EXTENDING NONNEGATIVE MATRIX FACTORIZATION—A DISCUSSION IN THE CONTEXT OF MULTIPLE FREQUENCY ESTIMATION OF MUSICAL SIGNALS

Stanisław Raczyński, Nobutaka Ono, Shigeki Sagayama

Graduate School of Information Science and Technology, The University of Tokyo
7-3-1, Hongo, Bunkyo-ku, Tokyo, 113-8656, Japan
phone: +(81)3-5841-6902, fax: +(81)3-5841-6902,
email: {raczynski,sagayama,onono}@hil.t.u-tokyo.ac.jp
web: http://hil.t.u-tokyo.ac.jp/index-e.html

ABSTRACT

Nonnegative Matrix Approximation (NNMA) is a very well known technique of multivariate data analysis. However, in its basic form it provides very little control over its behaviour. This article explores possible extensions to this method in the context of multiple frequency estimation: using a parameterized distortion measure, enforcing harmonic structure on the basis matrix and introducing additional regularizations. We provide the reader with three regularizations useful for multiple pitch estimation and propose an objective way of evaluating the performance of NNMA-based pitch estimators. Finally, we use this evaluation method to train the parameters of our regularized harmonic NNMA and present the results.

1. INTRODUCTION

Music Information Retrieval (MIR) is a dynamically growing interdisciplinary field of science that aims at retrieving useful information from musical signals, both acoustic and symbolic. One of the most important tasks that belong to this domain is the multipitch analysis—a task of estimating multiple simultanous pitches present in a musical signal. Unfortunately, pitched sounds in most cases consist of a fundamental tone and its overtones—tones with frequencies being integer multiples of the fundamental frequency. Presence of these multiple tones corresponding to multiple pitched sounds makes this task very difficult. Furthermore, in musical signals, it is very common to have rational proportions between fundamental frequencies of different pitched sounds, as such signals sound much more pleasant to the human ear, and, as a consequence, some pitched sounds share overtones between themselves, which further complicates the analysis process.

One of the most common ways of looking at the problem of multiple pitch estimation is as a matrix factorization, where the spectrogram matrix is approximated with a sum of N rank-1 matrices (or, equivalently, a product of two full-rank matrices). A common way of performing this factorization is by means of the Nonnegative Matrix Factorization (NMF, proposed in [9], with a fast and convinient algorithm given in [8]), also known as Nonnegative Matrix Approximation (NNMA) [13]. In this paper we will use the latter name, unless refering to papers that use the former one. This technique decomposes a nonnegative (having only nonnegative elements) magnitude or power spectrogram matrix $\mathbf{X}$ (later referred to as the data matrix) into a product of two, also nonnegative, matrices $\mathbf{A}$ and $\mathbf{S}$:

$$\mathbf{X} \cong \mathbf{AS} = \widetilde{\mathbf{X}}. \tag{1}$$

The decomposition is being done by minimizing a distortion measure between the data $\mathbf{X}$ and its approximation $\mathbf{AS}$. The original algorithm from [8] was designed to minimize Euclidean distance or I-divergence, which are still most commonly used. This method has gained a wide recognition among researchers in the Music Information Retrieval field—more than half of last year's (2008) entries to MIREX's task of multipitch analysis was based on Nonnegative Matrix Factorization [2].

NNMA approximates each column of the data matrix with a linear combination of basis vectors $\mathbf{a}_t$ (columns of $\mathbf{A}$):

$$\widetilde{\mathbf{x}}_t = \sum_n s_{n,t}\mathbf{a}_t, \tag{2}$$

where n is the basis vector number and t is the time index. Because the matrices are nonnegative, only additive mixtures of nonnegative basis vectors (interpreted as spectra of individual notes) are possible, which is consistent with the way individual note sounds are combined to form the musical signal we aim to analyse. Therefore, we call the coefficient matrix $\mathbf{S}$ the *note activity matrix*, since it is assumed to contain amplitudes of notes.

