FACTORS AFFECTING AUTOMATIC GENRE CLASSIFICATION: AN INVESTIGATION INCORPORATING NON-WESTERN MUSICAL FORMS

Noris Mohd Norowi, Shyamala Doraisamy, Rahmita Wirza
Faculty of Computer Science and Information Technology
University Putra Malaysia
43400, Selangor, MALAYSIA
`{noris,shyamala,rahmita}@fsktm.upm.edu.my`

ABSTRACT

The number of studies investigating automated genre classification is growing following the increasing amounts of digital audio data available. The underlying techniques to perform automated genre classification in general include feature extraction and classification. In this study, MARSYAS was used to extract audio features and the suite of tools available in WEKA was used for the classification. This study investigates the factors affecting automated genre classification. As for the dataset, most studies in this area work with western genres and traditional Malay music is incorporated in this study. Eight genres were introduced; *Dikir Barat, Etnik Sabah, Inang, Joget, Keroncong, Tumbuk Kalang, Wayang Kulit*, and *Zapin*. A total of 417 tracks from various Audio Compact Discs were collected and used as the dataset. Results show that various factors such as the musical features extracted, classifiers employed, the size of the dataset, excerpt length, excerpt location and test set parameters improve classification results.

Keywords: Genre Classification, Feature Extraction, Music Information Retrieval, Traditional Malay Music

1 INTRODUCTION

Improvements in audio compression along with increasing amounts of processing power, hard disk capacity and network bandwidth have resulted in increasingly large number of music files. Easier distribution of digital music through peer-to-peer file sharing has also made possible creation of large, digital personal music collection, typically containing thousands of popular songs. Users can now store large personal collections of digital music. These growing collection of digital audio data needs to be classified, sorted, organized and retrieved in order to be of any value at all.

At present, metadata such as the filename, date created, etc., are used at large but is very labor-intensive, costly and time consuming. A system that allows classification of music based on the audio characteristics is therefore highly sought.

Musical genre is used universally as a common metadata for describing musical content. Genre hierarchies are widely used to structure the large collections of music available on the Web. Musical genres are labels created and used by humans for categorizing and describing the vast universe of music [1]. Humans possess the ability to recognize and analyze sound immediately based on instrumentation, the rhythm and general tone. Furthermore, humans are able to draw connections to other songs that have a similar sound and feel. These commonalities make it possible for humans to classify music into different genres.

An automatic genre classification is a system that allows structuring and organization of the huge number of archived music automatically. The system should also be able to analyze and possibly extract the implicit knowledge of these musical files that infers the structure that underlies beneath the audio information into a comprehensible form. The analysis includes evaluation and comparisons of the feature sets that attempt to represent the musical content. The general framework for automatic genre classification includes feature extraction and classification. Various studies have been carried out toward the development [1,2,3,4].

Wold et al. [2] extensively discuss the essence of audio classification, search and retrieval. Many audio features were analysed and compared, such as rhythms, pitch, duration, loudness and instrument identification, in particular. Their work is not only limited to classification of music but includes classification of speech, laughter, gender, animal sounds and sound effects.

Tzanetakis [1] also found that although there has been significant work in the development of features for speech recognition and music-speech discrimination, relatively little work has been done on developing features to specifically design music signals. Here, he deduces three specific features for musical content; timbral texture, rhythmic structure and pitch content to be specific. This is similar to the conclusions from Aucoturier [3].

The features are not the only basis of a genre classification system. Kaminskyj [4] proves that classifiers or machine learning algorithms are just as crucial, where Gaussian Mixture Model was used. These approaches work conceivably well on western musical forms. However, it is very likely that the representation used for extraction of one genre is bad at describing other genres. Thus, a whole new set of features may be required in order to cater other musical forms.

Additionally, the classifier, test set parameters and dataset pre-processing issues are also taken into consideration in this study in order to get the real description of their effects on classification.

