

# SEMI-SUPERVISED CONVOLUTIVE NMF FOR AUTOMATIC PIANO TRANSCRIPTION

**Haoran Wu**  
 INSA Rennes, France  
 haoran.wu@insa-rennes.fr

**Axel Marmoret**  
 Univ Rennes, Inria, CNRS,  
 IRISA, France  
 axel.marmoret@inria.fr

**Jérémy E. Cohen**  
 Univ Lyon, INSA-Lyon, UCBL,  
 UJM-Saint Etienne, CNRS, Inserm,  
 CREATIS UMR5220, U1206, France  
 jeremy.cohen@cnrs.fr

## ABSTRACT

Automatic Music Transcription, which consists in transforming an audio recording of a musical performance into symbolic format, remains a difficult Music Information Retrieval task. In this work, which focuses on piano transcription, we propose a semi-supervised approach using low-rank matrix factorization techniques, in particular Convolutional Nonnegative Matrix Factorization. In the semi-supervised setting, only a single recording of each individual notes is required. We show on the MAPS dataset that the proposed semi-supervised CNMF method performs better than state-of-the-art low-rank factorization techniques and a little worse than supervised deep learning state-of-the-art methods, while however suffering from generalization issues.

## 1. INTRODUCTION

Automatic Music Transcription (AMT) is the task of transforming music recordings into symbolic format, such as scores or MIDI. It is a fundamental musical skill to acquire, taught from early age up to professional level in music schools and, given enough training, humans can be extremely accurate at transcription. Automatic music transcription aims at accelerating and improving time-consuming manual transcription and has applications in music tutoring and rehearsing, musicology analysis or in other music information retrieval tasks [1].

However, while audio generation from MIDI is rather mature, its counterpart AMT is still a very challenging task, even in scenarios involving a single multipitch instrument like a piano, which is our case study. As reported in the 2018 survey by Benetos *et. al.* [1], there are mainly two families of methods to perform AMT: 1) Methods based on low-rank factorizations of spectrograms, and in particular Nonnegative Matrix Factorization (NMF). These methods are mostly unsupervised [2–4]. 2) Deep Neural Networks (DNN) which are heavily supervised. They require registered symbolic-audio training data in a large amount, which can be hard to acquire [5–9].

A recent outbreak in the task of piano transcription (as well as other related tasks) is due to the release of the MAESTRO dataset [7], a large dataset of tightly matched MIDI and audio piano recordings of professional quality which improved the training quality of deep learning techniques. However, the supervised methods require extensive amounts of training data which may not be available for most instruments. The quality of the MAESTRO dataset comes from the existence of the Yamaha Disklavier™, which enables co-recording of audio and MIDI. This high level technology does not exist for most instruments, and building large training dataset for most polyphonic instruments would be extremely challenging on the practical side.

In contrast, since unsupervised factorization-based approaches do not require training data, they obviously solve the data frugality and generalization problems at the cost of being far less accurate than deep supervised approaches.

The goal of this paper is two-fold. On a first hand, leveraging training data available only in limited quantity. On another hand, deploying a variant of NMF, coined Convolutional NMF, in the context of transcription, to improve the transcription performance with respect to NMF. The most closely related work is surely the Attack Decay model [3], which also performs semi-supervision, and proposes a model reminiscent of CNMF. The major differences between the proposed CNMF framework and this work of Cheng *et. al.* are discussed in Section 2.2. Moreover, in Section 4, we show that the performance of the proposed approach are generally much higher and can reach the performance levels observed with Deep Learning at the cost of poor generalization properties. In [4], authors also consider CNMF for piano transcription but CNMF is not the main focus of their work.

This paper is organized as follows: in Section 2, we review the basics of NMF and CNMF for transcription. In Section 3, semi-supervised CNMF is introduced. In Section 4 we show experimental results on MAPS and MAESTRO. Section 5 is devoted to discussions and perspectives.

**Notations:** Matrices and higher-order arrays are denoted by capital letters,  $T_{ijk}$  is the element  $(i, j, k)$  in the three-way array  $T$ . To denote slices, we use semicolons, so that  $T_{i::}$  denotes for instance the slice of all elements of  $T$  on row  $i$ . Finally, we denote  $T_{[a:b]jk}$  elements  $(i, j, k)$  with  $i \in [a, b]$ .

