# EXTENDED PIPELINE FOR CONTENT-BASED FEATURE ENGINEERING IN MUSIC GENRE RECOGNITION

Tina Raissi<sup>†</sup> Alessandro Tibo<sup>‡</sup> Paolo Bientinesi<sup>\*†</sup>

<sup>†</sup> RWTH Aachen University, AICES, Schinkelstr. 2, 52062 Aachen, Germany <sup>‡</sup> University of Florence, Department of Information Engineering, Via S. Marta 3, 50139 Firenze, Italy

## ABSTRACT

We present a feature engineering pipeline for the construction of musical signal characteristics, to be used for the design of a supervised model for musical genre identification. The key idea is to extend the traditional two-step process of extraction and classification with additive stand-alone phases which are no longer organized in a waterfall scheme. The whole system is realized by traversing backtrack arrows and cycles between various stages. In order to give a compact and effective representation of the features, the standard early temporal integration is combined with other selection and extraction phases: on the one hand, the selection of the most meaningful characteristics based on information gain, and on the other hand, the inclusion of the nonlinear correlation between this subset of features, determined by an autoencoder. The results of the experiments conducted on GTZAN dataset reveal a noticeable contribution of this methodology towards the model's performance in classification task.

*Index Terms*— Musical signal, genre classification, feature extraction and selection, information gain, autoencoder

# 1. INTRODUCTION

One of the current subjects of research in Computer Science and Engineering concerns the enhancement of machines with abilities which are related to the human perception of the environment. Since recently, the term machine hearing [1] is used as an umbrella to unify all applications of speech, music and environment sounds processing under a general concept. The central goal of this field is to model the hearing apparatus and its internal functionality [2]. However, in the case of speech and music, it is not possible to achieve this goal without taking into account the representation of intrinsic characteristics of the sound, developed in relation to the relative learning task. In our case, the quality of the musical signal features is evaluated with respect to the task of genre classification, a well-established, and rather controversial topic of Musical Information Retrieval.

A quick look-back at remarkable works [3] shows that the

signal representation is obtained by the extraction of physical and perceptual features, in time, frequency and cepstral domains. This solution carries out the transformation of the original input signal to a new feature space. A reduction which results to be problematic, because the conservation of the relevant traits of the former is not guaranteed. This loss of information is caused by the applied analytical and computational model, and at the same time, by the rigid order of different steps of construction of the whole system. We shifted for this reason our attention from the analytical approach to a methodological one, by presenting a new feature extraction pipeline. For some of the stages, our pipeline builds on existing methods in its isolated stages. We expanded the content-based features by adding the bottleneck layer's features of an autoencoder, which forces to learn a low dimensional representation of the data. Moreover, the selection of the most predictive attributes, based on information gain criteria, is done by using a Random Forests classifier trained on an intermediate level feature vector. The entire process is not actually designed for recovering the lost information during the reduction, but for enriching the resulting feature vector with additive knowledge, useful for the specific classification task. The mean accuracy of the classifier, trained with the output dataset of this pipeline, is improved from the 78% to 91%.

In the following sections, we first introduce the general structure of this pipeline and describe every single process; we then present an evaluation of effectiveness of the features, at both final and intermediate stages.

## 2. GENERAL PIPELINE

For the most part, the proposed approaches in the literature for the automatic music genre identification, take into consideration a two-step process of extraction and classification, performed in consecutive order. In our work, the intermediate stages of the extraction gain autonomy, and can be reached from later stages. In this section we give a concise description of the process.

As shown in Figure 1, the input of the process is a continuous audio stream. The extraction of information from the digitalized audio signal is generally carried out by algorithms

<sup>\*</sup>Financial support from the Deutsche Forschungsgemeinschaft (DFG) through grant GSC 11 is gratefully acknowledged.

operating on prefixed time window or consecutive blocks of frames. In order to obtain a vector  $F \in \mathbb{R}^2$  relative to one feature, we should perform the stages windowing, feature extraction and temporal integration twice, by using the backtrack arrow (1). The first loop is responsible for the extraction of the *short-time features*, while the second loop is concerned with the extraction of *medium-time features* and the temporal integration over the whole signal length. In both cases we have an early temporal integration [4] which uses the *Mean-Var* model [5].

