

ADVERSARIAL LEARNING FOR IMPROVED ONSETS AND FRAMES MUSIC TRANSCRIPTION

Jong Wook Kim, Juan Pablo Bello

Music and Audio Research Lab, New York University

{jongwook, jpbello}@nyu.edu

ABSTRACT

Automatic music transcription is considered to be one of the hardest problems in music information retrieval, yet recent deep learning approaches have achieved substantial improvements on transcription performance. These approaches commonly employ supervised learning models that predict various time-frequency representations, by minimizing element-wise losses such as the cross entropy function. However, applying the loss in this manner assumes conditional independence of each label given the input, and thus cannot accurately express inter-label dependencies. To address this issue, we introduce an adversarial training scheme that operates directly on the time-frequency representations and makes the output distribution closer to the ground-truth. Through adversarial learning, we achieve a consistent improvement in both frame-level and note-level metrics over Onsets and Frames, a state-of-the-art music transcription model. Our results show that adversarial learning can significantly reduce the error rate while increasing the confidence of the model estimations. Our approach is generic and applicable to any transcription model based on multi-label predictions, which are very common in music signal analysis.

1. INTRODUCTION

Automatic music transcription (AMT) concerns automated methods for converting acoustic music signals into some form of musical notation [4]. AMT is a multifaceted problem and comprises a number of subtasks, including multi-pitch estimation (MPE), note tracking, instrument recognition, rhythm analysis, score typesetting, etc. MPE predicts a set of concurrent pitches that are present at each instant, and it is closely related to the task of note tracking, which predicts the onset and offset timings of every note in audio. In this paper, we address an issue in the recent approaches for MPE and note tracking, where the probabilistic dependencies between the labels are often overlooked.

A common approach for MPE and note tracking is through the prediction of a two-dimensional representation

that is defined along the time and frequency axes and contains the pitch tracks of notes over time. Piano rolls are the most common example of such representations, and deep salience [5] is another example that can contain more granular information on pitch contours. Once such representation is obtained, pitches and notes can be decoded by thresholding [23] or other heuristic methods [17, 25].

To train a model that predicts a two-dimensional target representation $\hat{\mathbf{Y}} \in \mathbb{R}^{P \times T}$ from an input audio representation \mathbf{X} , where P is the number of pitch labels and T is the number of time frames, a common approach is to minimize the element-wise sum of a loss function \mathcal{L} :

$$\text{minimize } \mathcal{L}(\hat{\mathbf{Y}}, \mathbf{Y}) = \sum_{p=1}^P \sum_{t=1}^T \mathcal{L}(\hat{\mathbf{Y}}_{pt}, \mathbf{Y}_{pt}), \quad (1)$$

where $\mathbf{Y} \in \mathbb{R}^{P \times T}$ is the ground truth. In a probabilistic perspective, we can interpret \mathcal{L} as the negative log-likelihood of the model parameters ϑ of a discriminative model $p_{\vartheta}(\mathbf{Y}|\mathbf{X})$:

$$p_{\vartheta}(\mathbf{Y}|\mathbf{X}) = e^{-\mathcal{L}(\hat{\mathbf{Y}}, \mathbf{Y})} = \prod_{p=1}^P \prod_{t=1}^T e^{-\mathcal{L}(\hat{\mathbf{Y}}_{pt}, \mathbf{Y}_{pt})} = \prod_{p=1}^P \prod_{t=1}^T p_{\vartheta}(\mathbf{Y}_{pt}|\mathbf{X}) \quad (2)$$

which indicates that each element of the label \mathbf{Y} is conditionally independent with each other given the input \mathbf{X} . This encourages the model to predict the average of the posterior, making blurry predictions when the posterior distribution is multimodal, e.g. natural images [10].

Music data is highly contextual and multimodal, and the conditional independence assumption does not hold in general. This is why many computational music analysis models employ a separate post-processing stage after sequence prediction. One approach is to factorize the joint probability using the chain rule and assume the Markov property:

$$p_{\vartheta}(\mathbf{Y}|\mathbf{X}) \approx \prod_{p=1}^P \prod_{t=1}^T p_{\vartheta}(\mathbf{Y}_{pt}|\mathbf{Y}_{\cdot(t-1)}, \mathbf{X}). \quad (3)$$

This corresponds to appending hidden Markov models (HMMs) [36] or recurrent neural networks (RNNs) [17, 39] to the transcription model. The Markov assumption is effective for one-dimensional sequence prediction tasks, such as chord estimation [35] and monophonic pitch tracking [32], but when predicting a two-dimensional representation, it still does not address the inter-label dependencies along the frequency axis.



There exist a number of models in the computer vision literature that can express inter-label dependencies in two-dimensional predictions, such as the neural autoregressive distribution estimator (NADE) [27], PixelRNN [43], and PixelCNN [42]. However, apart from a notable exception using a hybrid RNN-NADE approach [7], the effect of learning the joint posterior distribution for polyphonic music transcription has not been well studied.

