

High-resolution Piano Transcription with Pedals by Regressing Onsets and Offsets Times

Qiuqiang Kong, Bochen Li, Xuchen Song, Yuan Wan, Yuxuan Wang

Abstract—Automatic music transcription (AMT) is the task of transcribing audio recordings into symbolic representations such as Musical Instrument Digital Interface (MIDI). Recently, neural networks based methods have been applied to AMT, and have achieved state-of-the-art result. However, most of previous AMT systems predict the presence or absence of notes in the frames of audio recordings. The transcription resolution of those systems are limited to the hop size time between adjacent frames. In addition, previous AMT systems are sensitive to the misaligned onsets and offsets labels of audio recordings. For high-resolution evaluation, previous works have not investigated AMT systems evaluated with different onsets and offsets tolerances. For piano transcription, there is a lack of research on building AMT systems with both note and pedal transcription. In this article, we propose a high-resolution AMT system trained by regressing precise times of onsets and offsets. In inference, we propose an algorithm to analytically calculate the precise onsets and offsets times of note and pedal events. We build both note and pedal transcription systems with our high-resolution AMT system. We show that our AMT system is robust to misaligned onsets and offsets labels compared to previous systems. Our proposed system achieves an onset F1 of 96.72% on the MAESTRO dataset, outperforming the onsets and frames system from Google of 94.80%. Our system achieves a pedal onset F1 score of 91.86%, and is the first benchmark result on the MAESTRO dataset. We release the source code of our work at https://github.com/bytedance/piano_transcription.

Index Terms—Piano transcription, pedal transcription, high-resolution.

I. INTRODUCTION

Automatic music transcription (AMT) [1], [2], [3] is the task of transcribing audio recordings into symbolic representations [4], such as piano rolls, guitar fretboard charts and Musical Instrument Digital Interface (MIDI) files. AMT is an essential topic of music information retrieval (MIR), and is a bridge between audio based and symbolic based music understanding. An AMT system can benefit several MIR tasks, such as score following [5], audio to score alignment [6], and score-informed source separation [7]. In addition, AMT systems can be used in music education systems to assist music learners. For music production, AMT can be used to transcribe audio recordings to MIDI files for later editing. In addition, AMT can be used for symbolic based music information retrieval, and can be used to study unarchieved music, such as jazz improvisations.

Piano transcription is one problem of AMT to transcribe piano recordings into note events with pitch, onset, offset and velocity. Piano transcription is a challenging task due to the high polyphony in music pieces. Early works of piano

transcription [8], [9], [10] include using discriminative models, such as support vector machines to predict the presence or absence of notes in each audio frame [11]. A probabilistic spectral smoothness principle was proposed for multiple pitches estimation in [12]. A combination of frequency domain and time domain method has been proposed for piano transcription in [13]. Non-negative matrix factorizations (NMFs) has been proposed to decompose spectral to polyphonic notes [14]. An attack and decay system was proposed to model different state of piano onsets in [15]. An unsupervised learning method was used for piano transcription in [16]. A connectionist approach [17], and a fast convolutional sparse coding was proposed for piano transcription [18]. Recently, neural networks have been applied for piano transcription. A deep belief network was proposed to learn feature representations for music transcription in [19]. Fully connect neural networks, convolutional neural networks (CNNs) and recurrent neural networks (RNNs) [20], [21], [22], [23] have been proposed to learn regressions from audio input to labelled ground truths. Recently, onsets and frames systems [24], [25] have been proposed to predict both onsets and frame-wise pitches of notes, and have achieved state-of-the-art result in piano transcription. Several improved targets for piano transcription include SoftLoc [26] and non-discrete annotations [27], [28], [29].

Previous piano transcription systems such as [24], [25] first split audio recordings into frames. Then, piano rolls are created as targets for training, where a piano roll is a matrix containing a frame and a piano note dimension. Elements in a piano roll are assigned values of 1 or 0 indicating the presence or absence of onsets, offsets or frames in the piano roll. However, previous works such as [24] use a delta function as targets of onsets and offsets, where only one frame of an onset or offset is assigned a value of 1, with other frames assigned values of 0. There are several problems using this kind of target for piano transcription. First, the analysis from [15] shows that the attack of a piano note can last for several frames instead of only one frame. A note can be modeled in a more natural way with attack, decay, sustain and release (ADSR) states [30]. Second, the targets in [24] are sensitive to the misalignment between audio recordings and labels. For example, if an onset is misaligned by one or several frames, the training target [24] will be completely changed. Third, there are ambiguity when assigning labels for onsets and offsets events. For example, the offset time of a note can be ambiguous due to the reverb and fade-out effect.

In addition, previous piano transcription systems [24], [25] predict the presence or absence of onsets and offsets in frame-wise. So the transcription resolution is limited to the frame hop size. For example, a frame hop size of 32 ms is applied in [24],

Q. Kong, B. Li, X. Song, Y. Wan, Y. Wang are with ByteDance. (e-mail: kongqiuqiang@bytedance.com; bochenli@bytedance.com; xuchen.song@bytedance.com, wanyuan.0626@bytedance.com, wanyuxuan.11@bytedance.com). (Qiuqiang Kong is first author.)

so the system in [24] can not achieve higher transcription resolution than 32 ms. Furthermore, previous piano transcription systems [24], [25] do not contain sustain pedal transcription, which is an important part of piano transcription, and are frequently used by pianists. Liang et al. applied convolutional neural networks to detect sustain pedals [31], [32]. However, there is a lack of research on building piano transcription systems with both note and sustain pedal transcription.

To address the above problems, we propose a novel regression based high-resolution piano transcription system that can achieve arbitrary transcription resolution. In training, we propose continuous targets representing the time difference between the centre of a frame and its nearest onsets and offsets. This addresses the resolution loss problem caused by quantizing labels into piano rolls [24]. We show that our proposed continuous targets are robust to the misaligned onsets and offsets labels. We build both high-resolution note transcription and pedal transcription systems with convolutional recurrent neural networks. In inference, we propose an analytical algorithm to calculate the precise onsets and offsets times with arbitrary resolution. In evaluation, we investigate tolerances ranging from 10 ms to 20 ms for evaluation instead of using a fixed tolerance of 50 ms [24]. We show that our proposed high-resolution piano transcription system achieves state-of-the-art result on the MAESTRO dataset [25]. As far as we know, we are the first to evaluate sustain pedal transcription on the MAESTRO dataset.

This paper is organized as follows. Section II introduces neural network based piano transcription systems. Section III introduces our proposed high-resolution piano transcription system. Section IV introduces our proposed sustain pedal transcription system. Section V shows experimental results. Section VI concludes this work.