Copyright: © 2022 Haoran Wu *et al.* This is an open-access article distributed under the terms of the [Creative Commons Attribution 3.0 Unported License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

## 2. CNMF FOR TRANSCRIPTION

### 2.1 NMF and CNMF Formalisms

Given an element-wise nonnegative matrix  $M \in \mathbb{R}_+^{n \times m}$  indexed as  $M_{ft}$  with  $f \in [1, n]$ ,  $t \in [1, m]$ , Nonnegative Matrix Factorization (NMF) is a low-rank approximation technique that summarizes  $M$  as a sum of rank-one parts, such that

$$M_{ft} = \sum_{q=1}^r W_{fq} H_{qt} \quad (1)$$

where  $r \leq \min(n, m)$  is a user-defined parameter relating to the number of patterns underlying  $M$ , see Figure 1. In practice, when  $M$  is an amplitude spectrogram, such as in this work, NMF is computed approximately and boils down to solving a bi-level constrained optimization problem

$$\operatorname{argmin}_{W \in \mathbb{R}_+^{n \times r}, H \in \mathbb{R}_+^{r \times m}} D_{KL}(M, WH) \quad (2)$$

where  $D_{KL}(M, WH)$  is the element-wise Kullback-Leibler divergence between matrix  $M$  and its nonnegative low-rank approximation  $WH = \sum_{q=1}^r W_{:q} H_{q\cdot}$ . In AMT, parameter  $r$  often relates to the number of notes expected in the recording, and therefore is generally set to (sometimes a multiple of)  $r = 88$  for piano recordings [2].

Furthermore, factor matrices  $W$  and  $H$  are respectively related to pitch and time activation. More specifically, each column of  $W$  is expected to contain a spectral template characteristic of a single pitch on the instrument used in the recording, while each corresponding row in  $H$  is expected to provide the activation of that note in the recording [10], see Figure 2.

An immediate critic about applying NMF to AMT is that reducing a note to a single frequency template, even tailored for a given instrument, is too restrictive. In practice, frequency templates should evolve with both amplitude and time. While explicit amplitude dependence would break the principle of low-rank approximation underlying NMF, it is possible to extend NMF to include a time-dependence on the templates, which yields Convolutional NMF [12]:

$$\operatorname{argmin}_{W \in \mathbb{R}_+^{n \times \tau \times r}, H \in \mathbb{R}_+^{r \times m}} D_{KL}(M, \sum_{q=1}^r W_{::q} * H_{q\cdot}) \quad (3)$$

where  $[W_{::q} * H_{q\cdot}]_{:t} = \sum_{i=0}^{\tau-1} W_{:iq} H_{q(t-i)}$  is a discrete convolution and  $q \in [1, r]$ , see Figure 1 for an illustration. By convention, we set  $H_{q(t-i)} = 0$  whenever  $t-i \leq 0$ . Integer  $\tau$  is again a user-defined hyperparameter that dictates the size of the convolution window. To provide a different perspective, the element-wise noiseless CNMF also writes

$$M_{ft} = \sum_{q=1}^r \sum_{i=0}^{\tau-1} W_{fiq} H_{q(t-i)}. \quad (4)$$

In a nutshell, CNMF enriches NMF by allowing each note to have a full STFT matrix  $W_{:q}$  as a frequency template instead of a single column. Therefore, it may also capture time-dependent events such as echoes or non-uniform partials attenuation. It can also be interpreted as a

constrained NMF with large rank  $r \times \tau$  where each note is represented by  $\tau$  templates, and the corresponding  $\tau$  rows in  $H$  are constrained to be equal up to a shift. Other works have also considered enriching NMF with several templates per note albeit not using convolution, typically by fusing rows of the estimated  $H$  matrix a posteriori [13–15].

### 2.2 Comparing the Attack Decay Model With CNMF

A reader familiar with the work of Cheng *et al.* [3] will notice that our work is similar in several aspects with their proposed Attack Decay (AD) framework for music transcription, but let us properly compare the models. After some rewriting of the original AD (see additional material<sup>1</sup>), AD decomposes the data  $M$  into two terms

$$M_{ft} = \underbrace{\sum_{q=1}^r \sum_{i=0}^{2\tau} \left( \tilde{W}_{fq}^{\text{attack}} P_{-i} \right) H_{q(t-i)}}_{\text{one note attack}} \quad (5)$$

$$+ \underbrace{\sum_{q=1}^r \sum_{i=\tau}^{t+\tau-1} \left( \tilde{W}_{fq}^{\text{decay}} e^{-\alpha_q(i-\tau)} \right) H_{q(t-i)}}_{\text{one note decay}}. \quad (6)$$

It thus appears that the attack term is a CNMF with rank-one templates  $W_{fiq}^{(\text{CNMF})} = W_{fq}^a P_i$  which is therefore less general than the CNMF model. The decay term is also a CNMF with rank-one templates.