To this end, we propose a new approach for effectively leveraging inter-label dependencies in polyphonic music transcription. We pose the problem as an image translation task and apply an adversarial loss incurred by a discriminator network attached to the baseline model. We show that our approach can consistently and significantly reduce the transcription errors in *Onsets and Frames* [17], a state-of-the-art music transcription model.

2. BACKGROUND

2.1 Automatic Transcription of Polyphonic Music

Automatic transcription models for polyphonic music can be classified into frame- or note-level approaches. Frame-level transcription is synonymous with multi-pitch estimation (MPE) and operates on tiny temporal slices of audio, or frames, to predict all pitch values present in each frame. Note-level transcription, or note tracking, operates at a higher level, predicting a sequence of note events that contains the pitch, the onset time, and optionally the offset time of each note. Note tracking is typically implemented as a post-processing stage on the output of MPE [3], by connecting and grouping the pitch estimates over time to produce discrete note events. In this sense, we can say that MPE is at the core of polyphonic music transcription.

Two categories of approaches for MPE have been most successful in recent years: matrix factorization and deep learning. Factorization-based models for music transcription [40] use non-negative matrix factorization (NMF) [29] to factorize a time-frequency representation $\mathbf{X} \in \mathbb{R}^{F \times T}$ as a product of a dictionary matrix $\mathbf{D} \in \mathbb{R}^{F \times K}$ and an activation matrix $\mathbf{A} \in \mathbb{R}^{K \times T}$, where K is the number of pitch labels to be transcribed, e.g. 88 keys for piano transcription. This allows for an intuitive interpretation of each matrix, where each column of \mathbf{D} contains a spectral template for a pitch label, and each row of \mathbf{A} contains the activation of the corresponding pitch over time. Various extensions of factorization-based methods have been proposed to leverage sparsity [1], adaptive estimation of harmonic spectra [14, 44], and modeling of attack and decay sounds [2, 12]. In all of these approaches, an iterative gradient-descent algorithm is used to minimize an element-wise divergence function between the matrix factorization $\mathbf{D}\mathbf{A}$ and the target matrix \mathbf{X} [13].

Deep learning [28] methods for music transcription are increasingly popular [3], as larger labeled datasets and more powerful hardware become accessible. These approaches use neural networks (NN) to produce music transcriptions from the input audio. An early work [34] used deep belief networks [20] to extract audio features which




Figure 1. The Onset and Frames model. CNN denotes the convolutional acoustic model taken from [23], FC denotes a fully connected layer, and σ denotes sigmoid activation. Dotted lines mean stop-gradient, i.e. no backpropagation.

are subsequently fed to pitch-wise SVM-HMM pairs to predict the target piano rolls. More recent approaches are based on convolutional [5, 23] and/or recurrent neural networks [6, 17, 39], which are also optimized with gradient descent to minimize an element-wise loss of predicting the target time-frequency representations.

Onsets and Frames [17] is a state-of-the-art piano transcription model that we use as our baseline. It uses multiple columns of convolutional and recurrent layers to predict onsets, offsets, velocities, and frame labels from the Mel spectrogram input, as shown in Figure 1. Predicted onset and frame posteriors are then used for decoding the note sequences, where a threshold value is used to create binary onset and frame activations, and frame activations without the corresponding onsets are disregarded.

As discussed above, most NMF- and NN-based methods, including Onsets and Frames, use an element-wise optimization objective which does not consider the inter-label dependencies. This motivates the adversarial training scheme that is outlined in the following subsection.

2.2 Generative Adversarial Networks and pix2pix

Generative adversarial networks (GANs) [16] refer to a family of deep generative models which consist of two components, namely the generator G and the discriminator D . Given a data distribution $\mathbf{x} \sim p(\mathbf{x})$ and latent codes $\mathbf{z} \sim p(\mathbf{z})$, GAN performs the following minimax game:

$$\min_G \max_D \underbrace{\left[\mathbb{E}_{\mathbf{x}} \log D(\mathbf{x}) + \mathbb{E}_{\mathbf{z}} \log(1 - D(G(\mathbf{z}))) \right]}_{\mathcal{L}_{\text{GAN}}(G, D)}. \quad (4)$$

G and D are implemented as neural networks trained in an adversarial manner, where the discriminator learns to distinguish the generated samples from the real data, while the generator learns to produce realistic samples to fool the discriminator. GANs are most renowned for their ability to produce photorealistic images [22] and have shown promising results on music generation as well [9, 11, 45]. We refer the readers to [8, 15] for a comprehensive review of the techniques, variants, and applications of GANs.