II. NEURAL NETWORK BASED PIANO TRANSCRIPTION

A. Frame-wise transcription systems

Neural networks have been used for piano transcription in previous works [20], [21], [22], [24]. To begin with, audio recordings are transformed into log mel spectrograms as input features. Then, neural networks, such as fully connected neural networks, CNNs or RNNs are applied on the log mel spectrograms to predict the frame-wise prediction of piano notes. Usually, piano rolls are used as targets for training [21]. A piano roll is a matrix with a shape of the number of frames by the number of piano notes. We denote the input log mel spectrogram with a shape of $T \times F$ as X , where T is the number of frames, and F is the number of mel frequency bins. We denote the frame-wise target roll as I_{fr} with a shape of $T \times K$, where K is the number of piano notes, which is 88 in a piano. For polyphonic piano transcription, There can be several notes active in the same time. The elements of I_{fr} have values of 1 or 0, representing the presence or absence of piano notes. Neural network based methods [21] use a function f to map the log mel spectrogram of an audio recording to the frame-wise roll I_{fr} . We denote the predicted output with a shape of $T \times K$ as $P_{fr} = f(X)$. The function f is modeled by a neural network with a set of learnable parameters. The

following loss function was used to train the neural network [21]:

$$l_{fr} = \sum_{t=1}^T \sum_{k=1}^K l_{bce}(I_{fr}(t, k), P_{fr}(t, k)), \quad (1)$$

where l_{bce} is a binary cross-entropy defined as:

$$l_{bce}(y, p) = -y \ln p - (1 - y) \ln(1 - p), \quad (2)$$

where $y \in \{0, 1\}$ is a binarized target, and $y \in [0, 1]$ is a predicted probability. In (1), the neural network based methods [21] built a classification system to output the frame-wise predictions of piano notes.

In inference, the log mel spectrogram of an audio recording is calculated. Then, the log mel spectrogram is input to the trained neural network to calculate $f(X)$. Then, those predictions are post-processed to piano note events [21].

B. Onsets and frames transcription systems

One problem of the frame-wise transcription systems is that, the transcribed frames need to be elaborately post-processed to piano note events [21]. In addition, those frame-wise piano transcription systems [21] did not use onset targets which carry rich information of piano notes. To address this problem, an onsets and offsets dual objective system [24] was proposed to jointly predict onsets and frames. The onsets and frames predictions are modeled by individual acoustic models containing several convolutional layers and long short term memory (LSTM) layers. The predicted onsets are used to condition the frame-wise predictions. We denote the predicted onsets and frames with shapes of $T \times K$ as P_{on} and P_{fr} respectively, and denote the target onsets and frames with shapes of $T \times K$ as I_{on} and I_{fr} respectively. In [24], a joint frame and onset loss function was used for training the onsets and frames system:

$$l_{note} = l_{on} + l_{fr}, \quad (3)$$

where l_{fr} is the frame-wise loss defined in (1), and l_{on} is the onset loss defined as:

$$l_{on} = \sum_{t=1}^T \sum_{k=1}^K l_{bce}(I_{on}(t, k), P_{on}(t, k)). \quad (4)$$

One advantage of the onsets and frames system is that the onsets prediction can be used as extra information for frame-wise classification. The onsets and frames system have been a benchmark system for piano transcription.

III. HIGH-RESOLUTION PIANO TRANSCRIPTION

Previous piano transcription systems introduced in Section II have several limitations. First, the methods in Section II split audio recordings into frames, and predict the presence or absence of onsets and frames in each frame. The onsets and offsets predictions are calculated on the quantized frames, so the transcription resolution of those methods are limited to the frame hop size. For example, the system [24] applies a hop size of 32 ms, so the transcription resolution is limited 32 ms. Second, for each piano note, previous systems [24] only label one frame of an onset or offset as 1, with other frames labelled as 0. The first row of Fig. 1 shows the targets of onsets

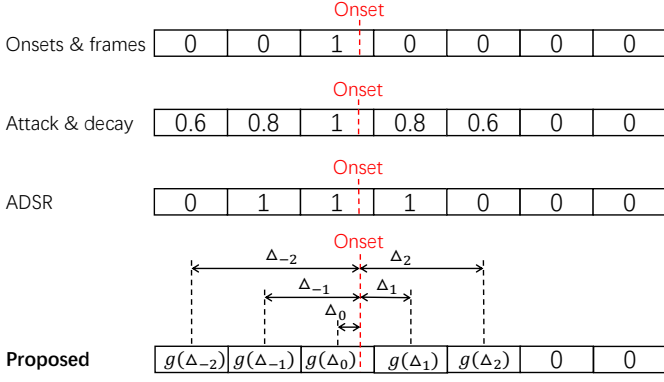


Fig. 1. Training targets of previous and our proposed piano transcription system.

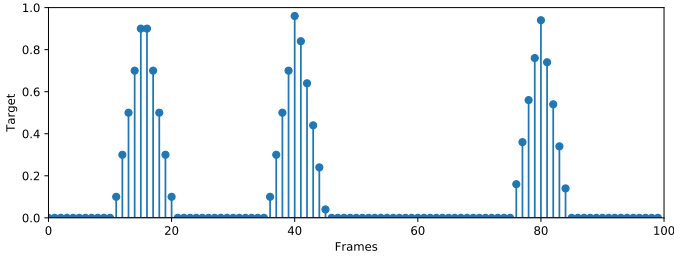


Fig. 2. High-resolution training targets of a piano note. There are three notes in this example.

used in [24], and is one column of I_{on} . The dashed vertical line shows the precise onset time of a note. The onsets and frames system [24] only assign one positive value to several consecutive frames indicating the onset of a piano note. This can be imprecise when an attack is longer than a frame. For example, a piano note was proposed to be modeled by attack, decay, sustain and release (ADSR) states [30]. In addition, the targets shown in the first row of Fig. 1 are sensitive to the misalignment of onset or offset labels. For example, shifting the onset by one or several frames will lead to a completely different target shown in the first row of Fig. 1. To address this problem, Cheng et al. [15] proposed attack and decay targets to model onsets of piano notes shown in the second row of Fig. 1. Instead of only label the frames containing onsets as 1, the neighbouring frames of onsets are also labelled with continuous values. This is to model that an attack or decay may last for several frames. In [30], ADSR states are used to model the onsets and offsets. That is, several neighbouring frames of an onset are labelled as 1, as shown in the third row of Fig. 1. In addition, the targets shown in the first to the third rows of Fig. 1 do not contain precise onsets time information. Those targets remain unchanged if the precise onsets times are shifted within a frame. Therefore, the transcription resolution of those methods are limited.