In Figure 1, the arrow (3) refers to the extraction of derivatives after the first windowing step. The feature vector extracted at this level can already be used for a classification task. In contrast to many existing approaches, we use this feature vector not for the final task, but for an intermediate classifier, built by random forests. The measure of the contribution of each attribute to the prediction of the target class is calculated by summing the information gain of each attribute, at every split. The attributes with positive contribution are then selected. In the pipeline, this step is indicated as preprocessing and is followed by another feature extraction step, connected via the backtrack arrow (2). This final extraction uses an autoencoder which maps the vector of selected features to itself. The features selected from the bottleneck layer are not going to substitute the original dataset, but are added to it, as an additional information about nonlinear correlation between features. As a last step before the training phase, the dataset is normalized.



**Fig. 1**. The general pipeline for content-based feature engineering, with backtracking and forward arrows.

# 3. FEATURE EXTRACTION AND TAXONOMY

# 3.1. Content-based features

We used three main categories of content-based features, based on a taxonomy presented in [6]: Time Domain Physical, Frequency Domain Physical and Cepstral Domain Perceptual. Following the three basic requirements of musical characteristics introduced by [4], we are interested in obtaining a single value or a low-dimensional vector, from several feature observations, known also as early temporal integration method. We use the MeanVar model, which calculates mean and standard deviation of observed values within a prefixed texture window [7]. The early temporal integration procedure is illustrated in Figure 2: first, 14 features are extracted by using an analysis frame [7] of 50 milliseconds with 50% overlapping. Namely, Compactness, Energy, Entropy of Energy, Root Mean Square (RMS), Zero Crossing, 26 Mel Frequency Cepstral Coefficients (MFCC), Chroma vector corresponding to 12 semitones, the standard deviation of values of Chroma vector, 10 Linear Prediction Coefficients (LPC), Spectral Centroid, Spectral Flux, Spectral Rolloff, Spectral Spread and Spectral Variability. The computation of derivatives and a temporal integration step over both feature values and their derivatives follow. The Fraction of Low Energy Windows (FoLEW) is then determined, and a conclusive temporal integration step is performed. Finally, from the Beat Histogram the Beat Sum, Strongest Beat (SBeat), Strength of Strongest Beat (SSBeat), and their derivatives are extracted. The computation of all these features is detailed in [8].



**Fig. 2**. The temporal integration procedure relative to the extraction of one feature vector.

#### 3.2. Early temporal Integration

As a first step, the signal is divided into  $n_f$  analysis frames of length L. For every feature, a vector  $\mathbf{s} \in \mathbb{R}^{n_f}$  of shorttime values is extracted. In the second step, the mean and the standard deviation of  $n_{int} = \frac{n_f}{n_w}$  values of the vector  $\mathbf{s}$  corresponding to a texture window of length W is calculated. Consequently, for every feature we have a matrix of medium-time values, having two rows derived from concatenation of  $\mathbf{m}_{\tau} = [\mu_{\tau}, \sigma_{\tau}]^T, \forall \tau = 1, \cdots, n_w$ . The feature vector  $\mathbf{F}_k \in \mathbb{R}^2$  is the average value of every row of the mediumtime matrix, corresponding to the k-th feature.

#### 3.3. Bottleneck layer's features

The linear version of *Principal Component Analysis* (PCA) is a common learning method for analyzing and giving a lowdimensional representation of an input space [9].

A generalization of PCA can be obtained by using an *autoencoder* [10], which is a pair of stacked *neural networks*: an encoder **E** and a decoder **D**. The encoder maps an input data  $\mathbf{x} \in \mathbb{R}^n$  into a hidden vector  $\mathbf{E}(\mathbf{x}) = \mathbf{h} \in \mathbb{R}^d$  with typically  $d \ll n$ . The decoder maps **h** to an output vector  $\mathbf{x}' \in \mathbb{R}^n$ . The *autoencoder* is trained to copy the input **x** to  $\mathbf{D}(\mathbf{E}(\mathbf{x})) = \mathbf{x}'$  by using binary cross-entropy loss function:

$$\mathcal{L}(\mathbf{x}, \mathbf{D}(\mathbf{E}(\mathbf{x}))) = - ig(\mathbf{x} \log(\mathbf{x}') + (\mathbf{1} - \mathbf{x}) \log(\mathbf{1} - \mathbf{x}')ig).$$

If the data lies in a low dimensional manifold, the *autoen*coder can actually learn such representation. Concerning music, we capture the evolution of the musical content in the signal by concatenating features referred to adjacent frames. The number of feature vectors that come from musically sound frames, is reasonably a subset of the whole feature space.