The second term in Equation 4 has near-zero gradients when $D(G(\mathbf{z})) \approx 0$, which is usually the case in early training. To avoid this, a non-saturating variant of GAN is suggested in [16] where the generator is trained with the following optimization objective instead:

$$\max_G \mathbb{E}_{\mathbf{z}} \log D(G(\mathbf{z})). \quad (5)$$

The non-saturating GAN loss is used more often than the minimax loss in Equation 4 and is implemented by flipping the labels of fake data while using the same loss function. Least-squares GAN [31] is an alternative method to address the vanishing gradient problem, which replaces the cross entropy loss in Equations 4-5 with squared errors:

$$\begin{aligned} \min_D \left[\mathbb{E}_{\mathbf{x}} (D(\mathbf{x}) - 1)^2 + \mathbb{E}_{\mathbf{z}} D(G(\mathbf{z}))^2 \right], \\ \min_G \mathbb{E}_{\mathbf{z}} (D(G(\mathbf{z})) - 1)^2. \end{aligned} \quad (6)$$

While the default formulation of GAN concerns unconditional generation of samples from $p(\mathbf{x})$, conditional GANs (cGAN) [33] produce samples from a conditional distribution $p(\mathbf{y}|\mathbf{x})$. To do this, the generator and the discriminator are defined in terms of the condition variable \mathbf{x} as well:

$$\min_G \max_D \underbrace{\left[\mathbb{E}_{\mathbf{x}, \mathbf{y}} \log D(\mathbf{x}, \mathbf{y}) + \mathbb{E}_{\mathbf{x}, \mathbf{z}} \log(1 - D(\mathbf{x}, G(\mathbf{x}, \mathbf{z}))) \right]}_{\mathcal{L}_{\text{cGAN}}(G, D)}. \quad (7)$$

`pix2pix` [21] is an image translation model that learns a mapping between two distinct domains of images, such as aerial photos and maps. A `pix2pix` model takes paired images (\mathbf{x}, \mathbf{y}) as its training data and minimizes the conditional GAN loss along with an additional L1 loss:

$$\mathbb{E}_{\mathbf{x}, \mathbf{y}, \mathbf{z}} \left\| \mathbf{y} - G(\mathbf{x}, \mathbf{z}) \right\|_1, \quad (8)$$

which encourages the conditional generator to learn the forward mapping from \mathbf{x} to \mathbf{y} . It can be thought that the GAN loss in Equation 7 is fine-tuning the mapping learned by the L1 loss in Equation 8, resulting in a predictive mapping that better respects the probabilistic dependencies within the labels \mathbf{y} .

In this paper, we adapt this approach to music transcription tasks and show that we can indeed improve the performance by introducing an adversarial loss to an existing music transcription model.

3. METHOD

We describe a general method for improving an NN-based transcription model G that performs prediction of a two-dimensional target \mathbf{Y} from an input audio representation \mathbf{X} . Say the original model G is trained by minimizing the loss $\mathcal{L}_{\text{task}}(G(\mathbf{X}), \mathbf{Y})$ between the predicted target $\hat{\mathbf{Y}} = G(\mathbf{X})$ and the ground-truth \mathbf{Y} . The main idea of our method is to adapt `pix2pix` [21] to this setup, by introducing an adversarial discriminator D during the training process. The adversarial training objective includes the conditional GAN loss $\mathcal{L}_{\text{cGAN}}$ (Equation 7):

$$\min_G \max_D \mathbb{E}_{\mathbf{X}, \mathbf{Y}} \left[\nu \mathcal{L}_{\text{task}}(G(\mathbf{X}), \mathbf{Y}) + \mathcal{L}_{\text{cGAN}}(G, D) \right], \quad (9)$$

where ν is a hyperparameter that controls how much the conditional GAN loss contributes to the gradient steps relative to the discriminative loss $\mathcal{L}_{\text{task}}$. Figure 2 illustrates how the two components are connected in the computation graph and how the loss terms are calculated.




Figure 2. A computation graph showing how a discriminator is appended to the original model. The appended parts are shown as dotted components.

Adversarial training with $\mathcal{L}_{\text{cGAN}}$ allows the model to learn the inter-label dependencies as desired, even when $\mathcal{L}_{\text{task}}$ is defined only in terms of element-wise operations between $\hat{\mathbf{Y}}$ and \mathbf{Y} , as in Equation 1. In the next subsection, we describe a neural network architecture for the cGAN discriminator that leverages prior knowledge on music.

3.1 Musically Inspired Adversarial Discriminator

Following `pix2pix`, we use a fully convolutional architecture [30] for the discriminator. By being fully convolutional, the discriminator has translation invariance not only along the time axis (as in HMMs and RNNs) but also along the frequency axis. Since the discriminator determines how realistic a polyphonic note sequence is, the translation invariance enforces that the decision does not depend on the musical key, but only on the relative pitch and time intervals between the notes. This effectively implements a music language model (MLM) [7, 39] and biases the transcription toward more realistic note sequences.

Unlike the image-to-image translation problem, the input representations (e.g. Mel spectrograms) and the output representations (e.g. piano rolls) of a music transcription model can have different dimensions. This makes combining \mathbf{X} and \mathbf{Y} in a fully convolutional manner difficult. For this reason, we make the discriminator a function of \mathbf{Y} only, simplifying the objective in Equation 7 to:

$$\mathcal{L}_{\text{cGAN}}(G, D) = \mathbb{E}_{\mathbf{Y}} \log D(\mathbf{Y}) + \mathbb{E}_{\mathbf{X}} \log(1 - D(G(\mathbf{X}))). \quad (10)$$

Note that \mathbf{z} is also omitted in Equation 10, as we follow [21] and implement the stochasticity of \mathbf{z} only in terms of dropout layers [41], without explicitly feeding random noises into the generator. This causes a mode collapse problem where the learned $p(\mathbf{Y}|\mathbf{X})$ is not diverse enough, but it does not harm our purpose of producing more realistic target representations.

