

NONNEGATIVE MATRIX FACTORIZATION BASED SELF-TAUGHT LEARNING WITH APPLICATION TO MUSIC GENRE CLASSIFICATION

Konstantin Markov

Human Interface Lab.,
University of Aizu,
Fukushima, Japan

Tomoko Matsui

Dep. of Statistical Modeling,
Institute of Statistical Mathematics,
Tokyo, Japan

ABSTRACT

Availability of large amounts of raw unlabeled data has sparked the recent surge in semi-supervised learning research. In most works, however, it is assumed that labeled and unlabeled data come from the same distribution. This restriction is removed in the self-taught learning approach where unlabeled data can be different, but nevertheless have similar structure. First, a representation is learned from the unlabeled data via non-negative matrix factorization (NMF) and then it is applied to the labeled data used for classification. In this work, we implemented this method for the music genre classification task using two different databases: one as unlabeled data pool and the other for supervised classifier training. Music pieces come from 10 and 6 genres for each database respectively, while only one genre is common for both of them. Results from wide variety of experimental settings show that the self-taught learning method improves the classification rate when the amount of labeled data is small and, more interestingly, that consistent improvement can be achieved for a wide range of unlabeled data sizes.

Index Terms— Music genre classification, Self-taught learning, non-negative matrix factorization, Transfer learning.

1. INTRODUCTION

A tremendous amount of music-related data has recently become available either locally or remotely over networks, and technology for searching this content and retrieving music-related information efficiently is demanded. This consists of several elemental tasks such as genre classification, artist identification, music mood classification, cover song identification, fundamental frequency estimation, and melody extraction. Essential for each task is the feature extraction as well as the model or classifier selection. Audio signals are conventionally analyzed frame-by-frame using Fourier or Wavelet transform, and coded as spectral feature vectors or chroma features extracted for several tens or hundreds of milliseconds

[1, 2, 3]. However, it is an open question how precisely music audio should be coded depending on the task kind and the succeeding classifier. Recently proposed non-negative matrix factorization approach [4] and its variants have also been applied in music processing with promising results [5, 6, 7].

For the classification, classical supervised pattern recognition approaches require large amount of labeled data which is difficult and expensive to obtain. On the other hand, in the real world, a massive amount of musical data is created day by day and various musical databases are newly composed. There may be no labels for some databases and musical genres may be very specific. Thus, recent music information retrieval research has been increasingly adopting semi-supervised learning methods where unlabeled data are utilized to help the classification task. Common assumption in this case is that both labeled and unlabeled data come from the same distribution [8] which, however, may not be easily achieved during the data collection. This restriction is alleviated in the transfer learning framework [9] which allows the domains, tasks, and distributions used in training and testing to be different. Utilizing this framework and the semi-supervised learning ideas, the recently proposed self-taught learning approach [10] is a good fit for the music genre classification task. First, an unsupervised learning algorithm finds high-level representation of the inputs given only unlabeled data. Then, using this representation, the labeled data which are assumed to be few are transformed into new feature vectors. Finally, classical supervised classifier is trained using the new features. There are several possible methods for unsupervised learning of the high-level representations and in our previous work [11] we used sparse representations [12, 13]. It can be considered that this method possibly captures sparse and key features for all classes.

In order to manage various types of new musical databases, in this work we use the self-taught learning algorithm based on non-negative matrix factorization for unsupervised high-level feature learning and investigate a genre classification method for new musical database using an existing one which has a well-balanced design for all genres.

2. THE SELF-TAUGHT LEARNING ALGORITHM

A classification task is considered with small labeled training data set $\mathcal{X}^l = \{\mathbf{x}_i^l, i = 1, \dots, M\}$ drawn i.i.d. from an unknown distribution \mathcal{D} . Each $\mathbf{x}_i^l \in \mathcal{R}^d$ is an input feature vector which is assigned a class label $y_i \in \mathcal{Y} = \{1, \dots, C\}$. In addition, a larger unlabeled training data set $\mathcal{X}^u = \{\mathbf{x}_i^u, i = 1, \dots, K\}$ is available, which is assumed only to be of the ‘‘same type’’ as \mathcal{X}^l and may not be associated with the class labels \mathcal{Y} and distribution \mathcal{D} . Obviously, in order \mathcal{X}^u to help the classification of the labeled data, it should not be totally different or unrelated.

