

# Tatum-Level Drum Transcription Based on a Convolutional Recurrent Neural Network with Language Model-Based Regularized Training

Ryoto Ishizuka, Ryo Nishikimi, Eita Nakamura, and Kazuyoshi Yoshii  
 Graduate School of Informatics, Kyoto University, Kyoto, Japan  
 E-mail: {ishizuka, nishikimi, enakamura, yoshii}@sap.ist.i.kyoto-u.ac.jp

**Abstract**—This paper describes a neural drum transcription method that detects from music signals the onset times of drums at the *tatum* level, where tatum times are assumed to be estimated in advance. In conventional studies on drum transcription, deep neural networks (DNNs) have often been used to take a music spectrogram as input and estimate the onset times of drums at the *frame* level. The major problem with such frame-to-frame DNNs, however, is that the estimated onset times do not often conform with the typical tatum-level patterns appearing in symbolic drum scores because the long-term musically meaningful structures of those patterns are difficult to learn at the frame level. To solve this problem, we propose a regularized training method for a frame-to-tatum DNN. In the proposed method, a tatum-level probabilistic language model (gated recurrent unit (GRU) network or repetition-aware bi-gram model) is trained from an extensive collection of drum scores. Given that the musical naturalness of tatum-level onset times can be evaluated by the language model, the frame-to-tatum DNN is trained with a regularizer based on the pretrained language model. The experimental results demonstrate the effectiveness of the proposed regularized training method.

## I. INTRODUCTION

Automatic drum transcription (ADT) is a challenging sub-task in automatic music transcription (AMT) that aims to estimate *symbolic* musical scores from music signals. This is an important task because the drum part forms the rhythmic backbone of popular music. In this paper, we focus on the three main drum instruments of a drum kit: bass drums (BD), snare drums (SD), and hi-hats (HH). In general, the estimated onset times of drums are represented at the frame level (in seconds) and very few ADT methods aim to estimate symbolic drum scores on regular time grids.

In ADT, deep learning and nonnegative matrix factorization (NMF) have been two major approaches to estimating the onset activations of drums from audio spectrograms *at the frame level* [1]. In particular, convolutional neural networks (CNNs) that can extract useful features from local time-frequency regions have shown good performances [2]–[4]. Recurrent neural networks (RNNs) have also been used for learning the frame-level long-term dependency of onset activations [5]–[7]. Because the spectrogram of a drum part consists of a number of repetitions of the same impulsive sounds, NMF has still been used extensively for ADT [8]–[12].

However, as these purely frame-level ADT methods have no mechanism to prevent the estimated onset times of drums

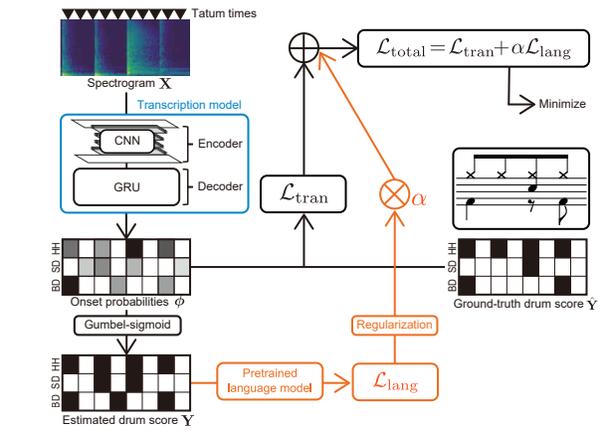


Fig. 1. Supervised training of a neural drum transcription model with musical naturalness-aware output regularization based on a pretrained language model.

from having a musically unnatural structure, the rhythmic and repetitive patterns of drum onsets are hard to be learned at the frame level. One solution to this problem is to use a *language model* (LM) that represents a probability distribution of drum onsets *at the tatum level* such that the musical naturalness of drum scores can be evaluated. Such an LM (drum score prior) has been integrated with an NMF-based acoustic model (drum score likelihood) in a Bayesian manner [10]. The performance of this method, however, remains unsatisfactory because of the limited expressive power of NMF and the time-consuming posterior inference of drum scores required at run-time.

LMs play an essential role in automatic speech recognition (ASR) for estimating a word sequence from a feature sequence such that the estimated word sequence is syntactically and semantically coherent. The classical yet effective approach to ASR is to combine a *word-level* language model (*e.g.*, n-gram model) representing the generative process of a word sequence with a *frame-level* acoustic model (*e.g.*, hidden Markov model (HMM)) representing the generative process of a feature sequence from a word sequence [13]. To infer a word sequence from a feature sequence using Bayes' theorem, a sophisticated decoder (*e.g.*, weighted finite-state transducer (WFST)) based on the language and acoustic models is used at run-time. The advantage of this approach lies in its modularity; the language and acoustic models can be trained from text data and paired

data (speech data with transcriptions), respectively.