3.2 TTUR and *mixup* to Stabilize GAN Training

Although an ideal GAN generator can fully reconstruct the data distribution at the global optimum [16], training of GANs in practice is notoriously difficult, especially for high-dimensional data [15]. This led to the inventions of a plethora of techniques for stabilizing GAN training, among which we employ the two-timescale update rule (TTUR) [19] and *mixup* [47]. TTUR means simply setting the generator’s learning rate a few times larger than that of the discriminator, which has been empirically shown to stabilize GAN training significantly.

The other technique, *mixup*, is an extension to empirical risk minimization where training data samples are drawn from convex interpolations between pairs of empirical data samples. For a pair of feature-target tuples $(\mathbf{X}_i, \mathbf{Y}_i)$ and $(\mathbf{X}_j, \mathbf{Y}_j)$ sampled randomly from the empirical distribution, their convex interpolation is given by:

$$\begin{aligned}\tilde{\mathbf{X}} &= \lambda \mathbf{X}_i + (1 - \lambda) \mathbf{X}_j \\ \tilde{\mathbf{Y}} &= \lambda \mathbf{Y}_i + (1 - \lambda) \mathbf{Y}_j\end{aligned}\quad (11)$$

where $\lambda \sim \text{Beta}(\alpha, \alpha)$, and α is the *mixup* hyperparameter which controls the strength of interpolation. When $\alpha = 0$, the Beta distribution becomes Bernoulli(0.5) which recovers the usual GAN training without *mixup*.

mixup is readily applicable to the binary classification task of GAN discriminators. In our conditional GAN setup, we have an additional advantage of having paired samples of a real label \mathbf{Y} and a fake label $\tilde{\mathbf{Y}} = G(\mathbf{X})$, which allow us to replace Equation 10 with:

$$\min_G \max_D \mathbb{E}_{\mathbf{X}, \mathbf{Y}, \lambda} \left[-\ell(D(\lambda \mathbf{Y} + (1 - \lambda)G(\mathbf{X})), \lambda) \right]. \quad (12)$$

where $\ell(p, y) = -y \log p - (1 - y) \log(1 - p)$ is the binary cross entropy (BCE) function. With this *mixup* setup, the discriminator now has to operate on the convex interpolation between the predicted representation and the corresponding ground truth. This makes the discriminator's task even more difficult when the prediction gets close to the ground truth, which is desirable because the discriminator should be inconclusive (i.e. $D = \frac{1}{2}$ everywhere) at the global optimum [16].

Algorithm 1 details the procedure of training the conditional GAN using *mixup*, based on Equations 10 and 12. Note that for training the generator network, we perform label flipping in $\mathcal{L}_{\text{cGAN}}^G$ similarly as in Equation 5. Also, to train a least-squares GAN (Equation 6) instead, we can simply replace ℓ with a mean squared error (MSE) loss.

Input: Generator $G_\vartheta(\mathbf{X})$ with initial parameters ϑ , learning rate η , and loss function $\mathcal{L}_{\text{task}}(\tilde{\mathbf{Y}}, \mathbf{Y})$, discriminator $D_\varphi(\mathbf{Y})$ with initial parameters φ , learning rate β , and loss function $\ell \in \{\text{BCE}, \text{MSE}\}$, batch size m , training data distribution $p(\mathbf{X}, \mathbf{Y})$, `pix2pix` weight ν , *mixup* strength α .

Output: Trained conditional generator $G_\vartheta(\mathbf{X})$.

```

while  $\varphi$  and  $\vartheta$  have not converged do
     $\{(\mathbf{X}_i, \mathbf{Y}_i)\}_{i=1, \dots, m} \leftarrow m$  samples from  $p(\mathbf{X}, \mathbf{Y})$ 
    for  $i = 1, \dots, m$  do
         $\tilde{\mathbf{Y}}_i \leftarrow G_\vartheta(\mathbf{X}_i)$ 
         $\lambda_i \leftarrow \text{sample from Beta}(\alpha, \alpha)$ 
         $\tilde{\mathbf{Y}}_i \leftarrow \lambda_i \mathbf{Y}_i + (1 - \lambda_i) \tilde{\mathbf{Y}}_i$ 
    end
     $\mathcal{L}_{\text{cGAN}}^D \leftarrow \sum_{i=1}^m \ell(D_\varphi(\tilde{\mathbf{Y}}_i), \lambda_i)$ 
     $\varphi \leftarrow \varphi - \beta \cdot \nabla_\varphi \mathcal{L}_{\text{cGAN}}^D$ 
     $\mathcal{L}_{\text{cGAN}}^G \leftarrow \sum_{i=1}^m \ell(D_\varphi(\tilde{\mathbf{Y}}_i), 1 - \lambda_i)$ 
     $\vartheta \leftarrow \vartheta - \eta \cdot \nabla_\vartheta \left[ \sum_{i=1}^m \nu \mathcal{L}_{\text{task}}(\tilde{\mathbf{Y}}_i, \mathbf{Y}_i) - \mathcal{L}_{\text{cGAN}}^G \right]$ 
end
    
```

Algorithm 1: Training of a *mixup* Conditional GAN.