A. Regress onsets and offsets times

We propose a high-resolution piano transcription system by predicting the continuous onsets and offsets times of piano notes. Instead of classifying the presence probabilities of onsets and offsets in each frame. This idea is inspired by

the *you look only once* (YOLO) [2] object detection method from computer vision. In YOLO, an image is split into grids. Then, each grid predicts a distance between the coordinate of the grid and the coordinate of an object. The distance to be predicted is a continuous value. Similarly, we propose a high-resolution piano transcription system as follows. We predict the time distance between the centre of a frame to the precise onset or offset time of a note. The bottom row of Fig. 1 shows the targets of our proposed high-resolution piano transcription system. We denote the frame hop size time as Δ , and the time difference between the centre of a frame and the onset time as Δ_i , where i is the index of a frame. Negative i and positive i indicate the past and future frame indexes of onsets and offsets respectively. Different from the targets of previous works [24], [15], [30] shown in the first to the third rows of Fig. 1, our proposed time difference Δ_i is able to capture precise onsets and offsets information with arbitrary resolution. In training, we encode the time difference Δ_i to targets $g(\Delta_i)$ by a function g :

$$\begin{cases} g(\Delta_i) = 1 - \frac{|\Delta_i|}{J\Delta}, & |i| \leq J \\ g(\Delta_i) = 0, & |i| > J, \end{cases} \quad (5)$$

where J is a hyper-parameter controlling the sharpness of the targets. Larger J indicates “smoother” target, and smaller J indicates “sharper” target. For example, when $J = 1$ and $\Delta_0 = 0$, then (5) is equal to the binarized target [24] shown in the first row of Fig. 1. Fig. 2 shows the visualization of onsets targets of a pitch out of 88 pitches in a piano with $J = 5$. There are three piano notes in Fig. 2. Different from the attack and decay targets [15] shown in the second row of Fig. 1, the targets $g(\Delta_i)$ in Fig. 2 contain precise onsets times information of piano notes. For example, by analysing the values of targets, we know that the onset of the first note in Fig. 2 is on the boarder of two adjacent frames, and the onsets of the second and third notes in Fig. 2 are close to the centre of the frames. In training, both onsets and offsets regression targets are matrices with shapes of $T \times K$. We denote onsets and offsets regression targets as G_{on} and G_{off} respectively, to distinguish them from the binarized targets I_{on} and I_{off} in Section II. We denote the predicted onsets and offsets regression values as R_{on} and R_{off} respectively, to distinguish them from P_{on} and P_{off} in Section II. All of regression based targets regression outputs have values between $[0, 1]$. We define the onset regression loss l_{on} and offset regression loss l_{off} are defined as:

$$l_{\text{on}} = \sum_{t=1}^T \sum_{k=1}^K l_{\text{bce}}(G_{\text{on}}(t, k), R_{\text{on}}(t, k)) \quad (6)$$

$$l_{\text{off}} = \sum_{t=1}^T \sum_{k=1}^K l_{\text{bce}}(G_{\text{off}}(t, k), R_{\text{off}}(t, k)). \quad (7)$$

Equation (6), and (7) are minimized when G_{on} equals R_{on} and G_{off} equals R_{off} . Because both prediction and ground truth values are between 0 and 1, so we can use binary cross-entropy loss in (6) and (7). The difference between (6) and (4) is that (6) applies regression targets instead of binarized targets.

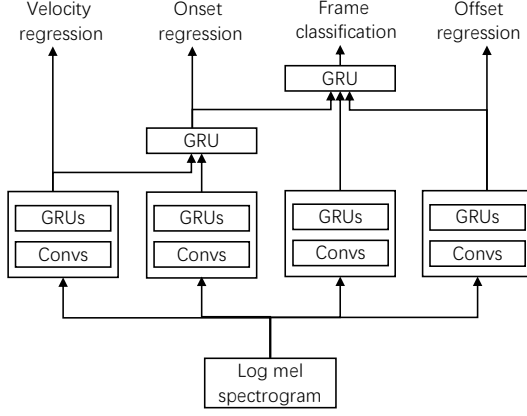


Fig. 3. High-resolution piano transcription system by regressing velocities, onsets, offsets and classifying frames.

B. Velocity estimation

Velocity is the loudness of a pitch being played. We build a velocity estimation submodule to estimate the velocity of each transcribed note. MIDI files represent velocities of notes using integers ranging from 0 to 128. Larger integers indicate louder notes, and smaller integers indicate quieter notes. To begin with, we normalize the dynamic range of velocities from $[0, 128)$ to $[0, 1)$. We denote the ground truth and prediction of velocities with shapes of $T \times K$ as I_{vel} and P_{vel} respectively. Then, we define the velocity loss as:

$$l_{\text{vel}} = \sum_{t=1}^T \sum_{k=1}^K I_{\text{on}}(t, k) \cdot l_{\text{bce}}(I_{\text{vel}}(t, k), P_{\text{vel}}(t, k)). \quad (8)$$

Equation (8) shows that the ground truth onsets $I_{\text{on}}(t, k)$ are used to modulate the velocity prediction. That is, we only predict velocities for piano notes and frames containing onsets. The reason is that, the onsets of piano notes carry rich information of their velocities, while the decay of the piano notes carry less information of velocities than onsets. Similar to (6) and (7), binary crossentropy is used to optimize (8), which is minimized when $P_{\text{vel}}(t, k)$ and $I_{\text{vel}}(t, k)$ are equal. In inference, we only predict velocities where onsets are detected. Finally, the predicted velocities are scaled from $[0, 1)$ back to integers of $[0, 128)$.

C. Entire system

Fig. 3 shows the framework of our proposed high-resolution piano transcription system. To begin with, an audio clip is transformed into a log mel spectrogram with a shape of $T \times F$ as input feature [24], where F is the number of mel frequency bins. There are four submodules in Fig. 3 from left to right: a velocity regression submodule, an onset regression submodule, a frame-wise classification submodule and an offset regression submodule. Each submodule is modeled by an acoustic model. In our system, we model each acoustic model with several convolutional layers followed by bidirectional gated recurrent units (biGRU) layers. The convolutional layers are used to extract high-level information from the log mel spectrogram, and the biGRU layers are used to summarize long time information

of the log mel spectrogram. Then, a time distributed fully connected layer is applied after the biGRU layer to predict the regression or classification result for each pitch along the time axis. The output of all acoustic models have dimensions of $T \times K$. We will describe the detailed configuration of the acoustic models in Section V-C.

Fig. 3 shows that the prediction of velocities are used to condition the prediction of onsets. This is because the detection of onsets and velocities can affect each other. For example, the velocity information of a piano note can be helpful to detect its corresponding onset. We concatenate the outputs of the velocity regression submodule and the onset regression submodule along their piano note dimension, and use this concatenation as input to a biGRU layer to calculate the final onset predictions. Similarly, we concatenate the outputs of the onset regression and offset regression submodules, and use this concatenation as the input to a biGRU layer to calculate the final frame-wise predictions. The total loss function for training consists of four parts:

$$l_{\text{note}} = l_{\text{fr}} + l_{\text{on}} + l_{\text{off}} + l_{\text{vel}}, \quad (9)$$

where l_{fr} , l_{on} , l_{off} and l_{velocity} are described in (1), (6), (7) and (8) respectively.

D. Inference

In inference, we input the log mel spectrogram of an audio recording to our trained piano transcription to calculate the frame-wise prediction, onset regression, offset regression and velocity regression outputs. Then, we propose an algorithm to process those outputs to high-resolution note events, where the note events can be represented by quadruples of $\langle \text{piano note}, \text{onset time}, \text{offset time}, \text{velocity} \rangle$.