This paper investigates the factors affecting classification results and expands the customary work scope of using western musical forms onto other non-western musical forms, i.e. traditional Malay musical forms. The remainder of this paper comprises six sections. The effect of combining different feature sets on classification is discussed in Section 2 followed by analysis of direct comparison between two classifiers in Section 3. Section 4 describes the experimental framework of this study. Results are revealed in Section 5. Finally, conclusion and future work are presented in Section 6.

2 FEATURE SETS

In music genre recognition, one challenge is the ability to differentiate between musical styles. Feature extraction is important for this. Feature extraction is the part where knowledge of music, psychoacoustics, signal processing and many other fields is considered. It is a process where a segment of an audio is characterized into a compact numerical representation. Audio data are fairly large to be stored and processing them can take a lot of space and time. This is the reason why typically, the features for audio and speech analysis algorithms are computed on a frame basis. By doing so, the amount of data that needed to be processed could be reduced.

Numerous audio features have been identified by various studies, each individual one or combination of them is best suited for classifying a certain audio class, be it between music and speech, between noise and music, male or female identifier, or musical genre classification. However, it is important to find the right features as any classification algorithm will always return with some kind of result, but a poor feature representation will only yield results that do not reflect the real nature of the underlying data. The extracted features will be useful during classification when standard machine learning techniques are used in the next step.

Works in Tzanetakis [1] and Aucoturier [3] proposed that for automatic musical genre classification, there are three prescriptive sets of feature that should be looked into thoroughly: timbral related features, rhythm related features and pitch related features. In order to exploit these features, it is vital to identify the feature space where all samples belonging to a particular genre must cluster closely. At the same time, clusters corresponding to different genres must have a large distance between them.

2.1 Timbral Related

The calculated features are based on the short time Fourier transform and are calculated for every short-time frame of the sound instead of calculating only one timbre value for the entire song although timbral feature is sometimes referred to global timbral.

2.1.1 Spectral Centroid

This is the gravity centre of the spectral distribution within a frame. The centroid measures the spectral shape. Higher centroid values indicate higher frequencies.

2.1.2 Spectral Roll-Off

The Roll-Off is another measure of spectral shape. It is the point where frequency that is below some percentage (usually at 85%) of the power spectrum resides.

2.1.3 Spectral Flux

This feature measures frame-to-frame spectral difference. In short, it tells the changes in the spectral shape. It is defined as the squared difference between the normalized magnitudes of successive spectral distribution.

2.1.4 Time Domain Zero Crossing

This refers to the number of time –domain zero crossings within a frame. Zero crossings occur when successive samples in a digital signal have different signs. In simple signals, ZCR is directly related to the fundamental frequency. This feature can be used as a measure of noisiness in a signal.

2.1.5 Mel-Frequency Cepstral Coefficients

MFCCs are based on the spectral information of a sound, but are modelled to capture the perceptually relevant parts of the auditory spectrum. It is the inverse of Fourier transform. They are thought to capture the perceptually relevant part of the auditory spectrum. Naturally, there are thirteen coefficients that have been found to be useful in speech representation. However, for genre classification, the first five correlations appear to be of any importance in terms of performance.

2.2 Rhythm Related

Other than the timbral related features, another feature set that can be made useful is the rhythm related features. The beat and rhythmic structure of a song is a good genre indicator. In Tzanetakis [1], a beat histogram is built from autocorrelation function of the signal. The beat histogram provides an outline of the strongness and complexity of the beat in the music. This is an especially remarkable feature in order to discriminate between an energetic genre such as rock and classical where the beat is not so accentuated.

Beat tracking and rhythm detection is a growing area of research nowadays despite it being remarkably difficult to be developed for automated systems. There are a number of reasons why this beat tracking appeared so difficult whilst humans manage to recognize the beat of the music effortlessly. Humans can easily determine the beat even if the tempo and metrical structure are not explicitly specified in the beginning of the song. In

addition, humans can also handle it when the tempo changes throughout the song as opposed to current systems where these are considered major and unsolved obstacles yet. Nevertheless, Kosina [5] gives a good overview on beat tracking methods.