4. EXPERIMENTAL SETUP

To verify the effectiveness of our approach, we compare Onsets and Frames [17], a state-of-the-art piano transcription model, with variants of the same model that are trained with the adversarial loss. We also aim to evaluate the choices of the GAN loss and the *mixup* strength α .

4.1 Model Architecture

We use the extended Onsets and Frames model [18] which increased the CNN channels to 48/48/96, the LSTM units to 256, and the FC units to 768. The extended model has total 26.5 million parameters. We do not use the frame loss weights described in [17] in favor of the offset stack introduced in the extended version (see Figure 1). During inference, we first calculate the posteriors corresponding to overlapping chunks of audio, with the same length as the training sequences, and perform overlap-add using Hamming windows to obtain the full-length posterior. This is because the effects of adversarial learning do not continue further than the training sequence length when we let the recurrent networks continue to predict longer sequences.

The input to the discriminator has two channels for the onset and frame predictions. The discriminator has 5 convolutional layers: c32k3s2p1, c64k3s2p1, c128k3s2p1, c256k3s2p1, c1k5s1p2, where the numbers indicate the number of output channels, the kernel size, the stride amount, and the padding size. At each non-final layer, dropout of probability 0.5 and leaky ReLU activation with negative slope 0.2 are used. The mean of the final layer output along the time and frequency axes is taken as the discriminator output.

4.2 Hyperparameters

Table 1 summarizes the hyperparameters used during the experiments, which are mostly taken directly from [17] and [21]. Also following [17], we use Adam [26] and apply learning rate decay of factor 0.98 in every 10,000 iterations, for both the generator and the discriminator. We examine two types of GAN losses, the non-saturating GAN ($\ell = \text{BCE}$) and the least-squares GAN ($\ell = \text{MSE}$). For each GAN loss, multiple values of *mixup* strengths are compared with $\alpha = 0$, i.e. no *mixup*. Training runs for one million iterations, and the iteration that best performs on the validation set are used for evaluation on the test set.

Hyperparameter	Values
Generator learning rate η	0.0006
Discriminator learning rate β	0.0001
Discriminator loss function ℓ	{BCE, MSE}
Batch size m	8
<code>pix2pix</code> weight ν	100
<i>mixup</i> strength α	{0, 0.2, 0.3, 0.4}
Activation threshold τ	0.5
Training sequence length	327,680

Table 1. Hyperparameters used during the experiments.

4.3 Dataset

We use the MAESTRO dataset [18], which contains Disklavier recordings of 1,184 classical piano performances. The dataset consists of 172.3 hours of audio, which are provided with 140.1, 15.3, and 16.9 hours of train/validation/test splits such that recordings of one composition only appear in the same split. We resample the audio to 16 kHz and down-mix into a single channel. Following [17], an STFT window of 2,048 samples is used for producing 229-bin Mel spectrograms, and a hop length of 32 ms is used. Training sequences sliced at random positions are used, unlike the official implementation which slices training sequences at silence or zero crossings.

4.4 Evaluation Metrics

The Onsets and Frames model perform both frame-level and note-level predictions, and their performance can be evaluated with the standard precision, recall, and F1 metrics. For multi-pitch estimation, we also report the error rate metrics defined in [36], which include total error, substitution error, miss error, and false alarm error. We use the *mir_eval* [37] library for all metric calculations. For the note-level metrics, we use the default settings of the library, which use 50 ms for the onset tolerance, 50 ms or 20% of the note length (whichever is longer) for the offset tolerance, and 0.1 for the velocity tolerance.

5. RESULTS

5.1 Comparison with the Baseline Metrics

Table 2 and 3 summarize the transcription performance, clearly showing a consistent improvement in the conditional GAN models over the Onsets and Frames baseline. Table 2 shows that both non-saturating GAN and least-squares GAN achieve the highest frame and note F1 scores when the *mixup* strength $\alpha = 0.3$ is used, and they both outperform the baseline. The binary piano rolls are easy to distinguish from the non-binary predictions, which may cause imbalanced adversarial training. *mixup* allows non-binary piano rolls to be fed to the discriminator, making its task more challenging and leading to higher performance.

Table 3 shows an important trend of the cGAN results compared to the baseline that cGAN trades off a bit of precision for a significant improvement in recall; this is a side effect of the cGAN producing more confident predictions, as will be discussed in the following subsections.