An example of onset regression output is shown in Fig. 2. However, those outputs are calculated on quantized frames with resolution limited to frame hop size. We propose to process those outputs to high-resolution transcription results. First, we detect the onset regression outputs with local maximum values shown in Fig. 2. If a local maximum is larger than an onset threshold, then we say there exist an onset or offset near the frame. Next, we analytically calculate the precise onset or offset times of the piano note. For each frame with a local maximum value, we take its past and future frame to constitute a triple of three frames shown in Fig. 4. We denote the coordinate of three frames as A , B and C . The horizontal and vertical axes represent time and predicted outputs respectively. The point B is the network output of the central frame, and A and C are the surrounding frames of B . We propose that the horizontal coordinate of G shown in Fig. 4 is the precise onset time to be estimated. The point G satisfies AG and CG are symmetric along the vertical line GI . We calculate the horizontal coordinate of G as follows.

Without loss of generality, we assume the output value of C is larger than A . We denote the coordinate of A , B and C as (x_A, y_A) , (x_B, y_B) and (x_C, y_C) respectively. We show that the time difference between B and G is:

$$BH = \frac{x_B - x_A}{2} \frac{y_C - y_A}{y_B - y_A} \quad (10)$$

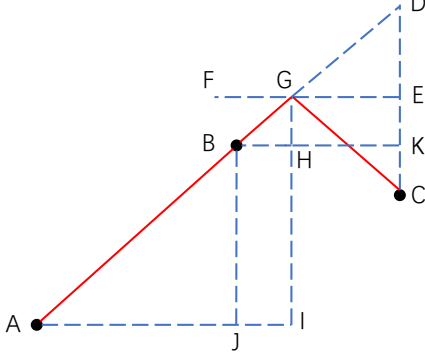


Fig. 4. Demonstration of calculating the precise onset or offset time of a note. The points A , B and C are predicted outputs of three frames. The point G is the calculated precise onset or offset time.

Proof. Extend AB to D where CD is a vertical line. Take the median point of CD as E . Draw a horizontal line EF cross AD at G . Then, we calculate BH . We have $\triangle BGH \sim \triangle ABJ$, so $BH = \frac{AJ \cdot GH}{BJ}$. We know that $AJ = x_B - x_A$, $BJ = y_B - y_A$ and $GH = DK - DE$, and we know $AJ = BK$, so $\triangle ABJ \cong \triangle BDK$, so $DK = y_B - y_A$. Then, we have $DE = \frac{CD}{2} = \frac{DK + CK}{2} = y_B - \frac{y_A + y_C}{2}$, so $GH = \frac{y_C - y_A}{2}$, so $BH = \frac{x_B - x_A}{2} \frac{y_C - y_A}{y_B - y_A}$. \square

In another case, if the output value of A is larger than C , then:

$$BH = \frac{x_C - x_B}{2} \frac{y_A - y_C}{y_B - y_C} \quad (11)$$

By this means, we are able to calculate the precise onsets and offsets times of piano notes. We describe the piano note transcription algorithm in Algorithm 1. The precise onsets are detected if they are over an onset threshold θ_{on} and is a local maximum value, and are refined to precise onset times by (10) or (11). The velocities of onsets are obtained by scaling the velocity regression back to a range of $[0, 128)$. For each detected onset, an offset is detected if the offset regression is over an offset threshold θ_{off} , or the frame prediction is lower than a frame threshold θ_{fr} .

IV. SUSTAIN PEDAL TRANSCRIPTION

Sustain pedals are important parts of a piano. When pressed, the sustain pedal sustains all the damped strings on a piano by moving all the dampers away from the strings, and allowing them to vibrate freely. All notes played will continue to sound until the pedal is released. However, many previous piano transcription systems [24], [15], [30] did not incorporate sustain pedal transcription into their piano transcription systems, and sustain-pedal transcription systems [31] did not transcribe piano notes. In [31], a convolutional neural network was used to detect piano pedals in frame-wise. However, there is a lack of benchmark sustain pedal transcription system on the MAESTRO dataset [25].

In this section, we propose a sustain pedal transcription system using our high-resolution transcription system. In the MIDI format, sustain pedals are represented with integer values ranging from 0 to 128. To simplify the sustain pedal

Algorithm 1 Piano notes transcription.

```

1: Inputs:  $R_{on}(t, k)$ ,  $R_{off}(t, k)$ ,  $P_{fr}(t)$ ,  $P_{vel}(t, k)$ ,  $\theta_{on}$ ,  $\theta_{off}$  and  $\theta_{fr}$ .
2: Outputs: Transcribed piano notes.
3: for  $k = 1, \dots, K$  do
4:   for  $t = 1, \dots, T$  do
5:     # Detect note onset.
6:     if  $R_{on}(t, k) > \theta_{on}$  and  $R_{on}(t, k)$  is local maximum
7:       then
8:         Note onset of pitch  $k$  is detected. The precise
9:         onset time is refined by (10) or (11).
10:        Calculate velocity of note by  $P_{vel}(t, k) \times 128$ .
11:      end if
12:     # Detect note offset.
13:     if ( $R_{off}(t, k) > \theta_{off}$  and  $R_{off}(t, k)$  is local maximum) or  $P_{fr}(t, k) < \theta_{fr}$  then
14:       Note offset of pitch  $k$  is detected. The precise
15:       offset time is refined by (10) or (11).
16:     end if
17:   end for
18: end for

```

transcription problem, we only classify the “on” and “off” states of sustain pedals. That is, MIDI values larger than 64 are regarded as “on”, and MIDI values smaller than 64 are regarded as “off”. We do not consider advanced sustain pedal techniques such as half pedals. We denote the pedal onset regression target, offset regression target and frame-wise target as $G_{ped_on} \in [0, 1]^T$, $G_{ped_off} \in [0, 1]^T$ and $I_{ped_fr} \in \{0, 1\}^T$ respectively. The onset and offset regression targets G_{ped_on} and G_{ped_off} are obtained by (5), and have continuous values between 0 and 1. The frame-wise target has binarized values of 0 or 1. We apply CRNN based acoustic models described in Section III-C to predict the onset regression, offset regression and the frame-wise classification of pedals. We denote the predicted onsets, offsets and frame-wise values as $R_{ped_on}(t)$, $R_{ped_off}(t)$ and $P_{ped_fr}(t)$ respectively. We write the loss functions as:

$$l_{ped_on} = \sum_{t=1}^T l_{bce}(R_{ped_on}(t), P_{ped_on}(t)) \quad (12)$$

$$l_{ped_off} = \sum_{t=1}^T l_{bce}(R_{ped_off}(t), P_{ped_off}(t)) \quad (13)$$

$$l_{ped_fr} = \sum_{t=1}^T l_{bce}(I_{ped_fr}(t), R_{ped_fr}(t)). \quad (14)$$

Then, the total loss for training the pedal transcription system is:

$$l_{ped} = l_{ped_fr} + l_{ped_on} + l_{ped_off} \quad (15)$$

In inference, we propose an algorithm shown in Algorithm 2 to process the regression outputs to high-resolution pedal events. Different from piano note onsets detection, the pedal onsets times do not need to be very precise, because the press of pedals do not make any sound until a piano note is pressed. A pedal onset is detected when the frame-wise prediction