The main idea of the self-taught learning approach is to use the unlabeled data to learn in an unsupervised way slightly higher level representation of the data [10]. In other words, to discover some hidden structures in the data which can be considered as basic building blocks. For example, if the data represent images, the algorithm would find simple elements such as edges, curves, etc., so that the image can be represented in terms of these more abstract, higher level features. Once learned, this representation is applied to the labeled data \mathcal{X}^l resulting in a new set of features which lighten the supervised learning task.

To learn the higher level representation, we use the non-negative matrix factorization method. It decomposes the unlabeled data matrix $\mathbf{X}^u = [\mathbf{x}_1^u, \mathbf{x}_2^u, \dots, \mathbf{x}_K^u] \in \mathcal{R}^{d \times K}$ as a product of two matrices $\mathbf{W}^u = [\mathbf{w}_1^u, \mathbf{w}_2^u, \dots, \mathbf{w}_n^u] \in \mathcal{R}^{d \times n}$ and $\mathbf{H}^u = [\mathbf{h}_1^u, \mathbf{h}_2^u, \dots, \mathbf{h}_K^u] \in \mathcal{R}^{n \times K}$ having only non-negative elements. The decomposition considered here is approximative in nature, i.e.,

$$\mathbf{X}^u \approx \mathbf{W}^u \mathbf{H}^u \quad (1)$$

or equivalently in a vector form $\mathbf{x}_i^u \approx \sum_{j=1}^n h_{i,j}^u \mathbf{w}_j^u$, where \mathbf{H}^u is the mixing (or activations) matrix and \mathbf{W}^u is the bases matrix. Since only additive combinations of these bases are allowed, the nonzero elements of \mathbf{W}^u and \mathbf{H}^u are all positive. Thus in such decomposition no subtractions can occur. For these reasons, the non-negativity constraints are compatible with the intuitive notion of combining components to form a whole signal, which is how NMF learns a parts-based (high level) representations.

In contrast to the sparse coding method, the NMF does not assume explicitly or implicitly sparseness or mutual statistical independence of components, however it can provide sparse decomposition [4].

3. NMF DECOMPOSITION

The most of known and used adaptive multiplicative algorithms for NMF are based on two cost functions: Square Euclidean distance expresses by Frobenius norm,

$$D_F(\mathbf{W}, \mathbf{H}) = \frac{1}{2} \|\mathbf{X} - \mathbf{W}\mathbf{H}\|_F^2 \quad (2)$$

which is optimal for Gaussian distributed approximation error, and the generalized Kullback-Leibler divergence,

$$D_{KL}(\mathbf{W}, \mathbf{H}) = \sum_{i,j} (x_{ij} \log(\frac{x_{ij}}{[\mathbf{W}\mathbf{H}]_{ij}}) - x_{ij} + [\mathbf{W}\mathbf{H}]_{ij}). \quad (3)$$

Although both functions are convex in \mathbf{W} and \mathbf{H} only, they are not convex in both variables together. Thus, we can only expect a maximization algorithm to find a local minimum. A good compromise between speed and ease of implementation have been proposed in [14] and consists of iterative application of the following update rules:

$$h_{ij} \leftarrow h_{ij} \frac{[\mathbf{W}^T \mathbf{X}]_{ij}}{[\mathbf{W}^T \mathbf{W}\mathbf{H}]_{ij}} \quad (4)$$

$$w_{ij} \leftarrow w_{ij} \frac{[\mathbf{X}\mathbf{H}^T]_{ij}}{[\mathbf{W}\mathbf{H}\mathbf{H}^T]_{ij}} \quad (5)$$

when Frobenius norm (Eq.2) is chosen as objective function.