Fig. 3. The architecture of a symmetric autoencoder.

The encoding and decoding functions in an *autoencoder* has respectively  $n_e$  and  $n_d$  hidden layers, and in our case  $n_e = n_d$ . In addition to just mentioned  $n_e + n_d$  hidden layers, we have also a middle layer known as *bottleneck layer* which takes as input  $\mathbf{E}(\mathbf{x})$ . An example of an autoencoder is depicted in Figure 3.

#### 3.4. Feature selection

The notion of *information gain* is relative to the average variation of information entropy due to current state's changes. Of particular interest is the application of this concept in *decision trees* [11]. Generally speaking, the process of construction of a decision tree is based on the choice of the attribute with highest IG, on whose values the dataset of every node is split. In the feature selection process based on IG, every example is considered as a vector of values containing a specific information for predicting the class. We can construct a decision tree and determine the  $IG_i$  of *i*-th attribute or feature by  $\sum_{j=1}^{|nodes|} IG_{ij}$ , where  $IG_{ij}$  is the information gain obtained by splitting on *i*-th attribute at *j*-th node. The idea of ranking  $IG_i$  values is equivalent to establishing an importance order of features, concerning their contribution to the prediction. Given a set of training examples S, with  $\frac{|S_{x_i=a}|}{|S_j|}$  we denote the fraction of examples of the *i*-th attribute having value a at *j*-th node. We define entropy as

$$H(\mathcal{S}) = -\sum_{c \in classes} p_c(\mathcal{S}) \log_2 p_c(\mathcal{S}),$$

where  $p_c(S)$  is the probability of a training example to belong to class c. Hence the information gain  $IG_{ij}$  obtained by splitting on attribute  $x_i$  at node j can be reformulated as

$$IG(\mathcal{S}_{j}, x_{i}) = H\left(\mathcal{S}_{j}\right) - \sum_{v \in values(x_{i})} \frac{\left|\mathcal{S}_{j(x_{i}=v)}\right|}{\left|\mathcal{S}_{j}\right|} H\left(\mathcal{S}_{j(x_{i}=v)}\right).$$

This function measures the difference of entropies before and after splitting on the specific attribute.

# 4. CLASSIFICATION

We consider two different classifiers as final models: the *Support Vector Machines* with both radial basis and linear kernels, and *Random Forests* for the feature selection step at intermediate stage of the pipeline. The choice of the latter, instead of a straightforward decision tree, derives from its well-known property of avoiding the overfitting problem.

#### 5. EXPERIMENT

The audio dataset is GTZAN [7], consisting of 1000 audio tracks, each 30 seconds long and associated to one of ten genres: blues, classical, country, disco, hiphop, jazz, metal, pop, reggae and rock. Exactly 100 tracks are associated to every genre. All tracks were converted in 22,050 Hz mono 16-bit wav format. We ran 10 experiments splitting the dataset into 900 samples of training set and 100 samples of test set, preserving the percentage of samples for each class. Feature selection and hyperparameter optimization of SVM are done by 10-fold cross validation on the training set. For the extraction of content-based features, we used pyAudioAnalysis and jAudio [12, 13]. The analysis frame for short-time features and the texture window for the medium-time features were set to 50 milliseconds and 1 second, respectively. In both cases we operated a 50% of overlap. After the feature selection in the preprocessing step, we dropped all feature components which resulted in no information gain. Table 1 lists the selected features with their dimensionality corresponding to the experiment with the best result.