Algorithm 2 Pedal events transcription.

```

1: Inputs:  $R_{\text{off}}(t)$ ,  $R_{\text{fr}}(t)$ ,  $\theta_{\text{ped\_off}}$ ,  $\theta_{\text{ped\_fr}}$ .
2: Outputs: Transcribed piano pedals.
3: for  $t = 1, \dots, T$  do
4:   # Detect pedal onset.
5:   if  $R_{\text{ped\_fr}}(t) > \theta_{\text{ped\_on}}$  and  $R_{\text{ped\_fr}}(t) > R_{\text{ped\_fr}}(t - 1)$ 
     then
6:     Pedal onset detected.
7:   end if
8:   # Detect pedal offset.
9:   if  $R_{\text{ped\_off}}(t) > \theta_{\text{ped\_off}}$  or  $P_{\text{ped\_fr}}(t) < \theta_{\text{ped\_fr}}$  then
10:    Pedal offset detected. The precise offset time is
        refined by (10) or (11).
11:   end if
12: end for

```

$R_{\text{ped_fr}}(t)$ is over a threshold $\theta_{\text{ped_on}}$. A pedal offset is detected is the pedal offset prediction $R_{\text{ped_off}}(t)$ is lower than an offset threshold $\theta_{\text{ped_off}}$, or frame-wise prediction $P_{\text{ped_fr}}(t)$ is lower than a threshold $\theta_{\text{ped_fr}}$.

V. EXPERIMENTS

A. Dataset

We use the MAESTRO dataset V2.0.0 [25], a large-scale dataset containing paired audio recording and MIDI files to train and evaluate our proposed piano transcription system. MAESTRO dataset contains piano recordings from the International Piano-e-Competition. Pianists perform on Yamaha Disklaviers concert-quality acoustic grand pianos integrated with high-precision MIDI capture and playback system. The MAESTRO dataset contains over 200 hours of piano solo recordings. Those audio recordings and MIDI files are aligned with a time resolution of around 3 ms. Each music recording contains meta information including the composer, title and year of the performance.

B. Preprocessing

All audio recordings are converted to monophonic, and are resampled to 16 kHz following [25]. This cut off frequency covers the frequency of the highest note C₈ on a piano of 4186 Hz. We split audio recordings into 10-second clips. Then, a short time Fourier transform with a Hanning window size 2048 is used to extract spectrogram. Mel banks with 229 banks and cut off frequencies of 30 Hz and 8000 Hz are used to extract log mel spectrogram [25]. We use a hop size of 10 ms between frames following the previous work [33]. For a 10-second audio clip, its corresponding log mel spectrogram has a shape of 1001×229 , where the extra one frame comes from the “center = True” argument during feature extraction. We set the hyper-parameter $J = 5$ in (5), that is, each onset or offset will affect the regression values of $2 \times J + 1 = 11$ frames. The first row in Fig. 5 shows an example of the log mel spectrogram of a 5-second audio clip. The second and third rows show the frame-wise targets and predictions of the audio clip. The fourth and fifth row show the regression onset targets and predictions of the audio clip. The sixth and seventh

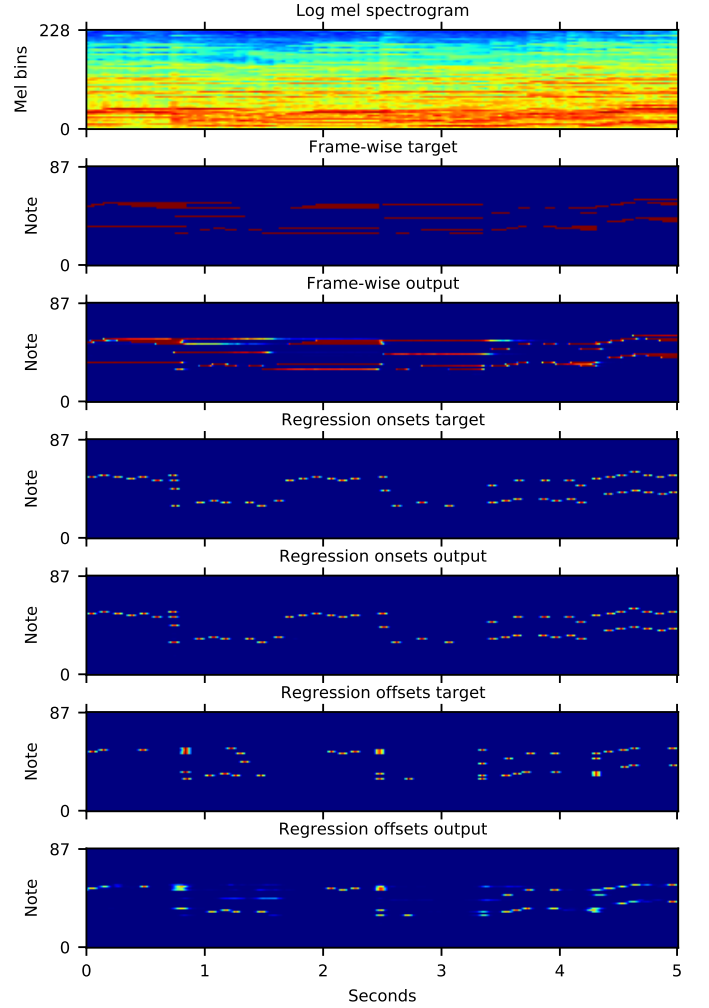


Fig. 5. From top to bottom: Log mel spectrogram of a 5-second audio clip, frame-wise targets and outputs, onset regression targets and outputs, offset regression targets and outputs

row show the regression offset targets and predictions of the audio clip.

C. Model architecture

After extracting the log mel spectrogram features, we apply a batch normalization [35] layer immediately on the individual frequency bins of the log mel spectrogram [33] to standardize the input. Then, acoustic models modeled by CRNNs are applied to predict the velocity regression, onset regression, frame-wise classification and offset regression as shown in Fig. 3. All acoustic models have the same architecture, where each acoustic model consists of four convolutional blocks and two bidirectional biGRU layers. Each convolutional block consists of two convolutional layers with kernel sizes 3×3 . Batch normalization [35] and ReLU nonlinearity [36] are applied after each linear convolutional operation to stabilize training and increase the nonlinear representation ability of the system. The four convolutional blocks have output feature maps of 48, 64, 92 and 128 respectively. After each convolutional block, feature maps are averaged pooled by a factor of 2 along the frequency axis to reduce the feature map sizes.

TABLE I
NOTE TRANSCRIPTION EVALUATED ON THE TEST SET OF MAESTRO DATASET.