With some further manipulations, one can see that it is possible to entirely recast the AD model as a CNMF model with rank-two templates, which may explain the performance gap between the two models observed in Section 4. Indeed in the semi-supervised setting, we seem to have enough data to learn unconstrained templates  $W^{\text{train}}$ , and the Attack-Decay structure on the templates may not be beneficial.

## 3. TEMPLATE LEARNING AND CNMF

### 3.1 Challenges in Unsupervised CNMF

In the context of music transcription, it is rarely discussed why NMF performs extremely well on simple dataset, but rather poorly on more complex ones. Saying that NMF, or CNMF, is a part-based representation with no destructive interferences between components does not explain this behavior. In fact, supposing the data indeed is generated reasonably well with a “ground-truth” NMF  $M = AB$  for some true frequency templates  $A \in \mathbb{R}_+^{n \times r}$  and activations  $B \in \mathbb{R}_+^{r \times m}$ , we need to ensure that computing an exact NMF  $M = WH$  will indeed yield  $A = W$  and  $B = H$ . In other words, the data  $M$  must admit a unique NMF.

Theoretically speaking, it is known that NMF will only enjoy this uniqueness property in particular cases, such as when sources are sufficiently scattered or when the data is very sparse [16, 17]. While this may hold for simple songs where notes do not overlap a lot, in the general case one should **not** expect that  $W$  and  $H$  behave as expected without restricting the set of solutions. Even worse, CNMF

<sup>1</sup> <https://github.com/cohenjer/TransSSCNMF>

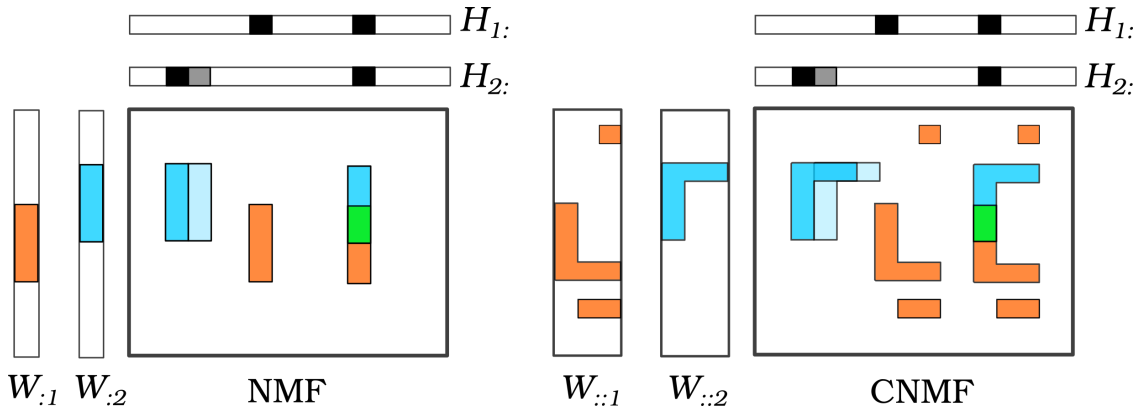


Figure 1. A visual comparison of NMF (left) and CNMF (right). CNMF allows to model complex time dependence while maintaining the number of templates low.




Figure 2. A toy example of transcription using NMF (adapted from [11]).

being a generalized NMF model, it is bound to have even weaker uniqueness properties than NMF (but nothing is known on CNMF identifiability to the best of our knowledge). Blind CNMF has been used with additional sparsity constraints for drums transcription, but dealing with drums typically yields much sparser and lower-rank data than pitched audio due to the temporal localization of percussive sounds.

Therefore, in general, unsupervised CNMF is not regularized enough to perform transcription. While some works focus on further regularization of NMF [18], we instead turn towards semi-supervision.