Different variations and extensions of the NNMA algorithm have been used for multipitch analysis: the regular NMF [12], penalized NMFs, such as the Nonnegative Sparse Coding (NNSC) [4, 3], or NMF with basis vectors extended to contain spectrotemporal signatures (a number of consequent data frames), such as Nonnegative Matrix Factor 2-D Deconvolution (NMF2D) and Sparse Nonnegative Matrix Factor 2-D Deconvolution (SNMF2D) [11]. NMF2D and SNMF2D use a single signature for every note (of a single instrument), making use of the shift-similarity of logarithmic frequency scale spectra of notes played on a single instrument. All of these methods, however, when used directly, do not guarantee that the results will be usable for multipitch analysis, i.e. that the basis matrix will contain note spectra and the coefficient matrix their activities. This might be true for a very unrealistic case, when note events occur independently and sparsely, note spectra do not change their harmonic structure over time and the number of basis vectors correspond to the number of different notes actually occurring in the analyzed signal, but even that there is no guarantee NNMA will perform multipitch analysis by itself. That is why we strongly believe that in order to use NNMA for that purpose, we need to develop additional constraints and penalties, which are specific to musical signals.

In [10] we have proposed a method we called Harmonic Nonnegative Matrix Factorization (HNNMA), which

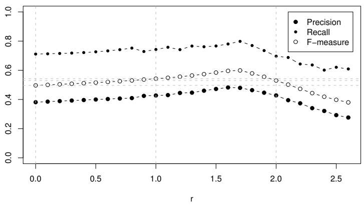

Figure 1: Multiple pitch estimation accuracy for different values of r. Gray dashed lines mark points that correspond to the Euclidean distance, the I-divergence and the Itakura-Saito divergence.

extended the regular NNMA by adding regularizations and a harmonic constraint on the basis matrix. In this paper we further investigate this approach and discuss a range of methods based on NNMA generalized to Bregman divergences, for which we propose new regularizations: a weighted mutual decorrelation penalty and a time smoothness objective. We then propose an objective method of evaluating the note activity in the context of multipitch estimation and use this measure to train the parameters of the proposed regularized harmonic NNMA.

2. NONNEGATIVE MATRIX APPROXIMATION

For the purpose of clarity, the following matrix notation is used throughout this article. $|\mathbf{A}| = \sum_{i,j} A_{i,j}$ is shorthand for sum of all matrix elements (1-norm in case of a nonnegative matrix), and $(\mathbf{A} \odot \mathbf{B})_{i,j} = A_{i,j}B_{i,j}$ is the Hadamard (element-wise) product between two matrices. Unless stated otherwise, all other matrix operations in this paper are also element-wise, in particular matrix division $\frac{\mathbf{A}}{\mathbf{B}}$ and power $\mathbf{A}^p$.

A Bregman divergence is defined as:

$$D_\varphi(\mathbf{X},\mathbf{Y}) = |\varphi(\mathbf{X}) - \varphi(\mathbf{Y}) - \varphi'(\mathbf{Y})(\mathbf{X}-\mathbf{Y})|. \quad (3)$$

where $\varphi : \mathscr{R} \to \mathscr{R}$ is a convex *generating function* with a continuous first derivative. NNMA is an optimization problem with the penalty function being a Bregman divergence between the data $\mathbf{X}$ and its approximation $\mathbf{AS}$ and with a constraint of nonnegativity [6]. It can be easily solved using the Karush-Kuhn-Tucker conditions [13]. By minimizing the divergence between the data $\mathbf{X}$ and its approximation $\mathbf{AS}$ we obtain a pair of multiplicative update rules that, used alternately, leads to the optimal factorization:

$$\mathbf{S} \leftarrow \mathbf{S} \odot \frac{\mathbf{A^T}\left(\mathbf{X} \odot \varphi''(\mathbf{AS})\right)}{\mathbf{A^T}\left(\mathbf{A} \odot \varphi''(\mathbf{AS})\right)}, \quad (4)$$

$$\mathbf{A} \leftarrow \mathbf{A} \odot \frac{(\mathbf{X} \odot \varphi''(\mathbf{AS}))\mathbf{S^T}}{(\mathbf{AS} \odot \varphi''(\mathbf{AS}))\mathbf{S^T}}. \quad (5)$$