Recently, the end-to-end approach to ASR has been actively investigated for directly inferring a word sequence from a feature sequence with a deep neural network (DNN). A popular choice is to use an encoder-decoder architecture with an attention mechanism, where the encoder and decoder are conceptually considered to have acoustic and language modeling capabilities, respectively [14]. While such a network is easy to implement and works fast at run-time, only paired data can be used for training the whole network, which means that massive text data cannot be used for improving the language modeling capability of the network. To use the knowledge of an LM trained on huge text data in an end-to-end recognizer, knowledge transfer techniques [15] have been investigated [16], [17].

In light of these circumstances, we propose an ADT method based on a convolutional RNN (CRNN) that directly infers a sequence of *tatum-level* onset times from a sequence of *frame-level* mel spectra (Fig. 1). We do not directly use an encoder-decoder architecture with an attention mechanism, which has widely been used for sequence-to-sequence learning in various applications, including ASR. In practice, the beat times can be estimated accurately for typical popular music with regular rhythmic structure (our main target) and attention-based alignment between long frame- and tatum-level sequences is hard to learn from a limited amount of training data. We thus use the estimated tatum times instead of using an attention mechanism for combining the frame-level convolutional layers (encoder) extracting useful features from mel spectra and tatum-level recurrent layers (decoder) learning the rhythmic and repetitive patterns of drum onsets.

To transfer the knowledge of an LM trained from a large number of drum scores, we train the CRNN in a regularized manner. More specifically, we aim to minimize the weighted sum of the transcription error  $\mathcal{L}_{\text{tran}}$  and the musical unnaturalness  $\mathcal{L}_{\text{lang}}$  computed for the CRNN output, where  $\mathcal{L}_{\text{tran}}$  is the cross entropy between the estimated *soft* drum score and the ground-truth score and  $\mathcal{L}_{\text{lang}}$  is the LM-based negative log-probability of the estimated *hard* (binarized) score. Note that the hard score is obtained by applying the gumbel-sigmoid trick [18] to the soft score in a differentiable manner for backpropagation-based optimization.

## II. RELATED WORK

This section reviews related work on AMT and ADT based on language models and knowledge transfer.

### A. Automatic Drum Transcription

Nonnegative matrix factorization (NMF) has often been used for decomposing a drum-part spectrogram into the spectra and temporal activations of drums [8], [11], [12]. To overcome the limited expressive power of NMF, CNNs have been used in ADT [2]–[4] for automatically extracting local features as well as in AMT [19]. RNNs have also been proposed for capturing the long-term temporal dependency at the frame level. Vogl *et al.* introduced RNNs [5] as well as multi-task learning [7] in ADT. In such ways, DNN-based transcription

methods, which are trained by paired data consisting of audio signals with annotations, have achieved high performances.

### B. Language model

One way of improving AMT and ADT methods is to introduce an LM that evaluates the musical naturalness of estimated scores. Such LMs have generally been formulated at the frame level. Raczyński *et al.* [20], for example, used a deep belief network for modeling a transition of chord symbols and improved the chord recognition system consisting of NMF. Sigia *et al.* [21] used an LM for estimating the most likely chord sequence from the chord posterior probabilities estimated by an RNN-based chord recognition system. As pointed out in [22], [23], however, LMs can be more effectively formulated at the tatum level for learning the musically meaningful structure.

Tatum-level LMs have recently been used for AMT and ADT. Korzeniowski *et al.* [24] used N-gram for a symbolic LM and improved the DNN-based chord recognition system. Korzeniowski *et al.* [25] also insisted that a frame-level LM can only smooth the onset probabilities of chord symbols, and they used an RNN-based symbolic LM with a chord duration model. Ycart *et al.* [26] investigated the predictive power of LSTM networks and demonstrated that a long short-term memory (LSTM) working at the level of 16th note timesteps could express the musical structure such as note transitions. Thompson *et al.* [27] used a template-based LM for classifying audio signals into a limited number of drum patterns with a support vector machine (SVM). Ueda *et al.* [10] proposed a Bayesian approach using a DNN-based LM as a prior of drum scores. Integration of a tatum-level LM into a DNN-based ADT system, however, has still been an open problem.

### C. Knowledge Transfer

Transfer learning aims to effectively use knowledge from a related domain [28], and has been used in various fields. This method can be used for labeled data as well as unlabeled data. In the student-teacher framework, some studies have attempted to train a student model that has the same capacity as a teacher model [29] for achieving higher accuracy [30]. Transfer learning using a compact student model is often called knowledge distillation [15].