	FRAME			NOTE			NOTE W/ OFFSET			NOTE W/ OFFSET & VEL.		
	P (%)	R (%)	F1 (%)	P (%)	R (%)	F1 (%)	P (%)	R (%)	F1 (%)	P (%)	R (%)	F1 (%)
Onsets & frames [25]	93.10	85.76	89.19	97.42	92.37	94.80	81.84	77.66	79.67	78.11	74.13	76.04
Adversarial onsets & frames [34]	93.10	89.80	91.40	98.10	93.20	95.60	83.50	79.30	81.30	82.30	78.20	80.20
Onsets & frames [reproduced]	86.63	90.89	88.63	99.52	89.23	93.92	80.43	72.27	75.99	79.51	71.48	75.14
Regress onsets	86.94	90.15	88.42	98.43	94.84	96.57	80.00	77.08	78.50	78.64	75.79	77.17
Regress onsets & offsets	88.91	90.28	89.51	98.53	94.81	96.61	83.81	80.70	82.20	82.36	79.33	80.79
Regress onsets & offsets & velocity	88.71	90.73	89.62	98.17	95.35	96.72	83.68	81.32	82.47	82.10	79.80	80.92
Onsets & frames (noisy labels)	80.93	90.93	85.54	65.59	93.06	76.52	44.40	63.36	51.92	40.41	57.64	47.25
Regress onsets & offsets & velocity (noisy labels)	84.65	91.36	87.79	98.65	94.30	96.39	80.59	77.09	78.77	77.35	74.02	75.62

We do not apply pooling along the time axis to remain the transcription resolution. After the convolutional layers, feature maps are flattened along the frequency and channel axes, and are input to a fully connected layer with 768 output units. Then, two biGRU layers with hidden sizes 256 are applied, followed by an additional fully connected layer with 88 sigmoid outputs. Dropout [37] with rates of 0.2 and 0.5 are applied after convolutional blocks and fully connected layers to prevent system from overfitting. Fig. 3 shows that the velocity and onset regression outputs are concatenated, and are input to a biGRU layer. The biGRU layer contains 256 hidden units, and is followed by a fully connected layer with 88 sigmoid outputs to predict the final regression onsets. Similarly, the onset regression, offset regression and frame-wise classification outputs are concatenated, and are input to a biGRU layer. The biGRU contains 256 hidden outputs, and is followed by a fully connected layer with 88 sigmoid outputs to predict the final frame-wise output. The training of piano note transcription applies the loss function (9). The pedal transcription submodules have the same acoustic model architectures as the note transcription submodules, except there are only one outputs instead of 88 outputs. The pedal onset regression, offset regression and frame-wise classification are modeled by individual acoustic models. The training of pedal transcription applies the loss function (15).

We use a batch size 12, and an Adam [38] optimizer with a learning rate of 0.0005 for training. The learning rate is reduced by a factor of 0.9 every 10k iterations in training. Systems are trained for 200k iterations. The training takes 4 days on a single Tesla-V100-PCIE-32GB GPU card. In inference, we set onset, offset, frame-wise and pedal thresholds to 0.3. The outputs are post-processed to MIDI events described in III-D.

D. Evaluation

We evaluate our proposed piano transcription system on the test set of MAESTRO dataset. The four types of evaluation including frame-wise evaluation, note evaluation with onset, note evaluation with onset and offset, and note evaluation with onset, offset and velocity. A tolerance of 50 ms is used for onset evaluation. A tolerance of 50 ms and an offset ratio of 0.2 are used for offset evaluation. A velocity tolerance of 0.1 is used for velocity evaluation. That is, estimated notes are considered correct if, after scaling and normalization to

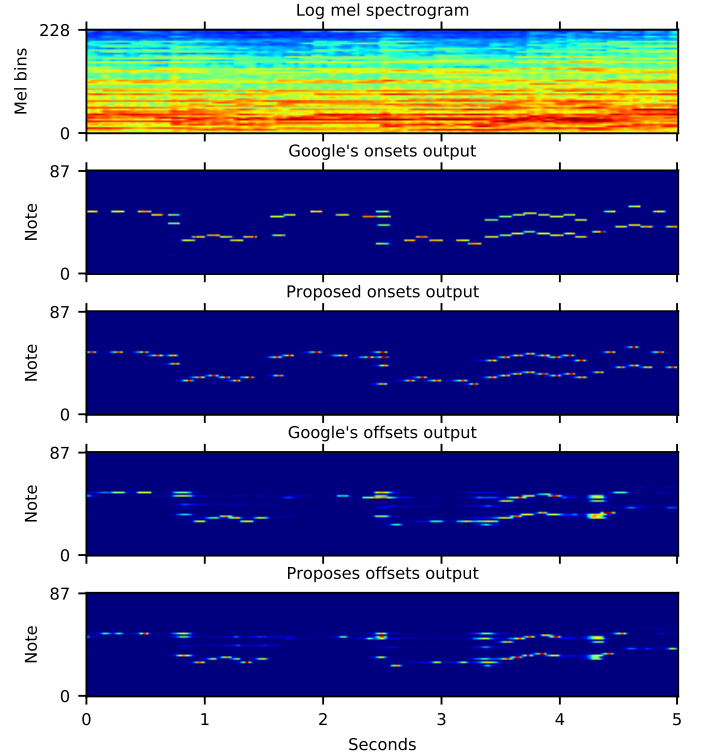


Fig. 6. Log mel spectrogram of a 5-second audio clip, Google's onset output, proposed onset output, Google's offset output, proposed offset output.

a range of 0 to 1, they are within velocity tolerance of a matched reference note. The first and second rows of Table I show the results of onsets and offsets systems [25], [34]. The third row shows our re-implemented onsets and offsets system for fair comparison with our proposed high-resolution piano transcription system. The fourth row shows that our regression based system without using velocity as condition and without offset regression improves the onsets and frames system note F1 score from 94.80% to 96.57%. The fifth row shows that our regression based system achieves an onset F1 score of 96.61%. The sixth row of Table I shows that using velocity regression as condition to onsets prediction further improves the note F1 score to 96.72%. Our proposed high-resolution system improves the note F1 score evaluated with offsets from 79.67% to 82.47%, and improves the note F1 score evaluated with offset and velocity from 76.04% to 80.92%. The first

TABLE II
ONSETS EVALUATION WITH DIFFERENT ONSET TOLERANCES

	ONSETS & FRAMES			PROPOSED		
	P (%)	R (%)	F1 (%)	P (%)	R (%)	F1 (%)
100 ms	99.54	89.24	93.94	98.24	95.42	96.79
50 ms	99.52	89.23	93.92	89.17	95.35	96.72
20 ms	99.14	88.92	93.58	97.22	94.45	95.79
10 ms	91.74	92.53	86.73	90.19	87.69	88.91
5 ms	63.89	57.69	60.53	66.22	64.47	65.32
2 ms	28.69	25.93	27.19	31.49	30.68	31.08

TABLE III
ONSETS AND OFFSETS EVALUATION WITH DIFFERENT OFFSET TOLERANCES

	ONSETS & FRAMES			REGRESS & ON. & OFF.		
	P (%)	R (%)	F1 (%)	P (%)	R (%)	F1 (%)
500 ms	93.71	84.13	88.49	93.25	90.59	91.88
200 ms	87.65	78.75	82.81	88.19	85.68	86.90
100 ms	82.80	74.41	78.24	84.48	82.08	83.25
50 ms	75.47	67.76	71.28	80.24	77.95	79.06
20 ms	51.66	46.19	48.68	69.96	67.97	68.94
10 ms	31.64	28.14	29.73	52.04	50.57	51.28

row of Fig. 5 shows the log mel spectrogram of an audio clip. The second and third rows show the frame-wise target and frame-wise system output. The fourth and fifth rows show the regression onset target and regression onset output. The sixth and seventh row show the regression offset target and regression offset output. The onsets and offsets regression targets are calculated by (5). Fig. 5 shows that our proposed system performs well on transcribing this 5-second audio clip.