### 3.2 Learning Note-Wise Templates

Our working hypothesis is that audio recordings of isolated pitches are available, similarly to what is used for virtual instruments, except that we only make use of one template per note. Each recording is processed as the module of its complex STFT, denoted  $V_{::q} \in \mathbb{R}_+^{n \times m_q}$  where  $m_q$  is the number of STFT frames for that recording. For a regular piano one needs 88 such templates. Apart from pitch knowledge, no registered MIDI information is required.

The goal of the learning phase here is to estimate  $W_{::q}$  for each  $q$  using each individual recording  $V_{::q}$ . We propose to compute an approximate rank-one CNMF of each  $V_{::q}$  to estimate  $W_{::q}$  and  $h_q^{\text{train}}$ , the latter being discarded after the training phase. From a theoretical perspective, rank-one CNMF is a constrained version of NMF of rank  $\tau$ , furthermore computed on a very simple dataset. Therefore it fulfills the qualitative NMF uniqueness criteria discussed above, and we expect the recovered  $W$  to contain adequate




Figure 3. Three trained templates from the AkPnCGdD synthetic piano in MAPS, using  $\tau = 10$  convolution size. Templates  $W_{::q}$  have been square rooted to better highlight higher frequencies.

note frequency templates.

Practically, we solve for each  $q \in [1, \tau]$  the following optimization problem

$$W_{::q}^{\text{train}}, h_q^{\text{train}} \in \underset{W \in \mathbb{R}_+^{n \times \tau \times 1}, h \in \mathbb{R}_+^{1 \times m}}{\operatorname{argmin}} D_{KL}(V_q, W * h) \quad (7)$$

using a recently proposed multiplicative algorithm [19] which alternates between  $W$  and  $h$  updates while preserving nonnegativity and ensuring cost decrease.

In spite of the rank-one approximation and the simple data, the optimization problem still proves challenging with many local minima. Therefore initialization plays an important role in the learning phase. Because it is reasonable to look for  $W_{::q}$  in the  $V_{::q}$  data itself, we set

$$W_{::q}^{\text{init}} = V_{:[t^*:t^*+\tau-1]q} \text{ and } t^* = \underset{t \leq m_q}{\operatorname{argmax}} \|V_{:[t:t+\tau-1]q}\|_1 \quad (8)$$

which amounts to finding the  $\tau$  consecutive columns with most energy for initialization. Then we fill  $h_q^{\text{init}}$  with zeros and place a one at  $t^*$ . Note that this initialization procedure mimics a recently proposed algorithm for separable CNMF<sup>2</sup> [20] but is less computationally intensive. A total of 500 outer iterations are performed to learn a single note template.

Once the training phase is over, for a single multipitch

<sup>2</sup> Separable CNMF is a computationally simpler variant of CNMF which looks for all matrices  $W_{::q}$  in the data itself.




Figure 4. An example of  $H^{\text{test}}$  computed through rank-one CNMF.

instrument, we have at our disposal the whole dictionary  $W^{\text{train}}$ , see Figure 3.

### 3.3 CNMF Transcription With Templates

Testing in the semi-supervised framework only consists of computing the time activations  $H$  for a given music excerpt  $M$  to transcribe, since  $W$  has been pre-trained. This makes the transcription task much easier since the problem

$$H^{\text{test}} \in \underset{H \in \mathbb{R}_+^{r \times m}}{\text{argmin}} D_{KL}(M, \sum_{q=1}^r W_{::q}^{\text{train}} * H_q) \quad (9)$$

is convex and therefore can be solved up to arbitrary precision with the algorithm proposed in [19]. In practice 100 iterations are used, which is generally enough to reach convergence. Initialization was carried out using a few iterations of NMF with  $W$  fixed as the first column of each trained template  $W_{::q}^{\text{train}}$ . An example output  $H^{\text{test}}$  is provided in Figure 4.

### 3.4 Post-Processing of Activations

The post-processing of  $H^{\text{test}}$  that produces a MIDI file matters a lot. Hopefully, prior works have already proposed quite efficient post-processing using an adaptive threshold [3]. We essentially use the same technique but simplified.

In short, activation values in each row of  $H^{\text{test}}$ , averaged over several consecutive frames, are added to a user-defined threshold  $\delta$ , defining an adaptive threshold. An onset is detected at the position where the signal is above this adaptive threshold, see Figure 5 for an illustration. Formally, an onset is detected at frame  $t$  for note  $q$  when

$$h_{qt} > \frac{1}{21} \sum_{j=-10}^{10} h_{q(t+j)} + \delta, \quad (10)$$

using zero-padding when necessary. The activations are typically very sparse, so we generally did not observe spurious double peaks using the adaptive threshold contrarily to what was observed in [3].