Recently, there have been some studies on using knowledge learned from an extensive collection of unpaired data for improving other models. In ASR, an LM was integrated into ASR systems to generate more syntactically or semantically word sequences. These methods, however, require an LM in decoding and take much time in the inference stage. More recently, an ASR system based on knowledge distillation was proposed, where an LM softened a probability distribution as a regularizer to transfer the knowledge of unpaired data. The method did not require an LM in the inference stage [16]. The idea of knowledge distillation was also used in ADT [31], where an NMF-based teacher model was applied to a DNN-based student model, and this method showed great potential to utilize unpaired data. Note that in the transfer learning, the

same or different datasets are used depending on the problem specification [15], [32], [33].

### III. PROPOSED METHOD

This section describes the proposed ADT method that estimates a drum score from the mel spectrogram of a music signal (Section III-A). As shown in Fig. 1, our method uses a CRNN-based transcription model for estimating the onset probabilities of drums at the tatum level (Section III-B). Given that a pretrained LM of drum scores can be used for evaluating the musical naturalness of a drum score (Section III-C), the transcription model is trained in a supervised manner with a regularization mechanism based on the pretrained language model (Section III-D).

#### A. Problem Specification

Our goal is to estimate a drum score  $\mathbf{Y} \in \{0, 1\}^{K \times M}$  from the mel spectrogram of a target musical piece  $\mathbf{X} \in \mathbb{R}_+^{F \times T}$ , where  $K$  is the number of drum instruments (BD, SD, and HH, *i.e.*,  $K = 3$ ),  $M$  the number of tatums,  $F$  the number of frequency bins,  $T$  the number of time frames. In this paper, we assume that all onset times are located on the tatum-level (quarter-beat-level) grid and the tatum times  $\mathbf{B} = \{b_m\}_{m=1}^M$  are estimated in advance.

#### B. Transcription Model

The transcription model is used for estimating the tatum-level onset probabilities  $\phi \in [0, 1]^{K \times M}$ , where  $\phi_{k,m}$  represents the posterior probability that drum  $k$  has an onset at tatum  $m$ . The estimated drum score  $\mathbf{Y}$  can be obtained by binarizing  $\phi$  with a threshold  $\delta \in [0, 1]$ . The transcription model is implemented as a CRNN consisting of a frame-level encoder based on convolutional layers and a tatum-level decoder based on GRU layers followed by a fully-connected layer (Fig. 2). The encoder converts the mel spectrogram  $\mathbf{X}$  to the latent features  $\mathbf{F} \in \mathbb{R}^{D \times T}$ , where  $D$  is the feature dimension. The frame-level features  $\mathbf{F}$  are then summarized to the tatum-level features  $\mathbf{G} \in \mathbb{R}^{D \times M}$  through a max pooling layer referring to the tatum times  $\mathbf{B}$  as follows:

$$G_{d,m} = \max_{\frac{b_{m-1}+b_m}{2} \leq t < \frac{b_m+b_{m+1}}{2}} F_{d,t}, \quad (1)$$

where  $b_0 = b_1$  and  $b_{M+1} = b_M$ . The decoder finally converts  $\mathbf{G}$  to the onset probabilities  $\phi$  while considering the temporal dynamics of drum scores.

#### C. Language Model

The LM is used for estimating the generative probability (musical naturalness) of a drum score. To achieve this, using an arbitrary existing drum score  $\tilde{\mathbf{Y}}^1$ , the LM should be trained beforehand in an unsupervised manner such that the following negative log-likelihood for  $\tilde{\mathbf{Y}}$  is minimized:

$$\mathcal{L}_{\text{lang}}(\tilde{\mathbf{Y}}) = - \sum_{m=1}^M \log p(\tilde{\mathbf{Y}}_{:,m} | \tilde{\mathbf{Y}}_{:,1:m-1}), \quad (2)$$

<sup>1</sup>For brevity, we assume that only one drum score is used as training data. In practice, a sufficient number of drum scores are used.

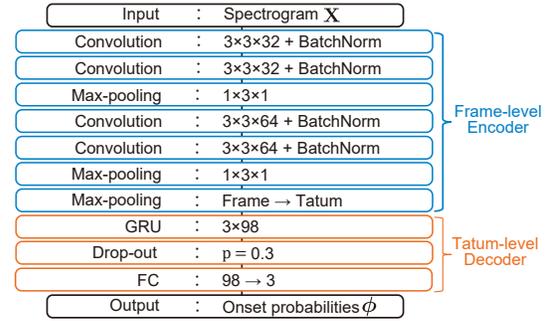


Fig. 2. Configuration of the transcription model.

where “ $i:j$ ” indicates a set of indices from  $i$  to  $j$  and “:” indicates all possible indices. In this paper, we propose two LMs: a *skip-type* bi-gram model and a neural language model.