To show that our proposed regression based piano transcription system is robust to the misalignment of labels, we randomly shift the onset and offset labels of notes in the time domain with a uniform distribution between -50 ms and +50 ms. The seventh row of Table I shows that the note F1 score of [24] decreases from 93.92% to 76.52% when trained with misaligned labels. One explanation is that, the system [24] trained with misaligned data is difficult to detect the precise onset or offset times. On the other hand, the bottom row of Table I shows that our proposed regression based system achieves a note F1 of 96.39%, compared to training with correct labels of 96.72%, and achieves an F1 of 75.62% when evaluated with offsets and velocities, compared to training with correct labels of 80.92%. These results indicate that our proposed system is robust to misalignment labels. The first row of Fig. 6 shows the log mel spectrogram of the same audio clip as in Fig. 5. The second and third row show the onset output of [24] and our proposed high-resolution system trained with misaligned labels. The fourth and fifth row show the offset output of [24] and our system. The outputs of our system contain information of precise onset and offset times, which can be calculated by (10) or (11). However, the onset and offset outputs of [24] are blurred shown in the second row and fourth row of Fig. 5.

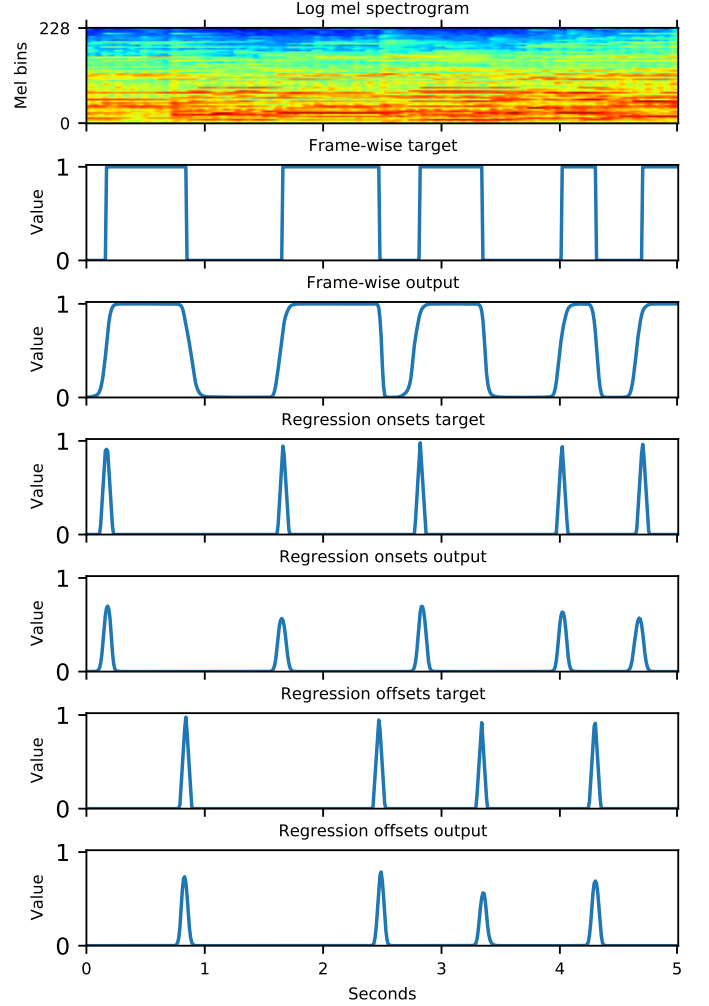


Fig. 7. Log mel spectrogram of a 5-second audio clip, onset target and output, offset target and output, frame target and output.

We investigate evaluating piano transcription performance with different onsets and offsets tolerance that has not been investigated in previous works [24], [25], [34]. In this section, we evaluate piano transcription with different onset and offset tolerances ranging from 10 ms to 500 ms. Table II shows that with an onset tolerance of 2 ms, our system achieves an onset F1 score of 31.08%. The F1 score increases to 88.91% when onset tolerance is 10 ms, and increases to 96.79% when onset tolerance is 100 ms. The F1 scores of our proposed system outperform the onsets and offsets system [24] in all onset tolerances. Table III shows the note F1 score evaluated with offset tolerances ranging from 10 ms to 500 ms with onset tolerance fixed to 50 ms. Our system achieves a note F1 score of 51.28% with an offset tolerance of 10 ms. The F1 score increases to 91.88% with a tolerance of 500 ms. Our proposed system outperforms the onsets and frames system [24] in all offset tolerances. These experiments shows that our proposed system can achieve higher transcription resolution than [24], [25].

We evaluate pedal transcription using our proposed high-resolution pedal transcription system. Table IV shows the pedal transcription result. Our system achieves an event based

TABLE IV
PEDAL TRANSCRIPTION EVALUATED ON THE TEST SET OF MAESTRO DATASET.

	FRAME			EVENT			EVENT W/ OFFSET		
	P (%)	R (%)	F1 (%)	P (%)	R (%)	F1 (%)	P (%)	R (%)	F1 (%)
Onsets & frames [reproduced]	94.30	94.42	94.25	93.20	90.26	91.57	86.94	84.28	85.47
Proposed	94.30	94.42	94.25	91.59	92.41	91.86	86.36	87.02	86.58
Onsets & frames (noisy labels)	93.62	94.14	93.77	92.71	85.48	88.69	83.17	77.03	79.78
proposed (noisy labels)	94.41	93.29	93.73	91.62	91.17	91.23	86.33	85.83	85.94

F1 of 91.86% evaluated with pedal onset tolerance of 50 ms, and achieves an event based F1 of 86.58% evaluated with both onset and offset tolerances of 50 ms and offset ratio of 0.2, outperforming our implemented onsets and frames system of 91.57% and 85.47%. As far as we know, we are the first to evaluate piano pedal transcription on the MAESTRO dataset. The first row of Fig. 7 shows log mel spectrogram of the audio clip that is the same in 5. The second and third rows show the frame-wise pedal target and system output. Values close to 1 indicate “on” states, and values close to 0 indicate “off” states. The fourth and fifth rows show the regression onset targets and outputs. The sixth and seventh rows show the regression offset targets and outputs. Fig. 7 shows that our pedal transcription system performs well on the 5-second audio clip example.

VI. CONCLUSION

We propose a high-resolution piano transcription system by regressing the precise onsets and offsets times of piano notes and pedals. In inference, we propose an analytical algorithm to calculate the precise onsets and offsets times. We show that our proposed system achieves state-of-the-art onset F1 score of 96.72% in piano note transcription, outperforming the onsets and frames system of 94.80%. We show that our system is robust to the misalignment of onsets and offsets labels. In addition, we investigate evaluating piano transcription systems with different onset and offset tolerances that are not evaluated in previous works. As far as we know, we are the first to evaluate pedal transcription on the MAESTRO dataset, and achieves a pedal event F1 score of 91.86%. In future, we will extend the high-resolution piano transcription system to transcribe other instruments.