In practice, the I-divergence version of NNMA is usually preferred as the distortion measure, as it yields sparser representations. However, it would be beneficial to investigate if it truly gives the best results in the context of multipitch analysis. We have defined a subset of the Bregman family of distributions, which includes all commonly used distortion measures. An *r-divergence* is a divergence generated by a function, which second derivative is of the following form:

$$\varphi''(x) = x^{-r}, \quad (6)$$

where r is the shape control parameter. The resulting divergence is in practice identical (when $r = 2 - \beta$) to the beta divergence proposed by Kompass in [7] with only two small small differences. Firstly, there is no single equation for an r-divergence because of the inherent ambiguity of integration. This means that it is slightly more general, e.g. we can get both KL-divergence $D(x,y) = x \log \frac{x}{y}$ and the I-divergence $D(x,y) = x \log \frac{x}{y} - x + y$ for $r = 1$ (and the resulting NNMA algorithm would be identical for both), but only the I-divergence for $\beta \to 1$. Secondly, we avoid definition problems for the limiting cases of $\beta \to 0$ and $\beta \to 1$. Four important divergences belong to that family: Euclidean distance for $r = 0$, the KL- and I-divergence for $r = 1$ and the Itakura-Saito divergence for $r = 2$, which are commonly used for signal analysis. For r-divergence the NNMA algorithm becomes simply:

$$\mathbf{A} \leftarrow \mathbf{A} \odot \frac{(\mathbf{X} \odot (\mathbf{AS})^{-r})\mathbf{S}^T}{((\mathbf{AS})^{1-r})\mathbf{S}^T}, \quad (7)$$

$$\mathbf{S} \leftarrow \mathbf{S} \odot \frac{\mathbf{S}^T(\mathbf{X} \odot (\mathbf{AS})^{-r})}{\mathbf{S}^T((\mathbf{AS})^{1-r})}, \quad (8)$$

and for $r = 0$ and $r = 1$ these equations become identical to the ones presented in [8].

A quick analysis of multipitch estimation results for different values of r is depicted in Fig. 1 and shows that it is not the commonly used divergences ($r \in 0,1,2$), for which best results are obtained. Instead, $r \cong 1.7$ yields the best F-measure for a regular NNMA, and $r = 0.948$ is found to be optimal for the proposed regularized NNMA (see section 5).

3. CONSTRAINING

Using an unconstrained basis matrix poses a series of problems. Basis vectors need to be analyzed and assigned to a particular pitch prior to the analysis of the note activity matrix, which introduces additional errors to the process (compare Figure 4). However, because note events do not occur sparsely and independently, and their spectra change greatly over time, using an unconstrained basis usually results in basis vectors that do not even have a harmonic structure, making the pitch estimation difficult or impossible. Furthermore, results for an unconstrained basis are very different each time the algorithm is run, and thus very difficult to compare and evaluate. That is why we firmly believe that a harmonic basis matrix with vectors constrained to harmonic structures strictly corresponding to notes (of, for instance, the diatonic scale) is a must when it comes to multipitch analysis. Analysis of the note activity matrix in this case is straightforward, as each row contains amplitudes of a single note.

Basis harmonicity can be achieved in three ways. We can either: use a fixed harmonic basis vectors (i.e. only use eq. 4), use a basis matrix pretrained on solo instrument data, or adapt the harmonic structure to the data. In the first approach we use an artificial harmonic spectra with partials' amplitudes decreasing exponentially with frequency. It would seem like an oversimplification, but, as we will see later, this method yields very good results, especially when additional penalties are used, and the overfitting present in the other two methods is avoided. In the second approach, we use averaged note spectra obtained from the recordings of piano taken from the RWC database, which gives better results than the first method, but the performance drops slightly when different instrument is used.