1) *Repetition-Aware Bi-Gram Model*: One possibility is to use a naive yet effective bi-gram model. Assuming that popular music tends to have the 4/4 time signature and the same drum patterns tend to be repeated for making the rhythmic backbone, we propose a *skip-type* bi-gram model representing the bar-level repetitive structure of  $\tilde{\mathbf{Y}}$  as follows:

$$\begin{aligned} p(\tilde{\mathbf{Y}}_{:,m} | \tilde{\mathbf{Y}}_{:,1:m-1}) &= \prod_{k=1}^K p(\tilde{Y}_{k,m} | \tilde{Y}_{k,m-16}) \\ &= \prod_{k=1}^K \pi_{\tilde{Y}_{k,m-16}, \tilde{Y}_{k,m}}, \end{aligned} \quad (3)$$

where  $\pi_{A,B}$  ( $A, B \in \{0, 1\}$ ) indicates the transition probability from  $A$  to  $B$ . Note that this model assumes the independence of the  $K$  drums.

2) *Gated Recurrent Unit Model*: Another possibility is to use a more powerful neural LM based on GRUs for directly representing  $p(\tilde{\mathbf{Y}}_{:,m} | \tilde{\mathbf{Y}}_{:,1:m-1})$  without assuming the independence of the  $K$  drums. This model is expected to implicitly represent different time signatures and consider a longer-range musically-meaningful temporal structure of drum scores.

#### D. Regularized Training

Given a ground-truth score  $\hat{\mathbf{Y}}$ , one can train the transcription model in a supervised manner such that the following modified negative log-likelihood for  $\hat{\mathbf{Y}}$  is minimized:

$$\begin{aligned} \mathcal{L}_{\text{tran}}(\phi | \hat{\mathbf{Y}}) &= - \sum_{k=1}^K \sum_{m=1}^M (\gamma \hat{Y}_{k,m} \log \phi_{k,m} + (1 - \hat{Y}_{k,m}) \log(1 - \phi_{k,m})), \end{aligned} \quad (4)$$

where  $\gamma > 0$  is a weighting factor compensating for the imbalance of the numbers of onset and non-onset tatums. Because  $\mathcal{L}_{\text{tran}}$  evaluates only the transcription incorrectness (cross entropy between  $\phi$  and  $\hat{\mathbf{Y}}$ ), the musical naturalness of the estimated score  $\mathbf{Y}$  obtained by binarizing  $\phi$  is not considered.

To solve this problem, we propose an LM-based regularized training method that minimizes

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{tran}}(\phi|\hat{\mathbf{Y}}) + \alpha\mathcal{L}_{\text{lang}}(\mathbf{Y}), \quad (5)$$

where  $\alpha > 0$  is a weighting factor. To use a backpropagation technique for optimizing the transcription model, the binary score  $\mathbf{Y}$  should be obtained from the soft representation  $\phi$  in a differentiable manner instead of simply binarizing  $\phi$  with a threshold. We thus use a differentiable sampler called the gumbel-sigmoid trick [18] as follows:

$$U_{k,m}^{(i)} \sim \text{Uniform}(0, 1), \quad (6)$$

$$V_{k,m}^{(i)} = -\log \left\{ -\log \left( U_{k,m}^{(i)} \right) \right\}, \quad (7)$$

$$Y_{k,m} = \sigma \left\{ \frac{\phi_{k,m} + V_{k,m}^{(1)} - V_{k,m}^{(2)}}{\tau} \right\}, \quad (8)$$

where  $i = 1, 2$ ,  $\tau > 0$  is a temperature, and  $\sigma(\cdot)$  is a sigmoid function ( $\tau = 0.2$  in this paper). Note that the pretrained LM (bi-gram or GRU model) is used as a fixed regularizer in the training phase and is not used in the prediction phase.

#### IV. EVALUATION

This section reports experiments conducted for validating the proposed LM-based regularized training of the neural transcription model for ADT.

##### A. Experimental Conditions

The RWC Popular Music Database [34] was used for evaluation. Among 89 songs having drum parts, we used 65 songs with correct ground-truth annotations. These songs were randomly split into training and testing data for 3-fold cross validation, where 15% of the training data was taken as validation data. To extract drum sounds from polyphonic music signals, we used a music separation method called Open-Unmix [35]. The music signals and the separated drum signals were used in the training phase and the separated drum signals were used in the prediction phase. The spectrogram of each music signal sampled at 44.1kHz was obtained using short-time Fourier transform (STFT) with a Hann window of 2048 points (46 ms) and a shifting interval of 441 points (10 ms). The mel-frequency spectrogram was calculated using a mel-filter bank with 80 bands from 20 Hz to 20,000 Hz.