Figure 5. The CNMF activation and adaptive peak-picking method shown for note F4 using a song from MAPS.

## 4. EXPERIMENTS ON MAPS AND MAESTRO

Although the proposed semi-supervised CNMF framework works in principle for transcribing any multipitch instrument, we only evaluate the performance for piano transcription as a proof on concept. Among the few existing open piano recordings dataset with registered audio and MIDI, in Section 4.2 we focused especially on MAPS [21] which has several kinds of individual notes recordings for several pianos, both virtual and acoustic. We also used MAESTRO [7] to evaluate generalization performance in Section 4.3. In our tests, we only considered the first 30 seconds of each song, as in [3]. Results are discussed in Section 5.

### 4.1 Experimental Setup

Let us briefly state the various experimental parameters required to reproduce the experiments<sup>3</sup>. All time signals are sampled at 44100Hz, the STFT is computed with windows of 80ms (3528 samples) with a hop-length of 20ms (882 samples). This results in  $n = 4097$  frequency bins and  $m = 1501$  time frames in the STFT for 30s of raw audio signal. No smoothing is applied to the STFT, and we set  $M$  as the amplitude spectrogram.

The  $\tau$  values are chosen among  $\tau = 5, 10, 20$ . We used one template for each piano note such that  $r = 88$ . Finally to fix the peak-picking threshold for  $H^{\text{test}}$ , two oracle strategies are used: 1) use the same threshold for all songs, and report results for the best value on the grid  $[0.01 : 0.01 : 0.4]$  2) perform transcription with a song-dependent threshold, optimized on the same grid. The first case corresponds to a scenario where the threshold is pre-trained for a category of recording (music genre, recording conditions) while the second case corresponds to a hand-tuned threshold for a specific song to transcribe.

We compare our method with the Attack Decay (AD) model presented in Section 3 which, to the best of our knowledge, is the current state-of-the-art for

<sup>3</sup> Python code to compute CNMF and reproduce all the experiments is available at <https://github.com/cohenjer/TransSSCNMF>

unsupervised/semi-supervised piano transcription. The results reported in Table 1 are the exact results from [3] (AD [3]), and the results of AD when applying our post-processing (AD\*). In both cases transcription is performed on  $H^{\text{attack}}$  as defined in [3]. Despite our efforts we were unable to exactly reproduce the original AD scores. In particular the original AD paper introduces smoothness in several aspects: the data spectrograms are locally averaged, and peak-fusion is performed in the post-processing. Consequently, the AD\* results enable comparison between the proposed CNMF and AD in the same pre/postprocessing conditions, while AD [3] are the best results achieved by Cheng *et. al.*. For completeness, we also report the transcription score from the state-of-the-art piano transcription network introduced in [8] which was trained on MAESTRO [7].

To measure performance, we compute a notewise score using the `mir_eval` [22] toolbox with a tolerance of 50ms. The offset detection problem is not tackled. Results are shown using only F-measure (F) and Accuracy (A) metrics (reported in percent), but full results including Precision and Recalls for all pianos are available in the complementary materials online.

#### 4.2 Transcription Performance on MAPS

The MAPS dataset contains classical piano music pieces recorded with different pianos and conditions: a Yamaha Disklavier<sup>TM</sup> in two settings 'ENSTDkCl' (EN1) and 'ENSTDkAm' (EN2), and six synthetic pianos 'AkPnBcht', 'AkPnBsdf' (AkB1-2), 'AkPnCGdD' (AkC), 'AkPnStgb' (AkS), 'SptkBGAm' (Sp), 'StbgTGd2' (St). For each piano/setting listed above, we train a template  $W^{\text{train}}$  using a rank-one CNMF as presented in Equation 7. Since there are many available single notes recordings in MAPS, we chose based on performance to use the Isolated notes (ISOL / NO) recorded at Medium intensity (M).

In a first experiment, we study the sensitivity of the proposed method to the selection of the convolution window size  $\tau$  and the choice of a threshold  $\delta$  tuned on the whole corpus versus on each song individually. We also compare our results to the Attack Decay model and the ByteDance supervised neural network. Results are shown in Table 1. From this experiment, we see that generally  $\tau = 10$  performs best, and that the song-wise threshold gives better results.