4. FEATURE GENERATION USING UNLABELED DATA

After learning the basis vectors \mathbf{w}_j^u from the unlabeled training data \mathcal{X}^u as described in the previous section, we can use them to obtain activations for the labeled data \mathcal{X}^l . This procedure can be viewed as a non-linear mapping or transformation of vectors $\mathbf{x}_i^l \in \mathcal{R}^d$ into vectors $\mathbf{h}_i^{l*} \in \mathcal{R}^n$, which will be further used as features in the supervised classification task. The new features are computed by solving the following optimization problem:

$$\mathbf{h}_i^{l*} = \arg \min_{\mathbf{h}} \|\mathbf{x}_i^l - \sum_{j=1}^n h_j \mathbf{w}_j^u\|_2^2 \quad (6)$$

This is a convex least squares task which is the same as the optimization of (2) with fixed bases \mathbf{w}_j and can be solved in the same way. Vectors \mathbf{h}_i^{l*} approximate labeled data \mathbf{x}_i^l as a linear combination of the bases which, however, are learned using the unlabeled data \mathbf{x}_i^u . Considering that basis vectors represent information discovered automatically from different data, and that activation vectors \mathbf{h}_i^{l*} in fact use this information to encode the labeled data, the self-taught learning approach can be viewed as an instance of *knowledge transfer* between tasks [9].

Finally, using labels y_i and the new feature vectors \mathbf{h}_i^{l*} we can train standard supervised classification models, such as SVMs. During testing, for each test vector an activation vector is computed using Eq.(6) which is then used as input to the trained classifier. Following three steps summarize the self-taught learning algorithm:

1. Compute basis vectors \mathbf{w}_j^u using unlabeled data \mathbf{x}_i^u by solving the optimization problem,

$$\mathbf{W}^u = \arg_{\mathbf{W}} \min_{\mathbf{W}, \mathbf{H}} D_F(\mathbf{W}, \mathbf{H}) \quad (7)$$

2. Fix the basis vectors and obtain activation vectors \mathbf{h}_i^{l*} for each labeled data vector \mathbf{x}_i^l using Eq.(6).
3. Use activation vectors \mathbf{h}_i^{l*} as new labeled features to train standard supervised classifier. Transform each test vector into an activation vector (using Eq.(6) again) and apply the classifier to obtain its label.

5. EXPERIMENTS

5.1. Databases

As unlabeled data we used the GTZAN collection of musical pieces [1]. It consists of 1000 30-second audio clips, each belonging to one of the following ten genres: Classical, Country, Disco, Hip-Hop, Jazz, Rock, Blues, Reggae, Pop and Metal. There are 100 clips per genre and all of them have been down-sampled to 22050 Hz. The other database which we used as labeled data is the corpus used in the ISMIR 2004 audio contest [15]. It contains of 729 whole tracks for training, but since the number of tracks per genre is non-uniform, the original nine genres are usually mapped in the following six classes: Classical, Electronic, Jazz-Blues, Metal-Punk, Rock-Pop and World. Another 729 tracks are used for testing. Note that the only common class between the two databases is the ‘‘Classical’’ genre.

Audio data from both databases are divided into 5 sec. pieces which were further randomly selected in order to make several training sets with different amount of data, keeping the same number of such pieces per genre. Table 1 summarizes the contents of the data sets we used in our experiments. For example, set GT-50 has 50 randomly selected 5 sec. pieces per genre, 500 pieces in total or 0.69 hours of music. All sets are constructed such that each larger set contains all the pieces from the smaller set. The test set consists of 250 pieces per genre randomly selected from the ISMIR 2004 test tracks.

Table 1. Data sets used in the experiments.

GTZAN database			ISMIR 2004 database		
Set	pieces	hours	Set	pieces	hours
GT-50	500	0.69	IS-20	120	0.17
GT-100	1000	1.39	IS-50	300	0.42
GT-250	2500	3.47	IS-250	1500	2.08
GT-500	5000	6.95	IS-500	3000	4.16

5.2. Audio signal preprocessing

When it comes to feature extraction for music information processing, in contrast to the case of speech where the MFCC

is dominant, there exists wide variety of approaches - from carefully crafted multiple music specific tonal, chroma, etc. features to single and simple ‘‘don’t care about the content’’ spectrum. In our experiments, we used spectral representation tailored for music signals, such as Constant-Q transformed (CQT) FFT spectrum. The main property of this transform is the log-like frequency scale where the consecutive musical notes are linearly spaced [16].

The CQT transform is applied to FFT spectrum vectors computed from 23.2ms (512 samples) frames with 50% overlap in a way that there are 12 Constant-Q filters per octave resulting in a filter-bank of 89 filters which covers the whole bandwidth of 11025 Hz. The filter-bank outputs of 20 consecutive frames are further stacked into a 1780 (89x20) dimensional super-vector which is used in the experiments. This is the same as to have a 20 frame time-frequency spectrum image. There is a overlap of 10 frames between such two consecutive spectrum images. This way, each 5 sec. music piece is represented by 41 spectrum images or super-vectors.