	<i>mixup</i> strength α	Frame Metrics							Note Metrics			Note Metrics with Offsets			Note Metrics with Offsets & Velocity		
		F1	P	R	E_{total}	E_{subs}	E_{miss}	E_{fa}	F1	P	R	F1	P	R	F1	P	R
Baseline		.899	.946	.857	.179	.013	.130	.036	.942	.990	.899	.802	.842	.765	.790	.830	.755
Non-Saturating GAN	0.3	.914	.931	.898	.156	.012	.089	.054	.956	.981	.932	.813	.835	.793	.802	.823	.782
Least-Squares GAN	0.3	.906	.942	.875	.167	.013	.113	.042	.950	.988	.916	.810	.841	.781	.799	.830	.771

Table 3. Summary of transcription performance using *mixup* strength $\alpha = 0.3$. The non-saturating GAN loss has the highest performance across all F1 metrics. The average metrics across the tracks in the MAESTRO test dataset are reported, and the model checkpoint where the average of frame F1 and note F1 is the highest on the validation dataset is used.

	Baseline	GAN type	<i>mixup</i> strength α			
			0	0.2	0.3	0.4
Frame F1	0.899	Non-Saturating	0.664	0.912	0.914	0.907
		Least-Squares	0.904	0.903	0.906	0.898
Note F1	0.942	Non-Saturating	0.717	0.953	0.956	0.951
		Least-Squares	0.944	0.947	0.950	0.943

Table 2. Frame and note F1 scores are the highest when the non-saturating GAN loss and $\alpha = 0.3$ are used.

While the percentage differences are moderate, our method achieves statistically significant improvements in F1 metrics on the MAESTRO test dataset ($p < 10^{-14}$ for all 4 metrics, two-tailed paired *t*-test). The distribution of per-track improvement in each F1 metric is shown in Figure 4, which indicates that the improvements are evenly distributed across the majority of the tracks. These improvements are especially promising, considering that Onsets and Frames is already a very strong baseline.

5.2 Visualization of Frame Activations

To better understand the inner workings of the conditional GAN framework, we visualize the frame posteriorgrams created by the baseline and the best performing conditional GAN model in Figure 3. In contrast to the baseline posteriorgrams which have many blurry segments, the posteriorgrams generated by our method mostly contain segments with solid colors, meaning that the model is more confident in its prediction. Figure 5 shows that the proportion of frame activation values in (0.1, 0.9) is noticeably higher in the baseline, thus making the output less sensitive to the threshold choice. This is because indecisive predictions are penalized by the discriminator, since they are easy to distinguish from the ground-truth which contains only binary labels. The generator is therefore encouraged to output the most probable note sequences even when it is unsure, rather than producing blurry posteriorgrams that might hamper the decoding process. This allows for an interpretation in which the GAN loss provides a prior for valid onset and frame activations, and the model learns to perform MAP estimation based on this prior.

5.3 Training Dynamics and The Generalization Gap

Figure 6 shows the learning curves for the frame F1 and note F1 scores, where the scores on the training dataset are plotted in dotted lines. It is noticeable in the figure

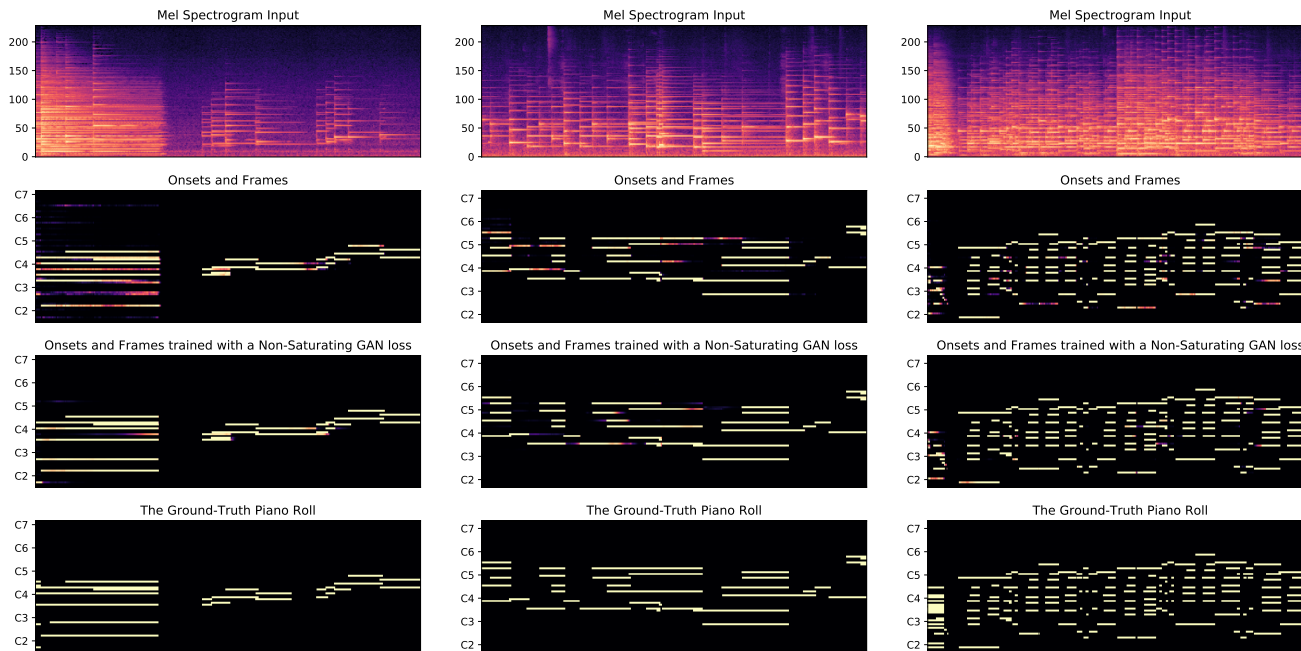


Figure 3. Comparisons of the frame activation posterior predicted by the baseline and our model ($\ell = \text{BCE}, \alpha = 0.3$), on three example segments. The input Mel spectrograms and the target piano rolls are shown together. The GAN version produces more confident predictions compared to the noisy baselines, leading to more accurate predictions.