In a second experiment, the templates for all other pianos are used to transcribe AkPnCGdD and ENSTDkCl to estimate the generalization capacities of the trained CNMF templates, see results in Table 2. Only CNMF with song-tuned threshold is shown and we set  $\tau = 10$ , to show the best results only.

Additionally, Table 3 reports the average running times when training the templates and performing transcription on the AkPnCGdD recordings. This test was run on a personal computer with AMD Ryzen 5 2600<sup>TM</sup> processor and 16GB RAM.

#### 4.3 Generalization on MAESTRO

A natural question regarding CNMF templates is how well they can be used outside their training context without any domain adaptation. While results shown in Table 2 already provide a partial answer, we also tried to apply CNMF to the MAESTRO dataset. However, since no individual notes recordings are publicly available for MAESTRO, we used the templates learnt from MAPS. We transcribed 20 songs from the MAESTRO test set randomly chosen.

The results are quite poor: even when choosing song-wise thresholds, for all templates, CNMF does not reach above 59% in F-measure (test results are available in the supplementary materials). For comparison, the state-of-the-art with supervised deep learning techniques reaches above 95% F-measure on MAESTRO. Its performance on MAPS with data augmentation are also state-of-the-art, around 89% F-measure on EN1, despite the training/testing mismatch.

### 5. DISCUSSION

In light of the experiments conducted in Section 4, let us discuss the strengths of the proposed CNMF. It exhibits a significant improvement with respect to the Attack Decay model, which as far as we know is state-of-the-art for semi-supervised piano transcription. This is even more true when using the same pre-processing and post-processing for AD and CNMF, the former being in particular prone to unstable activations which were not observed in the latter. We may therefore affirm that the improvement in performance is indeed due to the CNMF model design. In other words, CNMF with a semi-supervised setting is an efficient piano transcription method. From numerical results, it seems that a convolution window size  $\tau = 10$  is a good compromise between quality of transcription and transcription computation time.

The CNMF method does not perform better than the supervised state-of-the-art method we denoted as ByteDance DNN, which is expected given that this neural-network competitor is trained on MAESTRO which contains more than two hundred hours of perfectly aligned MIDI and audio piano recording of professional level. We still reach similar performances on some pianos such as EN1, EN2 and AkS. Nevertheless, the ByteDance DNN is not trained on MAPS contrarily to the proposed semi-supervised CNMF.

Moreover, the proposed semi-supervised setting only requires a handful of training dataset which are relatively easy to acquire. Indeed, only individual notes recordings are necessary, without any audio and MIDI registration. Compared to the very large amount of data currently required by state-of-the-art deep learning approaches, this is a huge advantage of the proposed approach applicable to any acoustic instrument with well-defined onsets readily available. Sadly our study is limited to piano transcription. A perspective of this work is to apply it to transcribe polyphonic instruments for which recording registered MIDI and audio is challenging.

Finally, while the performance does depend on the choice

thresh	$\tau$	EN1		EN2		AkB1		AkB2		AkC		AkS		Sp		St		
		F	A	F	A	F	A	F	A	F	A	F	A	F	A	F	A	
global	CNMF	5	78	65	70	55	88	80	75	62	83	72	80	69	81	70	75	61
		10	<b>85</b>	<b>75</b>	<b>77</b>	<b>64</b>	93	88	87	78	91	84	<b>88</b>	<b>79</b>	89	82	84	74
		20	83	72	76	63	<b>94</b>	<b>89</b>	<b>87</b>	<b>79</b>	<b>92</b>	<b>86</b>	87	79	<b>90</b>	<b>83</b>	<b>86</b>	<b>77</b>
	AD*	81	69	68	53	66	50	71	56	60	43	67	51	64	47	67	50	
song	CNMF	5	82	70	74	59	90	82	80	69	87	78	84	74	86	77	81	69
		10	<b>88</b>	<b>79</b>	<b>80</b>	<b>68</b>	<b>95</b>	<b>91</b>	<b>90</b>	<b>83</b>	94	89	<b>90</b>	<b>82</b>	<b>93</b>	<b>87</b>	89	80
		20	85	75	78	66	<b>95</b>	<b>91</b>	<b>90</b>	<b>83</b>	<b>94</b>	<b>90</b>	89	81	92	87	<b>89</b>	<b>81</b>
	AD*	82	70	69	54	68	52	73	59	61	45	69	54	66	50	70	54	
AD [3]		82	70	-	-	-	-	-	-	85	74	-	-	-	-	-	-	
ByteDance DNN [8]		<b>89</b>	<b>81</b>	77	65	<b>98</b>	<b>97</b>	<b>95</b>	<b>90</b>	<b>98</b>	<b>96</b>	87	77	<b>97</b>	<b>95</b>	<b>95</b>	<b>90</b>	