In the third approach, proposed by us in [10], we use an artificial harmonic basis from the first method and adapt it in a way that changes only the partials' amplitudes, leaving the overall harmonic structure intact. This can be easily achieved without modifying the existing algorithm, because zero-valued elements of basis vectors remain at zero throughout the learning process due to its multiplicative nature. We could therefore initialize the basis to have zeros everywhere but at the positions of fundamentals of notes from a specific range of the 12-TET (Twelve-tone Equal Temperament) scale and at their harmonics, thus constraining the solution space to harmonic factorizations only.

4. REGULARIZATIONS

NNMA can be extended to include additional penalties on both matrices. In this case, instead of minimizing a Bregman divergence, the following objective function is minimized

$$D_\varphi(\mathbf{X}, \mathbf{AS}) + \alpha(\mathbf{A}) + \beta(\mathbf{S}), \tag{9}$$

where α and β are the penalty functions. The update rules for NNMA become

$$\mathbf{A} \leftarrow \mathbf{A} \odot \frac{(\mathbf{X} \odot \varphi''(\mathbf{AS}))\mathbf{S^T}}{(\mathbf{AS} \odot \varphi''(\mathbf{AS}))\mathbf{S^T} + \nabla_\mathbf{A}\alpha(\mathbf{A})}, \tag{10}$$

$$\mathbf{S} \leftarrow \mathbf{S} \odot \frac{(\mathbf{X} \odot \varphi''(\mathbf{AS}))\mathbf{S^T}}{(\mathbf{AS} \odot \varphi''(\mathbf{AS}))\mathbf{S^T} + \nabla_\mathbf{S}\beta(\mathbf{S})}. \tag{11}$$

This allows the user to have greater impact on the resulting factorization. However, caution must be exercised when designing these additional penalty functions, as they might cause the solution to become negative and make the algorithm unstable. Nevertheless, in our experience, using only penalties with positive derivative led to a stable algorithm. Among the note activity matrix penalties used most successfully by us, are: the sparseness and the cross-correlation penalties, and the time smoothness objective.

To obtain sparser note activities we employ the l_p-norm with $p < 2$:

$$\beta_1(\mathbf{S}) = \mu_1|\mathbf{S}^p|, \tag{12}$$

$$\nabla_\mathbf{S}\beta_1(\mathbf{S}) = \mu_1 p\mathbf{S}^{p-1}. \tag{13}$$

The cross-correlation penalty can be used to decrease the crosstalk between activities of different notes. The penalty function is defined as:

$$\beta_2(\mathbf{S}) = \mu_2 \sum_{i,j,k} W_{i,j} S_{i,k} S_{j,k} = \mu_2|\mathbf{W} \odot (\mathbf{SS}^T)|, \tag{14}$$

where $\mathbf{W}$ is a weighting matrix. In order to penalize only cross-correlation between different notes, we set $W_{i,i} = 0$. Also, the weights will usually only depend on the interval between the notes and the weighting matrix will become circulant. In this case we simply get:

$$\nabla_\mathbf{S}\beta_2(\mathbf{S}) = 2\mu_2\mathbf{WS} \tag{15}$$

By using this penalty we can also decrease the number of the most common pitch detection errors: octave errors (by increasing all weights $W_i \equiv 0 \pmod{12}$), major third errors (by increasing all $W_i = 4$) and perfect fifth errors (by increasing all $W_i = 7$). An example of a weighting matrix constructed in this manner is presented in Fig. 3a.

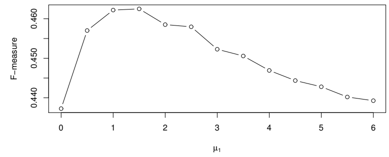

(a) Note activity sparsity, $r = 1$

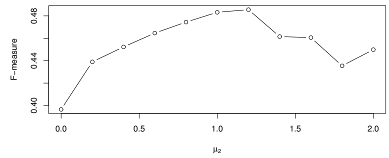

(b) Note activity decorrelation, $r = 1$

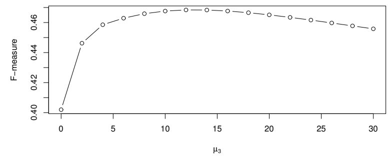

(c) Temporal smoothness, $r = 0$

Figure 2: Accuracy of multipitch analysis when using different additional objectives. A fixed harmonic basis was used and the all other regularization parameters were kept at zero.