Figure 4. F1 score improvements over the baseline, tested on the MAESTRO test tracks.




Figure 5. Distribution of frame activation values.




Figure 6. Learning curves showing the generalization gaps; training curves are dotted.

that the validation F1 scores for the baseline stagnate after 300k iterations, while the F1 scores of our model steadily grow until the end of 1 million iterations. Thanks to this, the generalization gap — the difference between the training and validation F1 scores — is significantly smaller for the conditional GAN model. This means that the GAN loss works as an effective regularizer that encourages the trained model to generalize better to unseen data, rather than memorizing the note sequences in the training dataset as LSTMs are known to be capable of [46].

6. CONCLUSIONS

We have presented an adversarial training method that can consistently outperform the baseline Onsets and Frames model, using the standard frame-level and note-level transcription metrics and visualizations that show how the improved model predicts more confident output. To achieve this, a discriminator network is trained competitively with the transcription model, i.e. a conditional generator, so that the discriminator serves as a learned regularizer that pro-

vides a prior for realistic note sequences.

Our results show that modeling the inter-label dependencies in the target distribution is important and brings measurable performance improvements. Our method is generic, and any model that involves predicting two-dimensional representation should be able to benefit from including an adversarial loss. These approaches are common not only in transcription models but also in speech or music synthesis models that predict spectrograms as an intermediate representation [24, 38].

Our results do not include the effects of using data augmentation [18], which is orthogonal to our approach and should bring additional performance improvements when applied. As discussed, the discriminator imposes the prior on the target domain whereas data augmentation enriches the input audio distribution. This implies that our method would be less effective when the majority of errors are due to the discrepancy in the audio distribution between the training and test datasets. How to apply adversarial learning for better generalization on the input distribution is a potential future research direction.

7. REFERENCES

- [1] Samer A Abdallah and Mark D Plumbley. Unsupervised analysis of polyphonic music by sparse coding. *IEEE Transactions on Neural Networks*, 17(1):179–196, 2006.
- [2] Emmanouil Benetos and Simon Dixon. Multiple-instrument polyphonic music transcription using a temporally constrained shift-invariant model. *The Journal of the Acoustical Society of America*, 133(3):1727–1741, 2013.
- [3] Emmanouil Benetos, Simon Dixon, Zhiyao Duan, and Sebastian Ewert. Automatic music transcription: An overview. *IEEE Signal Processing Magazine*, 36(1):20–30, 2019.
- [4] Emmanouil Benetos, Simon Dixon, Dimitrios Gianoulis, Holger Kirchhoff, and Anssi Klapuri. Automatic music transcription: challenges and future directions. *Journal of Intelligent Information Systems*, 41(3):407–434, 2013.
- [5] Rachel M Bittner, Brian McFee, Justin Salamon, Peter Li, and Juan Pablo Bello. Deep salience representations for f0 estimation in polyphonic music. In *Proceedings of the International Society for Music Information Retrieval (ISMIR) Conference*, pages 63–70, 2017.
- [6] Sebastian Böck and Markus Schedl. Polyphonic piano note transcription with recurrent neural networks. In *Proceedings of the IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 121–124, 2012.
- [7] Nicolas Boulanger-Lewandowski, Yoshua Bengio, and Pascal Vincent. Modeling temporal dependencies in high-dimensional sequences: Application to polyphonic music generation and transcription. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2012.
- [8] Antonia Creswell, Tom White, Vincent Dumoulin, Kai Arulkumaran, Biswa Sengupta, and Anil A Bharath. Generative adversarial networks: An overview. *IEEE Signal Processing Magazine*, 35(1):53–65, 2018.
- [9] Hao-Wen Dong, Wen-Yi Hsiao, Li-Chia Yang, and Yi-Hsuan Yang. MuseGAN: Multi-track sequential generative adversarial networks for symbolic music generation and accompaniment. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [10] Alexey Dosovitskiy and Thomas Brox. Generating images with perceptual similarity metrics based on deep networks. In *Advances in Neural Information Processing Systems*, pages 658–666, 2016.
- [11] Jesse Engel, Kumar Krishna Agrawal, Shuo Chen, Ishaan Gulrajani, Chris Donahue, and Adam Roberts. GANSynth: Adversarial neural audio synthesis. *arXiv preprint arXiv:1902.08710*, 2019.
- [12] Sebastian Ewert and Mark Sandler. Piano transcription in the studio using an extensible alternating directions framework. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 24(11):1983–1997, 2016.
- [13] Cédric Févotte and Jérôme Idier. Algorithms for non-negative matrix factorization with the β -divergence. *Neural computation*, 23(9):2421–2456, 2011.
- [14] Benoit Fuentes, Roland Badeau, and Gaël Richard. Harmonic adaptive latent component analysis of audio and application to music transcription. *IEEE Transactions on Audio, Speech, and Language Processing*, 21(9):1854–1866, 2013.
- [15] Ian Goodfellow. NIPS 2016 tutorial: Generative adversarial networks. *arXiv preprint arXiv:1701.00160*, 2016.
- [16] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014.
- [17] Curtis Hawthorne, Erich Elsen, Jialin Song, Adam Roberts, Ian Simon, Colin Raffel, Jesse Engel, Sageev Oore, and Douglas Eck. Onsets and frames: Dual-objective piano transcription. In *Proceedings of the International Society for Music Information Retrieval (ISMIR) Conference*, pages 50–57, 2018.
- [18] Curtis Hawthorne, Andrew Stasyuk, Adam Roberts, Ian Simon, Cheng-Zhi Anna Huang, Sander Dieleman, Erich Elsen, Jesse Engel, and Douglas Eck. Enabling factorized piano music modeling and generation with the MAESTRO dataset. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- [19] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, pages 6626–6637, 2017.
- [20] Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. A fast learning algorithm for deep belief nets. *Neural Computation*, 18(7):1527–1554, 2006.
- [21] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1125–1134, 2017.
- [22] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4401–4410.
- [23] Rainer Kelz, Matthias Dorfer, Filip Korzeniowski, Sebastian Böck, Andreas Arzt, and Gerhard Widmer. On the potential of simple framewise approaches to piano transcription. In *Proceedings of the International Society for Music Information Retrieval (ISMIR) Conference*, pages 475–481, 2016.
- [24] Jong Wook Kim, Rachel Bittner, Aparna Kumar, and Juan Pablo Bello. Neural music synthesis for flexible timbre control. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.