Table 1. CNMF, AD, and the ByteDance supervised network performance with respect to the choice of hyperparameter  $\tau$  and the choice of the peak-picking threshold, without training/testing mismatch for CNMF and AD. Only the first 30s of each songs were used. AD\* uses the same pre/post-processing as CNMF. Tolerance is 50ms.

AkC	F	EN2	AkB1	AkB2	AkS	Sp	St
		A	74	77	77	70	74
EN1	F	76	67	68	69	67	69
	A	62	50	52	53	52	53

Table 2. Transcription scores for CNMF, with training/testing mismatch.

$\tau$	Training			Transcription
	5	10	20	10
Av. time	56s	193s	634s	239s

Table 3. Average computation time for a learning pattern (Training) or transcribing 30s of a song (Transcription) for semi-supervised CNMF. Results are reported for the AkP-nCGdD piano in MAPS.

of a good activation threshold, CNMF still performs well using a global threshold over all songs in MAPS for each piano. Therefore extensively tuning the threshold hyperparameter is not essential to the success of CNMF here.

Despite these encouraging results, CNMF has a few issues which open interesting perspectives. First, it clearly has a significant generalization problem, or in other words, the learning stage overfits the training data. From Table 2, it appears that a mismatch between training and testing inside MAPS, while detrimental to transcription performance, is not as severe as a learning on MAPS and testing on MAESTRO. A tentative explanation is that the MAESTRO recordings are live performances with quite loud reverberation, while the MAPS recordings are drier. Looking for an audio transformation of the templates that minimizes recording conditions mismatch would therefore probably prove beneficial to generalize pre-recorded CNMF templates. Retraining a template library given few annotated data in the testing set could also be a possible solution. Whether this domain adaptation can be done fully blindly is still unclear however.

Second, despite performance not relying too much on the threshold level, the threshold selection method on the other hand is extremely important. Using a fixed threshold in-

stead of the adaptive peak-picking drastically decreased performances in our early tests. But this also means that the post-processing of activations can be further improved using more involved technique than thresholding each note individually, and this research direction should not be overlooked if transcription performances of CNMF are to be further improved.

Third, for simplicity only one template for each note was used for the transcription phase. However, most instruments sound quite differently depending on how they are played. The proposed semi-supervised framework currently does not account for this timbre variation with amplitude or technique, and adapting the current method to make use of several templates per notes is an interesting research direction.

Finally according to the results shown in Table 3, computation time is rather large even in the testing phase. With the current implementation, real-time processing is therefore prohibited. Using a CNMF solver dedicated to Kullback-Leibler divergence or working on a more efficient rank-one CNMF solver than [19] could nevertheless drastically reduce computation time.

## 6. CONCLUSION

The state-of-the-art for automatic piano transcription is undeniably nowadays detained by deep learning techniques. However these methods rely on very large audio and symbolic registered dataset which are potentially very hard to obtain. In this work, we propose a competitive semi-supervised matrix factorization model which only requires labeled recordings of each individual notes. We show that when there is no mismatch between the training data and the test data, our approach performs significantly better than semi-supervised state-of-the-art approaches, approaching supervised deep learning performance. Therefore, we believe that using CNMF instead of NMF is an important step towards learning more reasonable frequency templates in low-rank approximation techniques for piano transcription or other similar tasks. Further works should however be devoted to adapt pre-trained templates to reduce generalization error. Improving the onset detection

method, allowing timbre variation in templates and reducing computation time are other important research directions. Finally, the proposed semi-supervised approach should be tested with other instruments than the piano and in a multi-instrument setup.

### Acknowledgments

Jeremy E. Cohen and Axel Marmoret thank ANR JCJC LoRAiA ANR-20-CE23-0010 for supporting this work.