A very similar penalty can be used to encourage temporal smoothness in a way quite similar to the one presented in [14], but using less complicated penalty function:

$$\beta_3(\mathbf{S}) = -\mu_3 \sum_{i,j,k} V_{i,j} S_{k,i} S_{k,j} = -\mu_3|\mathbf{V} \odot (\mathbf{S}^T\mathbf{S})|, \tag{16}$$

where $\mathbf{V}$ is a weighting matrix. As with the note decorrelation penalty, using a circulant matrix with nullified main diagonal leads to a simple derivative:

$$\nabla_\mathbf{S}\beta_3(\mathbf{S}) = -2\mu_3\mathbf{SV}. \tag{17}$$

As mentioned before, using regularizations with negative derivative may lead to instability, so we used $\exp(\beta_3(\mathbf{S}))$ in place of 17, which should lead to equivalent solutions thanks to monotonicity of the exponential function. An example of weighting matrix $\mathbf{V}$ is depicted in Fig. 3b.

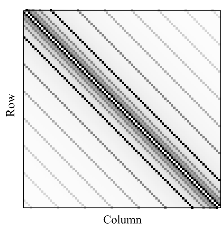

(a)

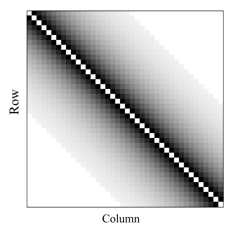

(b)

Figure 3: Circulant weighting decorrelation matrices: (a) a matrix that penalizes cross-correlation between activities of close notes and between notes in a common harmonic relation (octave 1:n, major third 5:4 and perfect fifth 3:2), (b) a matrix that encourages temporal smoothness; an exponential smoothness profile was used.

5. RESULTS

Different methods of evaluating multipitch analysis results have been proposed by researchers, but we would like to look at the problem in a somewhat narrower context of NNMA algorithms, and focus our attention only on the resulting note activity matrix. We therefore propose to directly compare this matrix with a ground truth matrix created from presumingly available MIDI data.

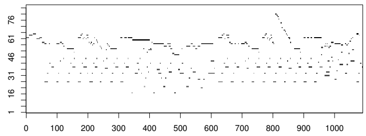

(a) Reference MIDI data

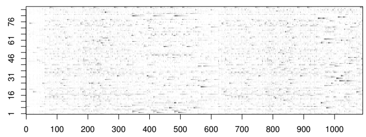

(b) Unconstrained and unpenalized NNMA

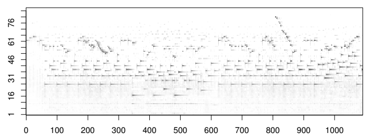

(c) Harmonically constrained NNMA

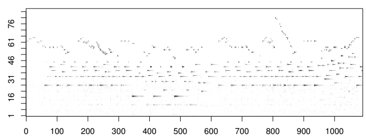

(d) Harmonically constrained and penalized NNMA

Figure 4: Original MIDI information and multiple pitch estimation results for different variants of spectrogram factorizations obtained for Chopin's *Nocturne in B flat minor, Op. 9, No. 1* with maximal polyphony of 6 simultanous notes. Horizontal axis correspond to time and the vertical axis to piano key number (exluding subfig. (d), for which there is no correspondance between basis vector number and piano key).

We have generalized precision (P) and recall (R) for real-valued data:

$$P = \left(\sum_{t,n} tp_{t,n}\right)\left(\sum_{t,n} tp_{t,n} + fp_{t,n}\right)^{-1}, \tag{18}$$

$$R = \left(\sum_{t,n} tp_{t,n}\right)\left(\sum_{t,n} tp_{t,n} + fn_{t,n}\right)^{-1}, \tag{19}$$