REFERENCES

- [1] C. Raphael, “Automatic transcription of piano music,” in *International Society for Music Information Retrieval (ISMIR)*, 2002.
- [2] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You only look once: Unified, real-time object detection,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 779–788.
- [3] 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, 2018.
- [4] E. Benetos, S. Dixon, D. Giannoulis, H. Kirchhoff, and A. Klapuri, “Automatic music transcription: challenges and future directions,” *Journal of Intelligent Information Systems*, vol. 41, no. 3, pp. 407–434, 2013.
- [5] B. Li and Z. Duan, “An approach to score following for piano performances with the sustained effect,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 12, pp. 2425–2438, 2016.
- [6] B. Niedermayer and G. Widmer, “A multi-pass algorithm for accurate audio-to-score alignment,” in *ISMIR*. Citeseer, 2010, pp. 417–422.
- [7] Z. Duan and B. Pardo, “Soundprism: An online system for score-informed source separation of music audio,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 6, pp. 1205–1215, 2011.
- [8] E. Scheirer, “Using musical knowledge to extract expressive performance information from audio recordings,” in *IJCAI Workshop on Computational Auditory Scene Analysis*, 1995, pp. 153–160.
- [9] S. Dixon, “On the computer recognition of solo piano music,” in *Proceedings of Australasian Computer Music Conference*, 2000, pp. 31–37.
- [10] M. Marolt, “Transcription of polyphonic piano music with neural networks,” in *Mediterranean Electrotechnical Conference. Information Technology and Electrotechnology for the Mediterranean Countries*, vol. 2, 2000, pp. 512–515.
- [11] G. E. Poliner and D. P. W. Ellis, “A discriminative model for polyphonic piano transcription,” *EURASIP Journal on Advances in Signal Processing*, 2006.
- [12] V. Emiya, R. Badeau, and B. David, “Multipitch estimation of piano sounds using a new probabilistic spectral smoothness principle,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 6, pp. 1643–1654, 2009.
- [13] J. P. Bello, L. Daudet, and M. Sandler, “Automatic piano transcription using frequency and time-domain information,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 6, pp. 2242–2251, 2006.
- [14] K. O’Hanlon and M. D. Plumbley, “Polyphonic piano transcription using non-negative matrix factorisation with group sparsity,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2014, pp. 3112–3116.
- [15] T. Cheng, M. Mauch, E. Benetos, and S. Dixon, “An attack/decay model for piano transcription,” in *International Society for Music Information Retrieval (ISMIR)*, 2016.
- [16] T. Berg-Kirkpatrick, J. Andreas, and D. Klein, “Unsupervised transcription of piano music,” in *Advances in Neural Information Processing Systems*, 2014, pp. 1538–1546.
- [17] M. Marolt, “A connectionist approach to automatic transcription of polyphonic piano music,” *IEEE Transactions on Multimedia*, vol. 6, no. 3, pp. 439–449, 2004.
- [18] A. Cogliati, Z. Duan, and B. Wohlberg, “Piano music transcription with fast convolutional sparse coding,” in *2015 IEEE 25th International Workshop on Machine Learning for Signal Processing (MLSP)*, 2015, pp. 1–6.
- [19] J. Nam, J. Ngiam, H. Lee, and M. Slaney, “A classification-based polyphonic piano transcription approach using learned feature representations,” in *International Society for Music Information Retrieval (ISMIR)*, 2011, pp. 175–180.
- [20] S. Böck and M. Schedl, “Polyphonic piano note transcription with recurrent neural networks,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2012, pp. 121–124.
- [21] S. Sigtia, E. Benetos, and S. Dixon, “An end-to-end neural network for polyphonic piano music transcription,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 5, pp. 927–939, 2016.
- [22] R. Kelz, M. Dorfer, F. Korzeniowski, S. Böck, A. Arzt, and G. Widmer, “On the potential of simple framewise approaches to piano transcription,” in *International Society for Music Information Retrieval (ISMIR)*, 2016.
- [23] R. Kelz, S. Böck, and C. Widnaer, “Multitask learning for polyphonic piano transcription, a case study,” in *2019 International Workshop on Multilayer Music Representation and Processing (MMRP)*. IEEE, 2019, pp. 85–91.
- [24] 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 *International Society for Music Information Retrieval Conference (ISMIR)*, 2018.
- [25] C. Hawthorne, A. Stasyuk, A. Roberts, I. Simon, C. A. Huang, S. Dieleman, E. Elsen, J. Engel, and D. Eck, “Enabling factorized piano music modeling and generation with the maestro dataset,” *International Conference on Learning Representations (ICLR)*, 2018.
 - [26] J. Schroeter, K. Sidorov, and D. Marshall, “Softloc: Robust temporal localization under label misalignment,” *Open Review at International Conference on Learning Representations (ICLR)*, 2019.
 - [27] A. Elowsson, “Polyphonic pitch tracking with deep layered learning,” *The Journal of the Acoustical Society of America*, vol. 148, no. 1, pp. 446–468, 2020.
 - [28] A. Elowsson, “Modeling music: studies of music transcription, music perception and music production,” *KTH Royal Institute of Technology*, 2018.
 - [29] A. Gkiokas and V. Katsouros, “Convolutional neural networks for real-time beat tracking: A dancing robot application,” in *International Society for Music Information Retrieval (ISMIR)*, 2017, pp. 286–293.
 - [30] R. Kelz, S. Böck, and G. Widmer, “Deep polyphonic ADSR piano note transcription,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 246–250.
 - [31] B. Liang, G. Fazekas, and M. Sandler, “Piano sustain-pedal detection using convolutional neural networks,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 241–245.
 - [32] B. Liang, G. Fazekas, and M. B. Sandler, “Detection of piano pedaling techniques on the sustain pedal,” in *Audio Engineering Society Convention*, 2017.
 - [33] Q. Kong, Y. Cao, T. Iqbal, Y. Wang, W. Wang, and M. D. Plumbley, “PANNs: Large-Scale Pretrained Audio Neural Networks for Audio Pattern Recognition,” *arXiv preprint arXiv:1912.10211*, 2019.
 - [34] J. W. Kim and J. P. Bello, “Adversarial learning for improved onsets and frames music transcription,” in *International Society for Music Information Retrieval (ISMIR)*, 2019, pp. 670–677.
 - [35] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *International Conference on Machine Learning (ICML)*, 2015, pp. 448–456.
 - [36] V. Nair and G. E. Hinton, “Rectified linear units improve restricted boltzmann machines,” in *Proceedings of International Conference on Machine Learning (ICML)*, 2010, pp. 807–814.
 - [37] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: a simple way to prevent neural networks from overfitting,” *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
 - [38] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *International Conference on Learning Representations (ICLR)*, 2015.