In terms of traditional Malay musical forms, this is seen as a very useful feature set indeed as the genre itself is made up of repeating rhythmic patterns.

2.3 Pitch Related

Ermolinskij [6] attempts at building a genre classifier for audio based on pitch features through the use of pitch histogram feature vectors. The histograms reveal some genre-specific information such genre with higher degree of harmonic variation tends to have a fair amount of well-pronounced peaks in comparison to those genre with lower degree of harmonic variation such as rock.

3 CLASSIFIERS

Classification relies on the basic assumption that each observed pattern belongs to a category. Individual signals may be different from one another, but there is a set of features that are similar to patterns belonging in same class and different patterns for a different class. The feature sets are the base that can be used to determine class membership. Classification is domain-independent and provides many fast, elegant and well-understood solutions that can be adopted for use in music recognition.

There are a number of key aspects in designing a good classification system. The classification system in itself must be robust, with tolerable complexity in terms of speed and memory usage when used in software.

Various number of classification techniques are available today. Some are better than others in detecting a pattern between the feature vectors. The dissimilarities caused by different underlying models are one of the problems. Likewise, the fact that classification only deals with the feature vector representation of the actual data makes it difficult to classify accurately owing to the feature values of signals that can vary considerably even when they belong in the same category. Another contributor towards the test of classification is the considerable variation that is caused by noise.

Creating a classifier usually means specifying its general form. The unknown parameters are estimated through training. Training can be defined as the process using sample data to determine the parameter settings of the classifier, and is essential in all real-world classification system [5]. Classification is often called supervised learning. It involves the first stage of training where models of a few musical genres are built with some manually labelled data. The second stage involves testing or recognition, where these models are used to classify unlabelled data. It is crucially important to ensure that the test data was not used in any way to create classifier during training. This means that the dataset should be divided in some way so that a fraction of it can be used for testing and the remaining part for training. In general, the larger the training sample the better the classifier and the larger the test sample, the more accurate the error estimate. In order to get a good classifier, a certain kind of arrangement must be made so that one does not jeopardize the other [7].

A machine-learning scheme called WEKA (Waikato Environment for Knowledge Analysis) was engaged to evaluate the computer audition applications using trained statistical pattern recognition classifiers. It enables pre-processing, classifying, clustering, attributes selections and data visualizing. WEKA is employed when applying a learning method to a dataset and during analysis of its output to extract information about the data. The learning methods are called *classifiers* [7].

An example of such classifier is the OneR classifier. This is one of the most primitive schemes. It produces simple rules based on one attribute only. Although it is a minimal form of classifier, it can be useful for determining a baseline performance as a benchmark for other learning schemes.

Another well-known classifier is J48. This is just one of the many practical learning schemes that can be applied to any dataset. J48 classifier forms rules from pruned partial decision trees built using C4.5's heuristics. C4.5 is Quinlan's most recent non-commercial tree-building algorithm. The main goal of this scheme is to minimize the number of tree levels and tree nodes, thereby maximizing data generalization. It uses a measure taken from information theory to help with the attribute selection process. For any choice point in the tree, it selects the attribute that splits the data so as to show the largest mount of gain in information.

The J48 classifier described above builds a C4.5 decision tree. Each time the Java virtual machine executes J48; it creates an instance of this class by allocating memory for building and storing a decision tree classifier. The algorithm, the classifier it builds, and a procedure for outputting the classifier, are all part of that instantiation of the J48 class.

The J48 class does not actually contain any code for building a decision tree. It includes references to instances of other classes that do most of the work. It also combines the divide-and-conquer strategy for decision tree and separate divide-and-conquer one for rule learning. Such approach adds flexibility and speed.