5.3. Bases learning

For each dataset given in Table 1 we learned several basis vector sets or dictionaries as they are often called in the literature. The sets sizes are: 100, 200, 300 and 500. Since the input vectors dimension is 1780, this is equivalent to a dimensionality reduction. This is justified by the fact that our super-vectors are highly redundant and that basis vectors actually represent higher level spectral image features, not just arbitrary projection directions.

Before bases learning, all data are pooled and randomly shuffled. Then, about 5000 iterations of the learning algorithm (see Section 2) are performed. Figure 1 shows the 100 basis vectors learned from GT-50 as 89x20 pixel spectrum images. Specific spectrum shapes and transitions are clearly captured by these bases.




Fig. 1. Example of learned basis vectors shown as spectrum images.

5.4. Supervised classification

After all labeled training data, sets IS-20, IS-50, IS-250 and IS-500, have been transformed into activation vectors for each dictionary learned from each unlabeled data set, we obtained in total 64 (4 labeled datasets x 4 dictionary sizes x 4 unlabeled datasets) labeled training data sets. Then, using the LIBSVM tool, we learned 64 SVM classifiers each consisting of 6 SVMs trained in one-versus-all mode. The SVM input vectors were linearly scaled to fit the [0,1] range. Linear kernel was used as distance measure and the SVMs were trained to produce probabilistic outputs.

During testing, each 5 sec. musical piece represented by 41 feature (activation) vectors is considered as a sample for classification. Outputs of all genre specific SVMs are aggregated (summed in the log domain) and the label of the maximum output is taken as the classification result.

5.5. Results

In order to assess the effect of the self-taught learning, we need a performance comparison with a system build under the same conditions but without unlabeled data. We will refer to this system as a baseline. In this case, the basis vectors are learned using labeled training data \mathcal{X}^l instead of the unlabeled \mathcal{X}^u . Then, the activations are obtained in the same way as if the bases were learned from the unlabeled data. Table 2 shows the baseline system performance for several dictionary sizes and different amounts of training data.

As can be expected, the more training data are available the better is the classification accuracy. It also increases with the size of the dictionary and this seems to be independent of the amount of training data.

Tables 3-6 show the absolute difference of the classification accuracy between the self-taught learning based system and the baseline when basis vectors are learned with different amounts of unlabeled data: GT-50 to GT-500 respectively.

The results clearly show that for the smallest labeled training set, IS-20, consistent improvement is achieved with dictionary sizes 300 and 500 no matter how much unlabeled data are used for bases learning. As the amount of labeled training data increases, the effect of the self-taught learning gradually

Table 2. Baseline classification accuracy. Bases are learned from the labeled training data.

Training Set	Dictionary size			
	100	200	300	500
IS-20	46.7%	45.7%	47.4%	50.3%
IS-50	50.4%	52.0%	53.7%	55.4%
IS-250	52.8%	54.2%	55.0%	57.0%
IS-500	50.8%	53.8%	55.7%	57.4%

Table 3. Absolute improvement wrt baseline when bases are learned from the unlabeled GT-50 dataset.

Training Set	Dictionary size			
	100	200	300	500
IS-20	0.47%	1.27%	-0.40%	-0.33%
IS-50	0.13%	-0.80%	-0.47%	0.13%
IS-250	-0.60%	-0.40%	0.13%	-0.27%
IS-500	-0.00%	-0.33%	-0.33%	0.07%

Table 4. Absolute improvement wrt baseline when bases are learned from the unlabeled GT-100 dataset.

Training Set	Dictionary size			
	100	200	300	500
IS-20	0.20%	1.67%	0.87%	-0.87%
IS-50	0.07%	-0.33%	-0.13%	0.73%
IS-250	-0.13%	0.07%	0.67%	-0.07%
IS-500	0.40%	0.13%	0.20%	-0.00%

Table 5. Absolute improvement wrt baseline when bases are learned from the unlabeled GT-250 dataset.

