting in 360 distinct tracks. Each track is approximately 1312.7 frames or 30.5 seconds on average. GuitarSet is used within a six-fold cross-validation schema, with the dataset splits representing the tracks produced by each respective guitarist. We limit training to only four splits during each fold, holding out one split for validation and one split for testing.

3.1.2 DadaGP

DadaGP [17] is a large collection of symbolic tablature encoded using the proprietary *GuitarPro* file format. The dataset features many popular songs from a variety of artists and spanning many musical styles, with a bias toward rock and metal music. The *GuitarPro* file format can store the transcription of multiple musical voices as tracks in a single file. Tracks corresponding to guitars have note labels which are always associated with an S/F combination.

We process all tracks corresponding to guitars in standard tuning within the *GuitarPro* files, ignoring duplicate⁴ files. Note that many files include more than one guitar track. Using the Python package *PyGuitarPro* [21], we carry out a series of steps⁵ where we assign an onset and offset to each note in the track to obtain tablatures in *JAMS* format [22]. Ultimately, we end up with 33967 pieces of symbolic tablature, most of which comes from full-length

⁴ A duplicate is defined as having a preexisting file name or the “copy” tag in the file name. In cases where there are duplicates with alternate file extensions, we keep the one with the more recent *GuitarPro* version.

⁵ Due to the many complexities of *GuitarPro*, we refer interested readers to the code for more information about these steps. It is also worth mentioning here that we treat slides and hammer-ons/off as two separate notes, and that for bends we only retain the original note.

songs. We use these to estimate pairwise likelihoods.

3.2 Metrics

We compute all of the same metrics⁶ as in [16]. These include precision, recall, and f-measure for all frames of tablature and the equivalent string-agnostic multipitch, as well as the tablature disambiguation rate (TDR), which indicates how well the model maps pitches to their respective strings. We also report the average inhibition losses L_{inh} and L_{inh}^+ using the standard and boosted inhibition weights, respectively. The inhibition losses serve as a proxy for how well the predictions of the model match the distribution of the symbolic tablature used to estimate the pairwise likelihoods. The lower the losses, the smaller the presence of inhibitory pairs in the predictions. Since the pairwise likelihoods were estimated using a large collection of real-world data, the inhibition losses also serve as a proxy for the feasibility of the predicted tablature. Note that the inhibition losses are calculated on the final predictions of a model. This is in contrast to the inhibition loss applied during training, which is computed on the activations after the sigmoid operation.

Additionally, we note that the inhibition weights tend to be very high for pairs on different strings corresponding to the same exact pitch $p \in \{1, \dots, P\}$, where P is the total number of unique pitches. A duplicate pitch error is a common mistake made by the network, since the acoustic profile of the same pitch on different strings is very similar. We report the average number of duplicate pitch errors per track, where the duplicate pitch error count $E_{d,p}$ for a track with N frames is expressed as

$$E_{d,p} = \sum_{n=1}^N \sum_{p=1}^P \max(0, m_{p,n}^{(y)} - m_{p,n}^{(t)}), \quad (10)$$

where $m_{p,n}^{(y)} \in \{0, \dots, 5\}$ and $m_{p,n}^{(t)} \in \{0, \dots, 5\}$ are the number of string-wise model predictions and ground-truth targets, respectively, at frame n corresponding to pitch p . Since it is extremely rare that a guitarist duplicates the same pitch in practice, this error count is a more explicit way of telling if the inhibition objective is effective. We also report the average number of false alarm errors,

$$E_{f.a.} = \sum_{n=1}^N \sum_{s=1}^6 \sum_{f=0}^F y_{s,f,n} \wedge \neg t_{s,f,n}, \quad (11)$$

indicating more generally all predictions made by the model which do not occur in the ground-truth. This allows us to compare the number of duplicate pitch errors to the total number of false alarm errors in each experiment. Note that here we do not consider incorrect silence predictions ($f = -1$) to be false alarm errors.