4 EXPERIMENTAL FRAMEWORK

A general methodology is formulated in this study with the aim of improving classification results by distinguishing the factors involved and also through parameter optimisation. As this problem falls into the category of supervised machine learning, it is a common approach to map the training data into feature vectors first. Once mapped, one or more classification techniques are applied on this data and a model for distribution underlying the data is created. This model, in the final stage will be used to estimate the likelihood of a particular category given the test data.

4.1 Experimental Set Up

Table 1 below sums up the five different experimental sets carried out during this study.

As can be seen in Table 1, five different experimental sets were used. In the first experimental set, the dataset size was used as the variable. This was done to examine whether dataset size play a major role in determining the classification results. The sizes were changed between a minimum of ten songs per genre and thirty songs per genre. Since not all genres contain these amounts of song, some genre had to be eliminated altogether. The second set focused on finding whether the track length has a significant role. Afterwards, the starting points of each dataset were altered to from starting after minute into the song to starting at point zero of the songs. The remaining of the sets dealt with classification parameters such as the number of cross-validation folds and classifiers utilized.

4.2 Dataset Outlook

In total, eight traditional Malay musical genres were used in this study namely *Dikir Barat, Etnik Sabah, Inang, Joget, Keroncong, Tumbuk Kalang, Wayang Kulit*, and *Zapin*. These eight genres were run against five common western genres (Blues, Classical, Jazz, Pop and Rock).

Table 1. Different experimental sets used when conducting the study

Set	Factors
1	Dataset Size
2	Dataset Track Length
3	Dataset Starting Point
4	Number of Cross-Validation Folds
5	Classifiers

The performing arts of Malaysia are mainly derivative [8], influenced by the initial overall Indian and Middle Eastern music during the trade era and later from colonial powers such as Thailand, Indonesia, Portuguese and British who introduce their own culture including dance and music.

The taxonomy of traditional Malay music depends on the nature of the theatre forms they serve and their instrumentations. The musical ensembles usually include *gendangs* or drums (membranophone) that is used to provide the constant rhythmic beat of the songs, gongs (idiophone) to mark the end of a temporal cycle at specific part of the songs, and some other traditional Malay instruments such as the *rebab* (chordophone), and *serunai* or *seruling* which resemble a lot like a wooden oboe and flute made up of bamboo respectively (aerophone). These instruments, in contrast to western music, which is based on the western tempered scale of twelve semitones, are without such standard [8]. However, two types of tonal systems are generally applicable; heptatonic seven tone scale and pentatonic five tone scale.

4.3 Dataset Treatment

417 dataset in total were used. Although this might seem like a passable amount of dataset, the number of songs per genre was not consistent. Western musicals were much easier to obtain and undeniably, huge portion of the dataset were made up of them. However, an adequate amount of dataset for traditional Malay songs managed to be collected.

The dataset were obtained from various sources, mainly from Audio Compact Discs and some were also downloaded via the Internet. While downloading from the Internet appeared a trivial task for western songs, the same could not be said for traditional Malay songs. As confirmed in [9], it is evident that traditional Malay musical culture is in the verge the corrosion, which signifies the complications in getting large dataset as originally intended. Various individuals and organizations provided these musical dataset including the Malaysia National Arts Academy, Sultan Salahuddin Abdul Aziz Shah's Cultural and Arts Centre at Universiti Putra Malaysia and also personal collections of audio CDs from many individuals.

The dataset came in an assortment of audio formats. These were later converted into wav format using standard audio editing tool. Wav format was chosen as it was the only format supported by the free feature extractor called MARSYAS, which will be discussed next.

After format conversion, these data were trimmed into a uniformed length each. A standard length for all was required because each song varied lengthwise and it was thought that with this could be one of the factors affecting classification results. In addition, based existing work of [1], the dataset used were also trimmed into a neat collection of thirty seconds each. Furthermore, using full-length music takes up too much space and is likely to increase computational load. These clips were used throughout the experiments, both for training and testing.