For the architecture of the autoencoder, summarized in Table 2, we used three stacked hidden layers of size 60, 20, 60 respectively, with PReLU [14] activation, initialized using *He normal initializer* [15]. The first two hidden layers were followed by a Dropout [16] layer with probability 0.2. Finally an output layer of size 190 with sigmoid activation, initialized using *He uniform initializer*, was stacked at the end of

the network. The model was trained by minimizing the binary cross-entropy error loss. The autoencoder was fed with 900 training feature vectors of size 190, rescaled in [0, 1]. We ran 100 epochs of the Adadelta [17] optimizer with learning rate 1.0,  $\rho = 0.95$ ,  $\epsilon = 1e - 08$  without decay factor on minibatches of size 32. Even though we did not have sufficient resources for an automatic hyperparameter optimization, the proposed structure minimizes the loss function on the training set. The final dataset consisted of content-based features of Table 1 augmented with bottleneck layer's features. Before the classification step the whole dataset was rescaled in [0, 1]. The best final classifier, trained over this dataset, turned out to be the SVM with radial basis kernel. The hyperparameters of SVM resulted to be  $\gamma = 2^{-6}$  and C = 4.

**Table 1.** The content-based features and their dimensionality. For every feature, mean (M), standard deviation (SD), mean of derivatives ( $\Delta$ M) and standard deviation of derivatives ( $\Delta$ SD) are reported. The symbol " $\star$ " indicates the elimination of feature components (see Section 3 for the original dimensions) and "-" means that no feature was extracted.

Features	Μ	SD	$\Delta \mathbf{M}$	$\Delta SD$	Total
Compactness	1	1	1	1	4
Energy	1	1	1	1	4
Entropy of Energy	1	1	-	-	2
FoLEW	1	1	1	1	4
RMS	1	1	1	1	4
Zero Crossing	1	1	1	1	4
SBeat	1	1	1	1	4
SSBeat	1	1	1	1	4
Beat sum	1	1	1	1	4
MFCC	22*	22*	15*	26	85
Chroma Vector	11*	11*	-	-	22
SD of Chroma	1	1	-	-	2
LPC	9*	9*	0*	10	28
Spectral Centroid	1	1	0*	1	3
Spectral Flux	1	1	1	1	4
Spectral Rolloff	1	1	1	1	4
Spectral Spread	1	1	1	1	4
Spectral Variability	1	1	1	1	4

Table 2. Autoencoder architecture.			
Layer	#Nodes		
Dense, PReLU, Dropout(0.2)	60		
Dense, PReLU, Dropout(0.2)	20		
Dense, PReLU	60		
Dense, Sigmoid	190		

#### 6. RESULTS

In the Music Information Retrieval Evaluation eXchange (MIREX) 2012, the state of the art systems for 10 genre cat-

egories achieved an accuracy of 50-80% [8]. Table 3 shows notable results in the literature. Regarding our approach, we measured the mean accuracy over ten experiments, at different stages of our pipeline: after the extraction of the contentbased features, as the output of two cycles of windowing, feature extraction and early temporal integration steps, we obtained a mean accuracy of 78%. The preprocessing step of feature selection increased this result to the 86.30%. At the third and final stage, by performing an additional feature extraction step with autoencoder, we achieved 91%, thus surpassing the accuracy of all other approaches.

**Table 3**. Notable classification accuracies achieved in the literature for musical genre classification.

Reference	Accuracy
Our approach	91.00%
Our approach, no bottleneck features	86.30%
Sturm et al. [18]	83.00%
Bergestra et al. [19]	82.50%
Li et al. [20]	78.50%
Panagakis et al. [21]	78.20%
Lidy et al. [22]	76.80%
Benetos et al. [23]	75.40%
Holzapfet et al. [24]	74.00%
Tzanetakis et al. [7]	61.00%

#### 7. CONCLUSIONS

We introduced a new feature engineering pipeline for the automatic identification of musical genre. In order to integrate the loss of information during the feature extraction, we reorganized the rigid waterfall scheme of the traditional two-step process of extraction and classification. We maintained continuity with respect to existing approaches in different fundamental aspects, like temporal integration method and feature taxonomy. Simultaneously we carried out a feature selection method based on information gain and extended the contentbased features with bottleneck layer's features of an autoencoder. Neither of the two methods has ever been applied on the feature vectors extracted from the musical signal. Thank to the feature selection stage, our system achieved 86.30% of accuracy. The extension of the preprocessed feature vectors with the bottleneck features had a further contribution of 4.7% to the accuracy of the final model, which resulted to be 91%.

If we take into account the semantic ambiguity every genre recognition system must deal with, the results we obtained open new insights for the future works on definition and construction of genre and sub-genre identification systems, with an accuracy which can reach human performance.