3.3 Experiments

We conduct a series of experiments to observe how the proposed output layer formulation compares to the base-

⁶ The original metrics were computed across all frames of all tracks in a track-agnostic manner, whereas we compute the metrics across all frames of each track and average results across all tracks. This is done to weight the influence of every track in the dataset equally.

line 6D Softmax formulation. All of our experiments are trained and evaluated on GuitarSet following the six-fold cross-validation schema laid out in [16]. There is one major difference in that we hold out one extra dataset split for validation. This means that per fold, only four splits are used for training, while the other two are used for validation and evaluation, respectively. Within each fold, the criterion for choosing the model checkpoint to evaluate on the evaluation split is the checkpoint with the maximum tablature f-measure on the validation split.

The purpose of first few experiments is to verify our reproduction of TabCNN and our experimental setup, and to investigate the difference in performance when using a validation set. The models in these experiments are trained with a batch size of 30 for 10000 iterations, where 200 consecutive frame groups within each track in the training set are sampled per iteration. Experiment (1) features our reproduction of TabCNN, with the standard 6D Softmax output layer formulation. In order to try to match the original TabCNN results as closely as possible, we do not perform validation in this experiment, and simply evaluate the final models within each fold at 10000 iterations. Next, in experiment (2), we run the same experiment but with the validation methodology outlined above. All of the remaining experiments utilize the same validation methodology.

The purpose of the remaining experiments is to compare the 6D Softmax and Logistic output layer formulations directly. As discussed in Sec. 2.1, an LSTM is inserted before the output layer of each model in these experiments as a simple improvement, and to observe what happens to the proposed metrics when a simple language model is added. In order to create a balance between the sequence length and the batch size when training the models with the LSTM, we also modify the training hyperparameters such that we train with a batch size of 50 with a sequence length of 125 frame groups for 50000 iterations. We increase the amount of training iterations to account for the additional model complexity of the LSTM.

Experiment (3) is the same as Experiment (2), except for the insertion of the LSTM layer and the new training hyperparameters. Experiment (4) features the Logistic output layer formulation detailed in Sec. 2.2, but with an ineffective inhibition objective, *i.e.* $\lambda = 0$. In the remaining experiments, $\lambda \neq 0$, and various different matrices of inhibition weights are employed. In Experiment (5), a set of weights which only cover the hard string constraints is used for the inhibition loss. The weights can be expressed as the following:

$$w(c_i, c_j) = \begin{cases} 1 & \text{if } |i - j| < (F + 2) \text{ and } i \neq j \\ 0 & \text{otherwise} \end{cases}. \quad (12)$$

This matrix of weights inhibits all of the pairs which are intrinsically impossible with the 6D Softmax formulation. A visualization of the complement to these weights is provided in Fig. 2. All other experiments employ matrices estimated using DadaGP [17] with the procedure detailed in Sec. 2.3. The standard inhibition weights (*DadaGP*) are used in Experiment (6). Experiment (7) and (8) use inhibition weights boosted with $b = 2^7$ (*DadaGP*⁺). The




Figure 2. Complement of the inhibition weights which enforce hard string constraints, described by Equation (12).

estimated pairwise likelihoods corresponding to each set of weights are presented in Fig. 1. Experiments (4-7) use an inhibition loss multiplier of $\lambda = 1$, whereas Experiment (8) uses $\lambda = 10$. These experiments make up a sort of ablation study w.r.t. the design choices of the inhibition objective within the Logistic formulation.

4. RESULTS & DISCUSSION

4.1 Transcription

The transcription results for the experiments outlined in Sec. 3.3 are presented in Table 1. The results of Experiment (1), our reproduction of TabCNN [16], are almost identical to what was originally reported. Regarding Experiment (2), validation lowers transcription performance slightly, as expected, but to a surprisingly small degree. This validates our decision to trade an extra dataset split for a more justifiable selection criterion. In Experiment (3), insertion of the LSTM significantly improves tablature transcription performance, mainly by increasing recall.