To pretrain the LMs (bi-gram and GRU models described in Section III-C), we used 512 external drum scores (Japanese popular songs and The Beatles), which have no overlap with the RWC Popular Music Database [34]. The GRU model we used consisted of 3 GRU layers with 64 hidden dimensions, which were experimentally determined by a Bayesian optimization method called Optuna [36] via 3-fold cross validation with the 512 scores.

We used madmom [37] for beat estimation and the performance was measured using the precision rate  $\mathcal{P}$ , the recall rate  $\mathcal{R}$ , and the F-measure  $\mathcal{F}$  given by

$$\mathcal{P} = \frac{N_c}{N_e}, \quad \mathcal{R} = \frac{N_c}{N_g}, \quad \mathcal{F} = \frac{2\mathcal{R}\mathcal{P}}{\mathcal{R} + \mathcal{P}}, \quad (9)$$

TABLE I  
THE RATIOS OF UNDETECTABLE ONSET TIMES IN THREE GROUPS.

Madmom			Ground-truth		
<i>conflict</i>	<i>far</i>	<i>conflict</i> $\cup$ <i>far</i>	<i>conflict</i>	<i>far</i>	<i>conflict</i> $\cup$ <i>far</i>
0.43%	0.23%	0.65%	1.19%	0.29%	1.48%

where  $N_e$ ,  $N_g$ , and  $N_c$  were the number of estimated beats, that of ground-truth beats, and that of correctly-estimated beats, respectively. The estimated beat was judged as correct if it was within 50 ms from the ground-truth beat. The mir\_eval library [38] was used for computing  $\mathcal{P}$ ,  $\mathcal{R}$ , and  $\mathcal{F}$ .

##### B. Justification of Tatum-Level Transcription

We validate the appropriateness of our tatum-level transcription approach because there are undetectable drum onsets if all the onset times of each drum are assumed to be exclusively located on tatum (quarter-beat) times. Such undetectable onsets are (doubly) categorized into two groups: *conflict* and *far*. In our experiment, to convert frame-level onset times (e.g., original ground-truth annotations) into a tatum-level score (e.g., estimation target  $\hat{\mathbf{Y}}$ ), each onset time was quantized to the closest tatum time. If multiple onset times are quantized into the same tatum time, only one onset time can be detected, i.e., the other onset times are undetectable and categorized into the *conflict* group. If actual onset times are not within 50 ms from the closest tatum times, they are categorized into the *far* group.

Table I shows the ratios of such undetectable onset times to the total number of actual onset times when the estimated or ground-truth beat times are used for quantization. The beat tracking method [37] achieved 96.4% for the 65 songs used for evaluation. This result justifies our approach at least for the majority of typical popular music because the total ratio of undetectable onset times was sufficiently low.

##### C. Evaluation of Language Modeling

We evaluated the performance of the pretrained LMs for the 65 songs. The perplexities obtained by the skip-type bi-gram model and the GRU model were 1.51 and 1.44, respectively (lower is better). The predictive capability of the GRU model was better than that of the bi-gram model because the bi-gram model assumed the 4/4 time signature with the simple repeating structure. We also confirmed that the perplexities were much smaller than the chance rate of 2. The LMs were expected to work as regularizers and guide the outputs of the transcription model into musically-natural drum patterns.

##### D. Evaluation of Drum Transcription

We evaluated the effectiveness of the proposed regularized training method based on the pretrained LMs. Our transcription model was inspired by the state-of-the-art ADT method [7] (Fig. 2). The encoder consisted of 4 convolutional layers with the kernel size of  $3 \times 3$  and the decoder consisted of 3 GRU layers with 98 hidden dimensions, followed by a drop-out layer ( $p = 0.3$ ). The weighting factor  $\gamma$  in the transcription loss (Eq. (4)) was set to  $\gamma = 0.46$  for the bi-gram model and  $\gamma = 0.61$  for the GRU model. The weighting factor  $\alpha$  in the