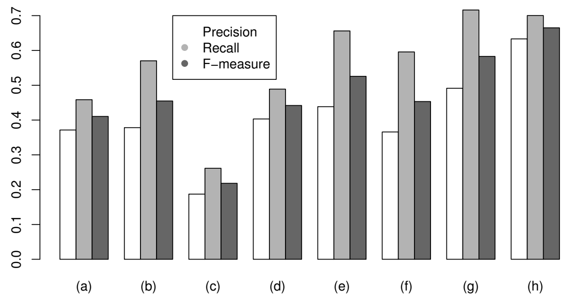

Figure 5: Results of multipitch analysis of Chopin's *Nocturne in B flat minor, Op. 9, No. 1* obtained for different harmonic bases and NNMA variants. (a) Fixed harmonic basis, (b) basis pretrained on RWC piano recordings, (c) adaptive harmonic basis, (d) adaptive harmonic basis with additional regularizations, (e) fixed harmonic basis with regularizations, (f) pretrained basis and regularizations, (g) geometric average of all methods, (h) averaged activities after simple median filtering.

where tp is the *true positive*, fp is the *false positive* and fn is the *false negative*. They are defined as:

$$tp_{n,t} = \begin{cases} \mathbf{S}_{n,t}/\mathbf{G}_{n,t} & \text{if } \mathbf{G}_{n,t} \neq 0 \\ 0 & \text{if } \mathbf{G}_{n,t} = 0 \end{cases}, \tag{20}$$

$$fp_{n,t} = \begin{cases} \mathbf{S}_{n,t} & \text{if } \mathbf{G}_{n,t} = 0 \\ 0 & \text{if } \mathbf{G}_{n,t} \neq 0 \end{cases}, \tag{21}$$

$$fn_{n,t} = \begin{cases} \mathbf{G}_{n,t} - \mathbf{S}_{n,t}/\mathbf{G}_{n,t} & \text{if } \mathbf{G}_{n,t} \neq 0 \\ 0 & \text{if } \mathbf{G}_{n,t} = 0 \end{cases}, \tag{22}$$

where n, t are the indices of elements from the note activity matrix corresponding to the note number and time, respectively.

All of above definitions require that both the note activity matrix $\mathbf{S}$ and the ground truth matrix were normalized to the range $[0,1]$. Intuitively, when we notice that $\sum_{t,n} tp_{t,n}$ is the amount of correctly identified note activity, $\sum_{t,n} tp_{t,n} + fp_{t,n}$ is the total detected note activity, and $\sum_{t,n} tp_{t,n} + fn_{t,n}$ is the amount of true note activity, it follows that the precision is a measure of how much of the detected note activity matches the ground truth data, and the recall is a measure of how well the note activity is detected. We can now use the standard F-measure definition:

$$F = \frac{2PR}{P+R}. \tag{23}$$

In our experiements we have used a dataset of 216 MIDI sequences recorded on an electric piano [5] and audio synthesized frmo them using a realistic and accurate Steinway Model-C grand piano sound font[1], which is a setup similar to the one used in the MIREX competition piano transcription task. The dataset was divided into 196 files used for training the parameters and 20 files used for evaluation. During the training, parameters values r, μ_1, μ_2 and μ_3 were randomized in each iteration from within the range of $[0,3]$, and the resulting note activity matrices were evaluated using the proposed F-measure. The best set of parameter values ($r = 0.948$, $\mu_1 = 4.643$, $\mu_2 = 0.185$ and $\mu_3 = 2.856$) was found and used in all further experiments.

Fig. 2 shows how the changes in the F-measure of the note activity matrix for different values of μ_1, μ_2 and μ_3. In

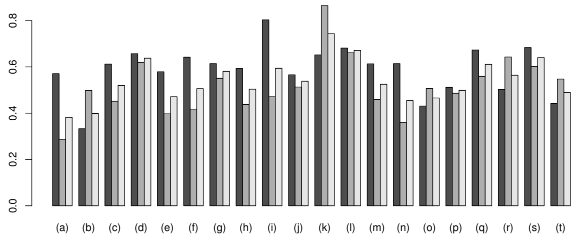

Figure 6: Results obtained for the randomly selected 20 pieces from the Chopin dataset and parameters obtained by training.