### 7. REFERENCES

- [1] E. Benetos, S. Dixon, Z. Duan, and S. Ewert, "Automatic music transcription: An overview," *IEEE Signal Processing Magazine*, vol. 36, no. 1, pp. 20–30, 2018.
- [2] E. Vincent, N. Bertin, and R. Badeau, "Harmonic and inharmonic nonnegative matrix factorization for polyphonic pitch transcription," in *2008 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2008, pp. 109–112.
- [3] T. Cheng, M. Mauch, E. Benetos, and S. Dixon, "An attack/decay model for piano transcription," in *ISMIR 2016-17th International Society for Music Information Retrieval*, 2016.
- [4] L. Gao, L. Su, Y.-H. Yang, and T. Lee, "Polyphonic piano note transcription with non-negative matrix factorization of differential spectrogram," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 291–295.
- [5] 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.
- [6] 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," *Proceedings of the 19th International Society for Music Information Retrieval Conference*, 2018.
- [7] C. Hawthorne, A. Stasyuk, A. Roberts, I. Simon, C.-Z. A. Huang, S. Dieleman, E. Elsen, J. Engel, and D. Eck, "Enabling factorized piano music modeling and generation with the MAESTRO dataset," in *International Conference on Learning Representations*, 2019.
- [8] Q. Kong, B. Li, X. Song, Y. Wan, and Y. Wang, "High-resolution piano transcription with pedals by regressing onsets and offsets times," *arXiv preprint arXiv:2010.01815*, 2020.
- [9] Y. Yan, F. Cwitkowitz, and Z. Duan, "Skipping the frame-level: Event-based piano transcription with neural semi-crfs," *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [10] P. Smaragdis and J. C. Brown, "Non-negative matrix factorization for polyphonic music transcription," in *2003 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*. IEEE, 2003, pp. 177–180.
- [11] N. Bertin, "Les factorisations en matrices non-négatives : approches contraintes et probabilistes, application à la transcription automatique de musique polyphonique," Ph.D. dissertation, 2009. [Online]. Available: <http://www.theses.fr/2009ENST0051>
- [12] P. Smaragdis, "Convolutional speech bases and their application to supervised speech separation," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 1, pp. 1–12, 2006.
- [13] T.-M. Wang, P.-Y. Tsai, and A. W. Su, "Score-informed pitch-wise alignment using score-driven non-negative matrix factorization," in *2012 International Conference on Audio, Language and Image Processing*. IEEE, 2012, pp. 206–211.
- [14] E. Benetos, A. Klapuri, and S. Dixon, "Score-informed transcription for automatic piano tutoring," in *2012 Proceedings of the 20th European Signal Processing Conference (EUSIPCO)*. IEEE, 2012, pp. 2153–2157.
- [15] D. Jeong and J. Nam, "Note intensity estimation of piano recordings by score-informed nmf," in *Audio Engineering Society Conference: 2017 AES International Conference on Semantic Audio*. Audio Engineering Society, 2017.
- [16] D. Donoho and V. Stodden, "When does non-negative matrix factorization give a correct decomposition into parts?" in *In Advances in Neural Information Processing 16*, 2003.
- [17] X. Fu, K. Huang, and N. D. Sidiropoulos, "On identifiability of nonnegative matrix factorization," *IEEE Signal Processing Letters*, vol. 25, no. 3, pp. 328–332, 2018.
- [18] V. Leplat, A. M. Ang, and N. Gillis, "Minimum-volume rank-deficient nonnegative matrix factorizations," in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 3402–3406.
- [19] D. Fagot, H. Wendt, C. Févotte, and P. Smaragdis, "Majorization-minimization algorithms for convolutional NMF with the beta-divergence," in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 8202–8206.
- [20] A. Degleris and N. Gillis, "A provably correct and robust algorithm for convolutional nonnegative matrix factorization," *IEEE Transactions on Signal Processing*, vol. 68, pp. 2499–2512, 2020.



- [21] V. Emiya, N. Bertin, B. David, and R. Badeau, “MAPS-a piano database for multipitch estimation and automatic transcription of music,” 2010.
- [22] C. Raffel, B. McFee, E. J. Humphrey, J. Salamon, O. Nieto, D. Liang, D. P. Ellis, and C. C. Raffel, “mir\_eval: A transparent implementation of common mir metrics,” in *In Proceedings of the 15th International Society for Music Information Retrieval Conference, ISMIR*, 2014.