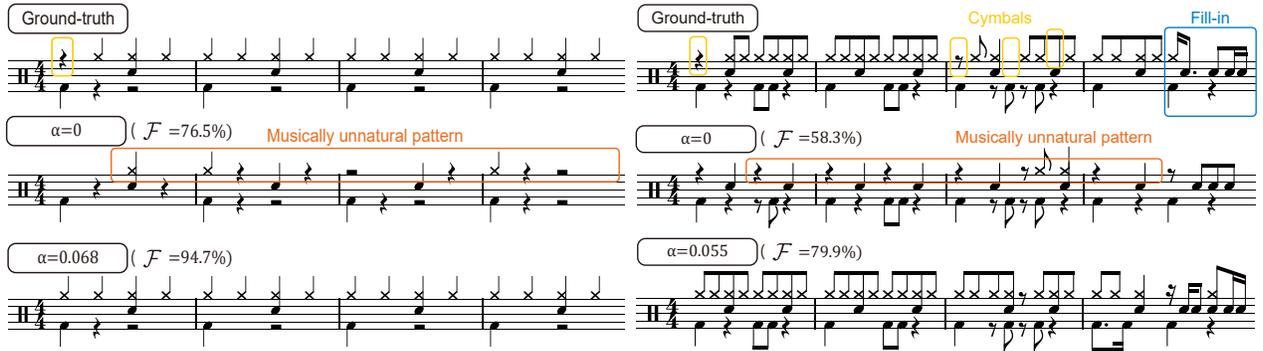


Fig. 3. Examples of transcribed drum scores. The left and right examples are excerpts from RWC-MDB-P-2001 No. 47 and No. 88, respectively. The F-measure computed for each song was shown. The top, middle, and bottom scores show the ground truth, the estimated score obtained by the unconstrained transcription model, and the estimated score obtained by the transcription model regularized by the bi-gram model (left) and GRU model (right).

total loss (Eq. (5)) was set to  $\alpha = 0.068$  for the bi-gram model and  $\alpha = 0.055$  for the GRU model. The influential hyperparameters, *i.e.*, the number of GRU layers, the hidden dimension,  $\gamma$ , and  $\alpha$  were optimized for the validation data with Optuna [36]. The weights of the convolutional and GRU layers were initialized based on [39], the fully connected layer was initialized by the sampling from Uniform(0, 1), and the biases were initialized to 0. We used AdamW optimizer [40] with the initial learning rate of  $10^{-3}$ , the weight decay of  $\lambda=10^{-4}$ ,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ , and  $\epsilon = 10^{-9}$ . The threshold for  $\phi$  was set to  $\delta = 0.2$ .

For comparison, we tested the state-of-the-art purely *frame-level* ADT method [7] based on a CRNN whose architecture was similar to our transcription model. This model was trained with the following frame-level cross entropy:

$$\begin{aligned} \mathcal{L}_{\text{tran}^*}(\phi^*|\hat{Y}^*) &= - \sum_{k=1}^K \sum_{t=1}^T (\beta \hat{Y}_{k,t}^* \log \phi_{k,t}^* + (1 - \hat{Y}_{k,t}^*) \log(1 - \phi_{k,t}^*)), \end{aligned} \tag{10}$$

where  $\phi^*, \hat{Y}^* \in \mathbb{R}^{K \times T}$  are the estimated onset probabilities and the ground-truth binary activations, respectively, and  $\beta > 0$  is a weighting factor ( $\beta = 8$  in this paper). For each drum  $k$ , a frame  $t$  was picked as an onset if

1.  $\phi^* = \max\{\phi_{k,t-w_1:t+w_2}^*\}$ ,
2.  $\phi^* \geq \text{mean}\{\phi_{k,t-w_3:t+w_4}^*\} + \hat{\delta}$ ,
3.  $t - t_{\text{prev}} > w_5$ ,

where  $\hat{\delta}$  was a threshold,  $w_{1:5}$  were interval parameters, and  $t_{\text{prev}}$  was the previous onset frame, which were set to  $\hat{\delta} = 0.2$ ,  $w_1 = w_3 = w_5 = 2$ , and  $w_2 = w_4 = 0$  as in [7]. To measure the tatum-level transcription performance, the estimated frame-level onset times were quantized at the tatum level referring to the estimated or ground-truth tatum times.

Table II shows the performances of the conventional frame-to-frame method [7] followed by the frame-to-tatum quantization (post-processing) and the proposed frame-to-tatum method when the estimated or ground-truth beat times were given.

TABLE II  
THE DRUM TRANSCRIPTION PERFORMANCES OF THE CONVENTIONAL AND PROPOSED METHODS (%).

	Madmom			Ground-truth		
	$\mathcal{F}$	$\mathcal{P}$	$\mathcal{R}$	$\mathcal{F}$	$\mathcal{P}$	$\mathcal{R}$
CRNN [7]	70.8	77.4	65.9	71.0	77.6	66.1
CRNN	78.9	86.3	73.1	79.3	86.7	73.3
+ Bi-gram ( $\alpha = 0.068$ )	81.4	84.7	79.1	80.8	83.7	78.8
+ GRU ( $\alpha = 0.055$ )	<b>81.6</b>	84.0	80.2	<b>81.1</b>	83.2	79.7

We confirmed that the regularization method was effective for improving the transcription model. The regularization with the GRU model improved the F-measure by a larger margin than that with the bi-gram model.