Without the inhibition objective, the logistic formulation in Experiment (4) yields essentially the same performance as the 6D Softmax formulation in Experiment (3). When using the inhibition objective with the string constrains in Experiment (5), the overall tablature transcription performance improves slightly. Since the weights only discourage activations on the same string from co-occurring, similar to the 6D Softmax formulation, this is an interesting result. With the *DadaGP* weights in Experiment (6), the inhibition objective lowers performance slightly by decreasing tablature and multipitch recall. This is not surprising, since the weights are very strict and inhibit almost everything besides perfect $4^{th}/5^{th}$ intervals. With the boosted weights *DadaGP+* in Experiments (7-8), recall improves and precision drops slightly for tablature and multipitch.




Figure 3. Estimated pairwise likelihoods for all S/F combinations computed using GuitarSet [20] with $b = 2^7$.

Finally, the stronger inhibition objective ($\lambda = 10$) in Experiment (8) lowers overall performance for multipitch estimation, while maintaining roughly the same tablature transcription performance w.r.t. Experiment (7).

The lack of an increase in tablature performance when using the inhibition objective can most likely be attributed to the small size of GuitarSet [20] and the presence of some noisy labels. Using the procedure detailed in Sec. 2.3, pairwise likelihoods were estimated from GuitarSet [20] for analysis. These are illustrated in Fig. 3. Clearly, there are co-occurrences for pairs which are highly inhibited by the *DadaGP* and *DadaGP+* weights. Upon inspection of the tracks which produced these artifacts, we found some instances of duplicate pitch errors and octave errors in the annotations. It could be that the inhibition variants actually avoided making these types of predictions, but received a lower score due to the likely annotation errors. Furthermore, the rock and metal bias in *DadaGP* [17] may have skewed the distribution of model predictions in the relevant models away from the distribution of GuitarSet [20], which contains genres such as Jazz and Bossa Nova.

4.2 Loss & Errors

The distribution and error measurements for the experiments outlined in Sec. 3.3 are presented in Table 2. Some key observations can be made about these results. First, the standard inhibition loss L_{inh} remains relatively consistent across experiments. As noted earlier, the *DadaGP* weights are very strict and inhibit most pairs, so the inhibition objective can conflict with the transcription objective. For this reason, it makes sense that among Experiments (3-8), Experiment (6) had the most influence on this metric, given that it used L_{inh} for training. We do observe a significant decrease in L_{inh}^+ for Experiments (7-8), which

Tablature Layer	p_{tab}	r_{tab}	f_{tab}	p_{pitch}	r_{pitch}	f_{pitch}	TDR
(1) <i>Reproduction</i>	0.809	0.692	0.742	0.910	0.762	0.825	0.903
(2) <i>Reproduction w/ Val.</i>	0.775	0.696	0.730	0.895	0.781	0.830	0.886
(3) <i>Reproduction w/ Val./Rec.</i>	0.783	0.757	0.768	0.879	0.835	0.854	0.905
(4) <i>Logistic (No Inhibition)</i>	0.782	0.757	0.767	0.878	0.836	0.854	0.902
(5) <i>Logistic w/ String Constraints</i>	0.789	0.761	0.773	0.881	0.836	0.856	0.907
(6) <i>Logistic w/ DadaGP</i>	0.787	0.743	0.763	0.880	0.821	0.847	0.902
(7) <i>Logistic w/ DadaGP⁺</i>	0.782	0.754	0.766	0.876	0.833	0.852	0.902
(8) <i>Logistic w/ DadaGP⁺ ($\lambda = 10$)</i>	0.781	0.755	0.766	0.867	0.829	0.845	0.907

Table 1. Average six-fold cross validation results on GuitarSet [20] for transcription metrics. Bold values indicate the highest observed result for each metric across experiments which follow the validation methodology outline in Sec. 3.3 (i.e., Experiments (2-8)). The break separates experiments without recurrence and experiments which used an LSTM.