each case one coefficient was changed while the other were fixed at zero. A gain in accuracy when using the proposed penalties is evident. Fig. 5 depicts results obtained for different variants of spectrogram factorizations performed on the same file. We had our greatest expectations in the adaptive harmonic basic approach, surprisingly however, the adaptive harmonic basis gave the poorest results if the additional regularizations were not used (Fig. 5c), and even with these regularizations (Fig. 5d) the penalized fixed harmonic approach (Fig. 5e) gave slightly better results. This poor performance can be explained as a result of overfitting. Combining all 3 approaches by taking a geometric mean of all 3 note activity matrices (geometric mean was chosen in order to boost the values common for all matrices and diminish random false positives) resulted in a 5% boost in the F-measure (Fig. 5g). Results obtained for the 20 pieces selected out of the dataset is presented in Fig. 6.

6. CONCLUSION

We have discussed different ways in which the NNMA can be extended to better fit the task of multipitch analysis: by using a parameterized distortion measure, by enforcing harmonic structure on the basis matrix and by introducing additional regularizations. These modifications have additional benefit of parameterizing the method, allowing for fine tuning to fit specific data and optimize its behavior. We have also proposed an objective method of evaluating the NNMA results directly, without having to perform the note detection, and used it to train the parameters of the proposed method on a dataset of piano MIDI data. The results obtained for the proposed regularized and harmonically-constrained NNMA were better than those of the other modifications of NNMA we have tested and can be further improved by postprocessing the note activities, e.g. with median filtering.

REFERENCES

[1] Steinway model-c grand piano soundfont, `http://www.pianosounds.com/`, December 2008.

[2] Mirex 2008 multiple fundamental frequency estimation and tracking results, `http://www.music-ir.org/mirex/2008/`, February 2009.

[3] S. Abdallah and M. Plumbley. Polyphonic Music Transcription by Non-negative Sparse Coding of Power Spectra. *Proc. 5th International Conference on Music Information Retrieval*, 510:318–325, 2004.

[4] S. Abdallah and M. Plumbley. Unsupervised Analysis of Polyphonic Music by Sparse Coding. *IEEE Trans. Neural Netw.*, 17(1):179–196, 2006.

[5] N. Carter. The ultimate chopin midi collection, `http://www.geocities.com/Vienna/2217/midi.htm`, February 2009.

[6] I. Dhillon and S. Sra. Generalized Nonnegative Matrix Approximations with Bregman Divergences. *Advances in Neural Information Processing Systems*, 18:283, 2006.

[7] R. Kompass. A generalized divergence measure for nonnegative matrix factorization. *Neural computation*, 19(3):780–791, 2007.

[8] D. Lee and H. Seung. Algorithms for Non-negative Matrix Factorization. *Advances in Neural Information Processing Systems*, pages 556–562, 2001.

[9] P. Paatero and U. Tapper. Positive Matrix Factorization: A Non-negative Factor Model with Optimal Utilization of Error Estimates of Data Values. *Environmetrics*, 5(2):111–126, 1994.

[10] S. Raczyński, N. Ono, and S. Sagayama. Multipitch analysis with harmonic nonnegative matrix approximation. In *Proceedings of the 8th International Conference on Music Information Retrieval (ISMIR 2007), Vienna, Austria*, 2007.

[11] M. Schmidt and M. Mørup. Sparse Non-negative Matrix Factor 2-D Deconvolution for Automatic Transcription of Polyphonic Music. In *Proc. 6th International Symposium on Independent Component Analysis and Blind Signal Separation*, 2006.

[12] P. Smaragdis and J. Brown. Non-negative matrix factorization for polyphonic music transcription. In *Applications of Signal Processing to Audio and Acoustics, 2003 IEEE Workshop on.*, pages 177–180, 2003.

[13] S. Sra and I. Dhillon. Nonnegative matrix approximation: Algorithms and applications. Technical report, Technical report, Dept. of Computer Sciences, The Univ. of Texas at Austin, 2006.

[14] K. Wilson, B. Raj, and P. Smaragdis. Regularized Non-negative Matrix Factorization with Temporal Dependencies for Speech Denoising.