4.2 Feature Extraction

For each of the song clips, the features were extracted to facilitate automatic music genre classification. This was done through MARSYAS; a free framework that enables the evaluation of computer audition applications. The Musical Research System for Analysis and Synthesis (MARSYAS) is a semi-automatic music classification system that is developed as an alternative solution for the existing audio tools that are incapable of handling the increasing amount of computer data [10].

When used for music genre classification, it performs notably better than by chance. It utilizes the three feature sets for representing the timbral texture, rhythmic content and pitch content of the music signals and uses

Table 2. Confusion Matrix for Set 1 (General)

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Acc
A	26									1			2		1		87
B		5								2							71
C	1		33														94
D				30								1					97
E					12												100
F						7						1					88
G	1						7	1					1				70
H	1		1		1			13		2						1	68
I	1			1					10				1	1	1		67
J		2		1						31							91
K			1		1						5						71
L	1			1		1			1			63					94
M			2	1			1	2				4	95				90
N			1											12			92
O												1			15	1	88
P													3	1		4	40
Rel	84	100	85	86	80	88	88	87	83	84	100	91	93	80	88	67	

A : Blues	E : *Etnik Sabah*	I : *Joget*	M : Rock
B : *Bongai*	F : *Gamelan*	J : *Keroncong*	N : *Tumbuk Kalang*
C : Classical	G : *Inang*	K : *Muzik Asli*	O : *Wayang Kulit*
D : *Dikir Barat*	H : Jazz	L : Pop	P : *Zapin*

trained statistical pattern recognition classifiers for evaluation. The feature extractor will end up in numerical results in the form of an ARFF file.

As of the time of the experiment was conducted, the exact same feature was used to extract audio features from the signals, i.e. the means and variances of multi-dimensional features such as Spectral Centroid, Spectral Flux, Spectral Roll-Off, Zero Crossings and the Low Energy were engaged in this experiment.

4.3 Classification

The ARFF files that were generated by the feature extractor containing representations of all the music files were used to train a J48 classifier in WEKA. It enables pre-processing, classifying, clustering, attributes selection and data visualizing. In order to achieve the objective of this study, the features that distinctively classify traditional Malay music must be identified and utilized.

Classification results were tested using stratified three-fold cross-validation, six-fold cross-validation and ten-fold cross validation respectively. Cross-validation is a standard evaluation technique in pattern classification, in which the dataset is split into *n* parts (folds) of equal size. *n*-1 folds are used to train the classifier. The *n*th fold that was held out is then used to test it.

5 RESULTS

The first part of this study was to observe the behaviour of the classification system, as this was a first attempt at incorporating non-western music. Table 2 shows the confusion

matrix where the columns correspond to the actual genre and the rows to the predicted genre. In this particular confusion matrix, the labels each correspond to a particular genre. The name of each label can be referred below the table along with the number of songs available per genre.

It can be clearly seen that the number of songs per genre are not consistent. Some may have as little as 7 songs per genre (*Bongai* and *Muzik Asli)*, which is especially true with traditional Malay music, whereas some can be over a hundred songs per genre.

The imbalance in terms of number might result in a biased classification. In order to avoid this, these genres had to be eliminated altogether.

The diagonal pattern starting from the top left hand corner of the table towards the bottom right hand of the table illustrate number of correct classification. The dispersed numbers outside the diagonal pattern tell the number of misclassifications. In this particular test set, the expected diagonal pattern is present, though the classification itself is not perfect, even when using the exact same dataset for both training and testing. This suggests that there are more factors that contribute towards developing an ideal system for automatic genre classification.

Since the result dataset were not pre-processed beforehand, and there is no standardized number of dataset for any particular genre, it was believed that with data pre-processing, classification results could be enhanced.

Classification was then performed using the modified dataset based on the factors listed in Table 1. Results were evaluated in terms of accuracy and reliability [11].

Classification accuracy indicates how many of the test samples were correctly classified whilst classification reliability discloses the confidence level that can be placed on the classifier results. The results are discussed below.