Tablature Layer	L_{inh}	L_{inh}^+	$E_{d.p.}$	$E_{f.a.}$
(1) <i>Reproduction</i>	8.87	0.132	21.4	359.8
(2) <i>Reproduction w/ Val.</i>	9.01	0.152	34.2	442.5
(3) <i>Reproduction w/ Val./Rec.</i>	9.27	0.158	24.3	489.6
(4) <i>Logistic (No Inhibition)</i>	9.27	0.154	20.0	503.3
(5) <i>Logistic w/ String Constraints</i>	9.25	0.155	19.5	485.8
(6) <i>Logistic w/ DadaGP</i>	9.19	0.147	12.0	481.8
(7) <i>Logistic w/ DadaGP⁺</i>	9.25	0.143	13.8	496.6
(8) <i>Logistic w/ DadaGP⁺ ($\lambda = 10$)</i>	9.26	0.132	10.6	504.6

Table 2. Average six-fold cross validation results on GuitarSet [20] for distribution and error metrics. Bold values indicate the lowest observed result for each metric across experiments which follow the validation methodology outline in Sec. 3.3 (i.e., Experiments (2-8)). The break separates experiments without recurrence and experiments which used an LSTM.

train with the boosted inhibition loss. In contrast to L_{inh} , even when L_{inh}^+ continues to go down, the transcription performance improves for Experiments (6-8). This agrees our hypothesis that the boosted inhibition weights are more suitable for use with the transcription objective.

We also notice that, despite a significant increase in Experiment (2) and subsequently Experiment (3), the average number of false alarm errors $E_{f.a.}$ remains relatively consistent across Experiments (3-8). It does seem to improve slightly with inhibition using string constraints and the *DadaGP* weights. The average number of duplicate pitch errors $E_{d.p.}$ improves significantly in Experiments (6-8). As discussed before, there are cases of duplicate pitches in the ground-truth of GuitarSet. The reason why the duplicate pitch error count is not lower may be because the models are trained to produce duplicate pitch predictions in some scenarios. Overall, we argue that the lower $E_{d.p.}$ and L_{inh}^+ suggests that models trained with *DadaGP* and *DadaGP⁺* produce tablature which is more feasible to play and more consistent with *DadaGP* [17].

4.3 Future Work

Although the inhibition objective was presented here in the context of improving guitar tablature transcription, it has several other potential uses. Inhibition shapes the distribution of model predictions by inhibiting unlikely S/F pairs. This means that one can estimate the pairwise likelihoods using a curated distribution, e.g. a collection of tablature

corresponding to a specific musician or genre. Similarly, collections based on playing difficulty could be curated to influence the model to produce tablature more suitable for specific users with varying proficiency.

Another usage could be within the context of tablature arrangement, where the inhibition objective could be applied to train a model to allocate a set of preexisting pitches to strings, such that the resulting fingerings are playable. This may even be useful in a two-stage approach to tablature transcription, where a generic multipitch estimation model feeds predictions into the arrangement system.

We also suggest several directions for improving the inhibition objective. First, one could explore various types of data augmentation for symbolic tablature. One example we refer to as capo augmentation, where a constant fret offset is added to all notes within a track. This could prevent issues related to a dataset’s lack of fretboard coverage. Another interesting direction would be the inclusion of a temporal inhibition objective at the note-level. This could prevent a model from generating predictions which shift around the fretboard too much. Lastly, it would be interesting to investigate higher-order S/F relationships (e.g., 3 notes or more), since the current method only takes pairwise relationships into account.

5. CONCLUSION

We propose a new output layer formulation for guitar tablature transcription which takes advantage of large collec-

tions of symbolic tablature data. The pairwise likelihood of concurrent activation for all possible notes on the guitar is estimated using a recently published dataset. The complement of the pairwise likelihood is used as a weight for an accompanying inhibitory objective during training. We compare the new formulation against the output layer formulation of a baseline tablature transcription model. The inhibition objective is shown to be effective in shaping the distribution of the output predictions and lowering the number of duplicate pitch errors. We also discuss alternative uses and future directions for the inhibition objective.

Acknowledgments

This work has been partially funded by the National Science Foundation grants IIS-1846184 and DGE-1922591.