- [25] Jong Wook Kim, Justin Salamon, Peter Li, and Juan Pablo Bello. CREPE: A convolutional representation for pitch estimation. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 161–165, 2018.
- [26] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015.
- [27] Hugo Larochelle and Iain Murray. The neural autoregressive distribution estimator. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 29–37, 2011.
- [28] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436, 2015.
- [29] Daniel D Lee and H Sebastian Seung. Algorithms for non-negative matrix factorization. In *Advances in Neural Information Processing Systems*, pages 556–562, 2001.
- [30] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3431–3440, 2015.
- [31] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2794–2802, 2017.
- [32] Matthias Mauch and Simon Dixon. pYIN: A fundamental frequency estimator using probabilistic threshold distributions. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 659–663. IEEE, 2014.
- [33] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- [34] Juhan Nam, Jiquan Ngiam, Honglak Lee, and Malcolm Slaney. A classification-based polyphonic piano transcription approach using learned feature representations. In *Proceedings of the 12th International Society for Music Information Retrieval (ISMIR) Conference*, pages 175–180, 2011.
- [35] Yizhao Ni, Matt McVicar, Raul Santos-Rodriguez, and Tjil De Bie. An end-to-end machine learning system for harmonic analysis of music. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(6):1771–1783, 2012.
- [36] Graham E Poliner and Daniel PW Ellis. A discriminative model for polyphonic piano transcription. *EURASIP Journal on Advances in Signal Processing*, 2007(1):048317, 2006.
- [37] 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 *Proceedings of the International Society for Music Information Retrieval (ISMIR) Conference*, 2014.
- [38] Jonathan Shen, Ruoming Pang, Ron J Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Zhifeng Chen, Yu Zhang, Yuxuan Wang, Rj Skerrv-Ryan, et al. Natural TTS synthesis by conditioning WaveNet on Mel spectrogram predictions. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4779–4783. IEEE, 2018.
- [39] Siddharth Sigtia, Emmanouil Benetos, and Simon Dixon. An end-to-end neural network for polyphonic piano music transcription. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 24(5):927–939, 2016.
- [40] Paris Smaragdis and Judith C Brown. Non-negative matrix factorization for polyphonic music transcription. In *2003 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, pages 177–180, 2003.
- [41] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [42] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with PixelCNN decoders. In *Advances in Neural Information Processing Systems*, pages 4790–4798, 2016.
- [43] Aäron Van Den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 1747–1756, 2016.
- [44] Emmanuel Vincent, Nancy Bertin, and Roland Badeau. Adaptive harmonic spectral decomposition for multiple pitch estimation. *IEEE Transactions on Audio, Speech, and Language Processing*, 18(3):528–537, 2010.
- [45] Li-Chia Yang, Szu-Yu Chou, and Yi-Hsuan Yang. MidiNet: A convolutional generative adversarial network for symbolic-domain music generation. In *Proceedings of the International Society for Music Information Retrieval (ISMIR) Conference*, pages 324–331, 2017.
- [46] Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. Recurrent neural network regularization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015.
- [47] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.