1) Larger dataset size is favourable

Whilst appointing J48 classifier at 10 songs per genre, 84% of correct classification had been achieved. The examination was then continued by setting a higher number of minimum songs per genre. At 30 songs per genre, some genres had to be removed from the list, leaving only six genres altogether (Blues, Classical, *Dikir Barat*, *Etnik Sabah*, Pop, Rock). Again, the parameters were kept as above and ultimately, there was a slight increase (90%) in the classification result. This shows consistent numbers of songs per genre, results in better and unbiased classification accuracy. Figure 1 further illustrates the results of the two dataset sizes.


Figure 1. Classification performance between two different dataset sizes

2) Excerpt length need not be long.

It was originally thought that the longer the duration of a musical excerpt, the better the classification accuracy. 30 seconds was chosen as the default value for classification in all the experimental set in this study as used by previous researchers in their respective research [1]. Maintaining the same parameters as before, tracks of 10 seconds, 30 seconds and 60 seconds were tested.

In Figure 2, it can be seen that for 90% accuracy and reliability were achieved with 30 seconds tracks and 93% when used with 60 seconds tracks. Although results appear better when the track length is extended to 60 seconds, the increase is minute in comparison to the storage cost required and the heavier computational load, as audio data are huge in size. Interestingly, when tested with 10 seconds tracks, the result was better than the previous two. This outcome implies that the standard 30 seconds used by many researchers may not be the best length for genre classification.


Figure 2. Classification performance between three varying track lengths

3) The first few seconds of a song are crucial genre indicator

It was hypothesized that in order to get a better 'feel' of the genre of a song, it is better to process the tracks halfway into the song where all the actions and instruments can be heard. In this set, the exact tracks were processed twice; the first batch was trimmed at one minute into the song while the second batch was taken from the beginning. The length of both batches was kept at 30 seconds long.

It is astounding to find that classification performed better on tracks that started from the beginning compared to tracks that are halfway into the song (Figure 3). When a song reaches the middle part, the chorus comes in, all the instruments are played, the song usually becomes more dynamic and alive, and the overall energy of the audio signal gets higher. This explains why sometimes classical can be misclassified as rock. It is possible that due to these reasons, it is easier to obtain higher classification results with tracks that began from the beginning of the song.
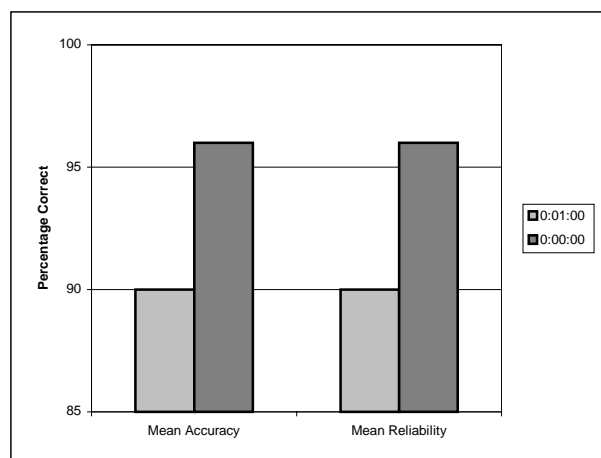

Figure 3. Classification performance between two different starting points

4) Higher Cross Validation Folds may not always attain better results

Investigation on the effect of increasing number of folds in cross-validation were done to uncover whether it really is better to have at least ten or more folds as ten seems to be the default value used by many [7].

At three folds, the accuracy and reliability of the result was at 72%. When increased to six folds, the performance also increased to 75%. When increased further at ten folds, the performance dropped slightly to 74%. The number remained at 74% when further tested at 15 folds.

It appears that after a certain point, increasing the number of folds would not be of any aid in classification performance. With this particular dataset, six-folds is seen to be the optimal folds (see Figure 4). However, this does not mean that six folds is best for every classification problem, and that perhaps ten folds is also a safe number to rely on as it does not effect the results in a major way.