Training Set	Dictionary size			
	100	200	300	500
IS-20	1.07%	0.93%	0.53%	-1.47%
IS-50	0.33%	-0.73%	-0.13%	-0.40%
IS-250	-0.53%	0.40%	0.33%	-0.47%
IS-500	0.27%	-0.47%	0.53%	-0.73%

Table 6. Absolute improvement wrt baseline when bases are learned from the unlabeled GT-500 dataset.

Training Set	Dictionary size			
	100	200	300	500
IS-20	0.73%	1.27%	0.60%	-0.27%
IS-50	0.20%	-0.60%	-0.53%	0.13%
IS-250	-0.00%	0.60%	0.80%	0.13%
IS-500	0.47%	-0.27%	-0.07%	-0.20%

decreases, which is an expected result.

6. CONCLUSION

In this study, we investigated the effect of the NMF based self-taught learning algorithm when applied to the music genre classification task. Results of the experiments conducted under various conditions confirmed that when small size of labeled training is available, higher level features extracted from a bigger unlabeled dataset in an unsupervised manner via sparse coding can indeed help the supervised classification task.

7. REFERENCES

- [1] G. Tzanetakis and P. Cook, "Musical genre classification of audio signals," *IEEE Trans. ASLP*, vol. 10, no. 5, pp. 293–302, 2002.
- [2] R. Typke, F. Wiering, and R. C. Veltkamp, "A survey of music information retrieval systems," in *Proc. International Conference on Music Information Retrieval*, 2005, pp. 153–160.
- [3] M. Casey, R. Veltkamp, M. Goto, M. Leman, C. Rhodes, and M. Slaney, "Content-based music information retrieval: Current directions and future challenges," *Proceedings of the IEEE*, vol. 96, no. 4, pp. 668–696, 2008.
- [4] D. Lee and H. Seung, "Learning the parts of objects by non-negative matrix factorization," *Nature*, vol. 401, pp. 788–791, 1999.
- [5] P. Smaragdis and J.C. Brown, "Non-negative matrix factorization for polyphonic music transcription," in *Proc. of IEEE Workshop on Applications of Signal Proc. to Audio and Acoustics*, New Paltz, USA, 2003.
- [6] Tuomas Virtanen, "Monaural sound source separation by non-negative matrix factorization with temporal continuity and sparseness criteria," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 3, 2007.
- [7] C. Fvotte, N. Bertin, and J.-L. Durrieu, "Nonnegative matrix factorization with the itakura-saito divergence: With application to music analysis," *Neural Computation*, vol. 21, no. 3, pp. 793–830, 2009.
- [8] K. Nigam, A. McCallum, S. Thrun, and T. Mitchell, "Text classification from labeled and unlabeled documents using EM," *Machine Learning*, vol. 39, no. 2-3, pp. 103–134, 2000.
- [9] S. Pan and Q. Yang, "A survey on transfer learning," *IEEE Trans. Knowledge and Data Engineering*, vol. 22, no. 10, pp. 1349–1359, 2010.
- [10] R. Raina, A. Battle, H. Lee, B. Packer, and A. Ng, "Self-taught learning: Transfer learning from unlabeled data," in *Proc. ICML*, 2007.
- [11] K. Markov and T. Matsui, "Music genre classification using self-taught learning via sparse coding," in *Proc. ICASSP*, 2012, pp. 1929–1932.
- [12] Plumbley M., Blumensath T., Daudet L., Gribonval R., and Davies M., "Sparse representations in audio and music: From coding to source separation," *Proc. IEEE*, vol. 98, no. 6, pp. 995–1005, 2010.
- [13] M. Henaff, K. Jarrett, K. Kavukcuoglu, and Y. LeCun, "Unsupervised learning of sparse features for scalable audio classification," in *Proc. ISMIR*, 2011.
- [14] D. Lee and H. Seung, "Algorithms for non-negative matrix factorization," in *Advances in Neural Information Processing*, vol. 13. MIT Press, 2001.
- [15] P. Cano, E. Gomes, F. Gouyon, P. Herrera, M. Koppenberger, B. Ong, X. Serra, S. Streich, and N. Wack, "ISMIR 2004 audio description contest," Tech. Rep. MTG-TR-2006-02, Universitat Pompeu Fabra, 2006.
- [16] C. Schoerhuber and A. Klapuri, "Constant-Q transform toolbox for music processing," in *Proc. 7th. Sound and Music Computing Conference*, 2010.