Fig. 3 illustrates two examples of transcribed drum scores, which show the positive effect of the language model-based regularization. In both examples, the non-regularized transcription model often yielded musically-unnatural drum patterns, while the regularized model effectively avoided such patterns. The regularized model, however, yielded extra onset times of hi-hats because the other kinds of percussive instruments (crash cymbals in both cases) were used instead of hi-hats. We also found that the regularization mechanism was effective to estimate regular drum patterns, but tended to oversimplify highly-sophisticated non-regular drum patterns (*e.g.*, fill-ins).

## V. CONCLUSION

This paper described a tatum-level ADT method based on a CRNN trained with an LM-based regularization mechanism. This network consists of a frame-level convolutional encoder extracting the latent features of music signals and a tatum-level recurrent decoder considering musically-meaningful structure. The experimental results showed that the regularized training significantly improves both the correctness and musical naturalness of estimated drum scores.

Extending this approach, we plan to deal with sophisticated and/or non-regular drum patterns (*e.g.*, fill-ins) played by various kinds of percussive instruments (*e.g.*, cymbals and toms). Considering that beat and downbeat times are closely related to drum patterns, it would be beneficial to integrate beat tracking into ADT in a multi-task learning framework.

## ACKNOWLEDGMENT

This work is partially supported by JST ACCEL No. JPM-JAC1602 and JSPS KAKENHI No. 16H01744, No. 19K20340, and No. 19H04137.