5) Classifiers must be chosen according to specific needs to ensure higher results accuracy

Two different classifiers were tested; OneR and J48 to demonstrate this point. With OneR classifier, results were just above satisfactory at 67% accuracy and 44% reliability. J48 classifier, on the other hand, boosted the classification performance to 75% overall. This just shows that some classifiers are suitable in performing a classification problem while others may not.

6 CONCLUSION

Results show that the audio classification can be improved by taking into consideration the factors such as dataset size, track length and track location, the number of cross-validation folds and utilizing the suitable classifiers.

Overall, it is best to use tracks that are no more than 30 seconds in length, starts from the beginning of the track and apply J48 classifier for categorical classification. In addition, this study supports the theory that states larger dataset are favourable in order to come up with a better classification result.

It is also clear that audio classification of traditional Malay music is possible with existing genre classification tools which thus far has been used only investigated using western musical genres. Expanding the scope of automatic genre classification beyond western musical forms proves that classification of other non-western music is also achievable and can be used just as good.

Future work includes investigating the specific features that improve classification performance traditional Malay music and to work on a larger dataset.
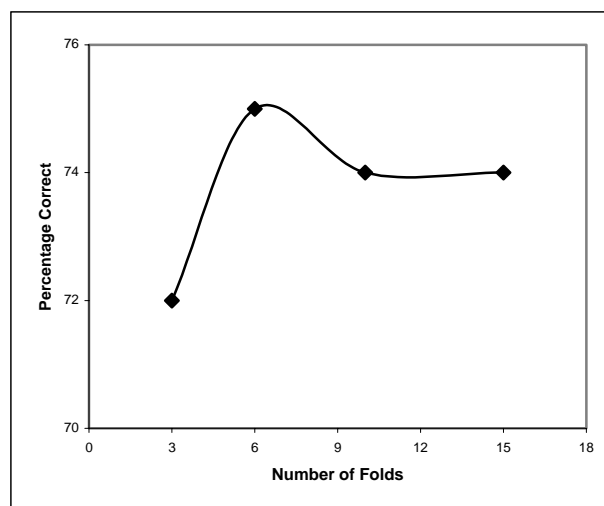

**Figure 4.** Classification performance between different numbers of cross-validation folds

Figure 5. Classification performance between two different classifiers

REFERENCES

[1] Tzanetakis, George and Perry Cook "Musical Genre Classification of Audio Signals", *IEEE Transactions on Speech and Audio Processing* , 10(5), July 2002.

[2] Wold, E., Blum, T., Keislar, D., and Wheaton, J. "Content-based Classification, Search, and Retrieval of Audio." IEEE Multimedia, Vol. 3, No.3, pp. 27-36, Fall 1996.

[3] Aucoturier, J., and Pachet, F. "Representing Musical Genre: A State of the Art". Journal of New Music Research. Vol. 32, No. 1, Nov. 1993, pp. 83-93.

[4] Kaminskyj, I. "Multi-feature Musical Instrument Sound Classifier." Proceedings ACMA Conference, pp. 46-54.

[5] Kosina, K. Music Genre Recognition. Diploma Thesis, Hagenberg University, June 2002.

[6] Ermolinskiy, A., Cook, P., and Tzanetakis, G. "Musical Genre Classification based on the analysis of harmonic content features in audio and midi. *Work Report,* Princeton University.

[7] Witten, I., and Frank E. Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations. Morgan Kauffman Publisher, 1999.

[8] Mohd Ghouse Nasuruddin, The Malay Traditional Music. Kuala Lumpur: Dewan Bahasa dan Pustaka, (1992).

[9] Shriver, R. "Digital Stereo Recording of Traditional Malaysian Musical Instruments. Audio Engineering Society Convention Paper". March 2003.

[10] George Tzanetakis and Perry Cook "MARSYAS: A Framework for Audio Analysis" *Organized Sound, Cambridge University Press* 4(3), 2000

[11] Pandya A. S., and Macy, R. B. "Pattern Recognition with Neural Networks in C++", CRC Press, 1996.