# 8. REFERENCES

- Richard F Lyon, "Machine hearing: An emerging field [exploratory dsp]," *IEEE signal processing magazine*, vol. 27, no. 5, pp. 131–139, 2010.
- [2] David Gerhard, Audio signal classification: History and current techniques, Department of Computer Science, University of Regina, 2003.
- [3] Bob L Sturm, "A survey of evaluation in music genre recognition," in *International Workshop on Adaptive Multimedia Retrieval*. Springer, 2012, pp. 29–66.
- [4] Cyril Joder, Slim Essid, and Gaël Richard, "Temporal integration for audio classification with application to musical instrument classification," *IEEE Transactions* on Audio, Speech, and Language Processing, vol. 17, no. 1, pp. 174–186, 2009.
- [5] Anders Meng, Peter Ahrendt, Jan Larsen, and Lars Kai Hansen, "Temporal feature integration for music genre classification," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 5, pp. 1654– 1664, 2007.
- [6] Francesc Alías, Joan Claudi Socoró, and Xavier Sevillano, "A review of physical and perceptual feature extraction techniques for speech, music and environmental sounds," *Applied Sciences*, vol. 6, no. 5, pp. 143, 2016.
- [7] George Tzanetakis and Perry Cook, "Musical genre classification of audio signals," *IEEE Transactions on speech and audio processing*, vol. 10, no. 5, pp. 293– 302, 2002.
- [8] Alexander Lerch, An introduction to audio content analysis: Applications in signal processing and music informatics, John Wiley & Sons, 2012.
- [9] Ian T Jolliffe, "Principal component analysis and factor analysis," in *Principal component analysis*, pp. 115– 128. Springer, 1986.
- [10] DE Rummelhart and JL McClelland, "Parallel distributed processing, vol. 1: Foundations," 1986.
- [11] Tom M Mitchell, "Machine learning. 1997," Burr Ridge, IL: McGraw Hill, vol. 45, no. 37, pp. 870–877, 1997.
- [12] Theodoros Giannakopoulos, "pyaudioanalysis: An open-source python library for audio signal analysis," *PloS one*, vol. 10, no. 12, pp. e0144610, 2015.
- [13] Cory McKay, Ichiro Fujinaga, and Philippe Depalle, "jaudio: A feature extraction library," in *Proceedings* of the International Conference on Music Information Retrieval, pp. 600–3.

- [14] Andrew L Maas, Awni Y Hannun, and Andrew Y Ng, "Rectifier nonlinearities improve neural network acoustic models," in *Proc. ICML*, 2013, vol. 30.
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Delving deep into rectifiers: Surpassing humanlevel performance on imagenet classification," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1026–1034.
- [16] Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting.," *Journal of machine learning research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [17] Matthew D Zeiler, "Adadelta: an adaptive learning rate method," arXiv preprint arXiv:1212.5701, 2012.
- [18] Bob L Sturm, "On music genre classification via compressive sampling," in 2013 IEEE International Conference on Multimedia and Expo (ICME). IEEE.
- [19] James Bergstra, Norman Casagrande, Dumitru Erhan, Douglas Eck, and Balázs Kégl, "Aggregate features and adaboost for music classification," *Machine learning*, vol. 65, no. 2-3, pp. 473–484, 2006.
- [20] Tao Li, Mitsunori Ogihara, and Qi Li, "A comparative study on content-based music genre classification," in Proceedings of the 26th annual international ACM SI-GIR conference on Research and development in informaion retrieval. ACM, 2003, pp. 282–289.
- [21] Ioannis Panagakis, Emmanouil Benetos, and Constantine Kotropoulos, "Music genre classification: A multilinear approach," in *ISMIR*, 2008, pp. 583–588.
- [22] Thomas Lidy, Andreas Rauber, Antonio Pertusa, and José Manuel Inesta, "Mirex 2007 combining audio and symbolic descriptors for music classification from audio," *MIREX 2007 - Music Information Retrieval Evaluation eXchange*, 2007.
- [23] Emmanouil Benetos and Constantine Kotropoulos, "A tensor-based approach for automatic music genre classification," in *Signal Processing Conference*, 2008 16th European. IEEE, 2008, pp. 1–4.
- [24] Andre Holzapfel and Yannis Stylianou, "Musical genre classification using nonnegative matrix factorizationbased features," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 16, no. 2, pp. 424–434, 2008.