## REFERENCES

- [1] Chih-Wei Wu, Christian Dittmar, Carl Southall, Richard Vogl, Gerhard Widmer, Jason Hockman, Meinard Müller, and Alexander Lerch. A review of automatic drum transcription. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 26(9):1457–1483, 2018.
- [2] Céline Jacques and Axel Roebel. Automatic drum transcription with convolutional neural networks. In *DAFx*, pages 80–86, 2018.
- [3] Nicolai Gajhede, Oliver Beck, and Hendrik Purwins. Convolutional neural networks with batch normalization for classifying hi-hat, snare, and bass percussion sound samples. In *Proceedings of the Audio Mostly 2016*, pages 111–115, 2016.
- [4] Carl Southall, Ryan Stables, and Jason Hockman. Automatic drum transcription for polyphonic recordings using soft attention mechanisms and convolutional neural networks. In *ISMIR*, pages 606–612, 2017.
- [5] Richard Vogl, Matthias Dorfer, and Peter Knees. Recurrent neural networks for drum transcription. In *ISMIR*, pages 730–736, 2016.
- [6] Ryan Stables, Jason Hockman, and Carl Southall. Automatic drum transcription using bi-directional recurrent neural networks. In *ISMIR*, pages 591–597, 2016.
- [7] Richard Vogl, Matthias Dorfer, Gerhard Widmer, and Peter Knees. Drum transcription via joint beat and drum modeling using convolutional recurrent neural networks. In *ISMIR*, pages 150–157, 2017.
- [8] Chih-Wei Wu and Alexander Lerch. Drum transcription using partially fixed non-negative matrix factorization with template adaptation. In *ISMIR*, pages 257–263, 2015.
- [9] Axel Roebel, Jordi Pons, Marco Liuni, and Mathieu Lagrangey. On automatic drum transcription using non-negative matrix deconvolution and itakura saito divergence. In *ICASSP*, pages 414–418, 2015.
- [10] Shun Ueda, Kentaro Shibata, Yusuke Wada, Ryo Nishikimi, Eita Nakamura, and Kazuyoshi Yoshii. Bayesian drum transcription based on nonnegative matrix factor decomposition with a deep score prior. In *ICASSP*, pages 456–460, 2019.
- [11] Christian Dittmar and Daniel Gärtner. Real-time transcription and separation of drum recordings based on nmf decomposition. In *DAFx*, pages 187–194, 2014.
- [12] Jouni Paulus and Anssi Klapuri. Drum sound detection in polyphonic music with hidden markov models. *EURASIP Journal on Audio, Speech, and Music Processing*, 2009:1–9, 2009.
- [13] Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N.Sainath Sainath, and Kingsbury Brian. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Processing Magazine*, 29(6):82–97, 2012.
- [14] Jan Chorowski, Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. End-to-end continuous speech recognition using attention-based recurrent NN: First results. In *NIPS Workshop on Deep Learning*, 2014.
- [15] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv*, 2015.
- [16] Gernot Kubin and Zdravko Kacic. Learn spelling from teachers: Transferring knowledge from language models to sequence-to-sequence speech recognition. In *Interspeech*, pages 3795–3799, 2019.
- [17] Yen-Chun Chen, Zhe Gan, Yu Cheng, Jingzhou Liu, and Jingjing Liu. Distilling the knowledge of bert for text generation. *arXiv*, 2019.
- [18] Yi-Hsuan Tsai, Ming-Yu Liu, Deqing Sun, Ming-Hsuan Yang, and Jan Kautz. Learning binary residual representations for domain-specific video streaming. In *AAAI*, pages 7363–7370, 2018.
- [19] Jan Schlüter and Sebastian Böck. Improved musical onset detection with convolutional neural networks. In *ICASSP*, pages 6979–6983, 2014.
- [20] Stanisław A Raczynski, Emmanuel Vincent, and Shigeki Sagayama. Dynamic bayesian networks for symbolic polyphonic pitch modeling. *IEEE Transactions on Audio, Speech, and Language Processing*, 21(9):1830–1840, 2013.
- [21] Siddharth Sigtia, Nicolas Boulanger-Lewandowski, and Simon Dixon. Audio chord recognition with a hybrid recurrent neural network. In *ISMIR*, pages 127–133, 2015.
- [22] Filip Korzeniowski and Gerhard Widmer. On the futility of learning complex frame-level language models for chord recognition. In *ISMIR*, pages 10–17, 2017.
- [23] Adrien Ycart, Andrew McLeod, Emmanouil Benetos, Kazuyoshi Yoshii, et al. Blending acoustic and language model predictions for automatic music transcription. In *ISMIR*, pages 454–461, 2019.
- [24] Filip Korzeniowski and Gerhard Widner. Automatic chord recognition with higher-order harmonic language modelling. In *EUSIPCO*, pages 1900–1904, 2018.
- [25] Filip Korzeniowski and Gerhard Widmer. Improved chord recognition by combining duration and harmonic language models. In *ISMIR*, pages 10–17, 2018.
- [26] Adrien Ycart, Emmanouil Benetos, et al. A study on lstm networks for polyphonic music sequence modelling. In *ISMIR*, pages 421–427, 2017.
- [27] Lucas Thompson, Matthias Mauch, Simon Dixon, et al. Drum transcription via classification of bar-level rhythmic patterns. In *ISMIR*, pages 187–192, 2014.
- [28] Karl Weiss, Taghi M Khoshgoftaar, and DingDing Wang. A survey of transfer learning. *Journal of Big data*, 3(1):9, 2016.
- [29] Hossein Mobahi, Mehrdad Farajtabar, and Peter L Bartlett. Self-distillation amplifies regularization in hilbert space. *arXiv*, 2020.
- [30] Tommaso Furlanello, Zachary C Lipton, Michael Tschannen, Laurent Itti, and Anima Anandkumar. Born again neural networks. In *ICML*, pages 1602–1611, 2018.
- [31] Chih-Wei Wu and Alexander Lerch. Automatic drum transcription using the student-teacher learning paradigm with unlabeled music data. In *ISMIR*, pages 613–620, 2017.
- [32] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *ICLR*, 2017.
- [33] Junho Yim, Donggyu Joo, Jihoon Bae, and Junmo Kim. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *CVPR*, pages 4133–4141, 2017.
- [34] Masataka Goto, Hiroki Hashiguchi, Takuichi Nishimura, and Ryuichi Oka. Rwc music database: Popular, classical and jazz music databases. In *ISMIR*, pages 287–288, 2002.
- [35] Fabian-Robert Stöter, Stefan Uhlich, Antoine Liutkus, and Yuki Mitsufuji. Open-unmix - A reference implementation for music source separation. *J. Open Source Softw.*, 4(41):1667, 2019.
- [36] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *ACM-SIGKDD*, pages 2623–2631, 2019.
- [37] Sebastian Böck, Filip Korzeniowski, Jan Schlüter, Florian Krebs, and Gerhard Widmer. Madmom: A new python audio and music signal processing library. In *ACM international conference on Multimedia*, pages 1174–1178, 2016.
- [38] Colin Raffel, Brian McFee, Eric J Humphrey, Justin Salamon, Oriol Nieto, Dawen Liang, Daniel PW Ellis, and C Colin Raffel. mir\_eval: A transparent implementation of common mir metrics. In *ISMIR*, pages 367–372, 2014.
- [39] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *ICCV*, pages 1026–1034, 2015.
- [40] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019.