6. REFERENCES

- [1] E. Benetos, S. Dixon, Z. Duan, and S. Ewert, "Automatic music transcription: An overview," *IEEE Signal Processing Magazine*, vol. 36, no. 1, pp. 20–30, 2019.
- [2] X. Fiss and A. Kwasinski, "Automatic real-time electric guitar audio transcription," in *Proceedings of ICASSP*, 2011.
- [3] L. Alcabasa and N. Marcos, "Automatic guitar music transcription," in *Proceedings of the International Conference on Advanced Computer Science Applications and Technologies (ACSAT)*, 2012.
- [4] C. Kehling, J. Abeßer, C. Dittmar, and G. Schuller, "Automatic tablature transcription of electric guitar recordings by estimation of score and instrument-related parameters," in *Proceedings of DAFx*, 2014.
- [5] I. Barbancho, L. J. Tardon, S. Sammartino, and A. M. Barbancho, "Inharmonicity-based method for the automatic generation of guitar tablature," *IEEE Transactions on Audio, Speech, and Language Processing (TASLP)*, vol. 20, no. 6, pp. 1857–1868, 2012.
- [6] G. Burlet and I. Fujinaga, "Robotaba guitar tablature transcription framework," in *Proceedings of ISMIR*, 2013.
- [7] K. Yazawa, D. Sakaue, K. Nagira, K. Itoyama, and H. G. Okuno, "Audio-based guitar tablature transcription using multipitch analysis and playability constraints," in *Proceedings of ICASSP*, 2013.
- [8] K. Yazawa, K. Itoyama, and H. G. Okuno, "Automatic transcription of guitar tablature from audio signals in accordance with player's proficiency," in *Proceedings of ICASSP*, 2014.
- [9] A. M. Barbancho, A. Klapuri, L. J. Tardon, and I. Barbancho, "Automatic transcription of guitar chords and fingering from audio," *IEEE Transactions on Audio, Speech, and Language Processing (TASLP)*, vol. 20, no. 3, pp. 915–921, 2011.
- [10] G. Hori, H. Kameoka, and S. Sagayama, "Input-output HMM applied to automatic arrangement for guitars," *Information and Media Technologies*, vol. 8, no. 2, pp. 477–484, 2013.
- [11] G. Burlet and A. Hindle, "Isolated guitar transcription using a deep belief network," *PeerJ Computer Science*, vol. 3, p. e109, 2017.
- [12] J. Michelson, R. Stern, and T. Sullivan, "Automatic guitar tablature transcription from audio using inharmonicity regression and bayesian classification," *Journal of the Audio Engineering Society (AES)*, 2018.
- [13] J. M. Hjerrild and M. G. Christensen, "Estimation of guitar string, fret and plucking position using parametric pitch estimation," in *Proceedings of ICASSP*, 2019.
- [14] J. M. Hjerrild, S. Willemsen, and M. G. Christensen, "Physical models for fast estimation of guitar string, fret and plucking position," in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2019.
- [15] E. J. Humphrey and J. P. Bello, "From music audio to chord tablature: Teaching deep convolutional networks to play guitar," in *Proceedings of ICASSP*, 2014.
- [16] A. Wiggins and Y. Kim, "Guitar tablature estimation with a convolutional neural network," in *Proceedings of ISMIR*, 2019.
- [17] P. Sarmiento, A. Kumar, C. Carr, Z. Zukowski, M. Barthelet, and Y.-H. Yang, "DadaGP: A dataset of tokenized GuitarPro songs for sequence models," in *Proceedings of ISMIR*, 2021.
- [18] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [19] C. Hawthorne, E. Elsen, J. Song, A. Roberts, I. Simon, C. Raffel, J. Engel, S. Oore, and D. Eck, "Onsets and frames: Dual-objective piano transcription," in *Proceedings of ISMIR*, 2018.
- [20] Q. Xi, R. M. Bittner, J. Pauwels, X. Ye, and J. P. Bello, "GuitarSet: A dataset for guitar transcription," in *Proceedings of ISMIR*, 2018.
- [21] S. Abakumov, "PyGuitarPro," [Online], available at: <https://github.com/Perlence/PyGuitarPro>.
- [22] E. J. Humphrey, J. Salamon, O. Nieto, J. Forsyth, R. M. Bittner, and J. P. Bello, "JAMS: A JSON annotated music specification for reproducible MIR research," in *Proceedings of ISMIR*, 2014.