

## MUSICAL INSTRUMENT IDENTIFICATION IN CONTINUOUS RECORDINGS

*Arie A. Livshin*

Analysis/Synthesis Team  
Ircam, Paris, France  
livshin@ircam.fr

*Xavier Rodet*

Analysis/Synthesis Team  
Ircam, Paris, France  
rod@ircam.fr

### ABSTRACT

Recognition of musical instruments in multi-instrumental, polyphonic music, is a difficult challenge which is yet far from being solved. Successful instrument recognition techniques in solos (monophonic or polyphonic recordings of single instruments) can help to deal with this task.

We introduce an instrument recognition process in solo recordings of a set of instruments (bassoon, clarinet, flute, guitar, piano, cello and violin), which yields a high recognition rate. A large and very diverse solo database (108 different solos, all by different performers) is used in order to encompass the different sound possibilities of each instrument and evaluate the generalization ability of the classification process.

First we bring classification results using a very extensive collection of features (62 different feature types), and then use our GDE feature selection algorithm to select a smaller feature set with a relatively short computation time, which allows us to perform instrument recognition in solos in real-time, with only a slight decrease in recognition rate.

We demonstrate that our real-time solo classifier can also be useful for instrument recognition in duet performances.

### 1. INTRODUCTION

Most works on instrument recognition have dealt with classification of separate musical tones taken from professional sound databases, e.g. McGill, Studio Online, etc.

Instrument recognition in solo performances (monophonic or polyphonic musical phrases performed by a single instrument) is different and more complicated than dealing with separate note databases, as the time evolution of each sound (attack, decay, sustain, release) is not well defined, the notes are not separated, there are superpositions of concurrent sounds and room echo, different combinations of playing techniques, etc. Marques and Moreno [1] classified 8 fairly different instruments (bagpipes, clarinet, flute, harpsichord, organ, piano, trombone and violin) using one CD per instrument for learning and one for classification. They compared 3 feature types using 2 different classification algorithms, and achieved 70% recognition rate. Brown Houix and McAdams [2] classified 4 wind instruments (flute, sax, oboe and clarinet), compared 4 feature types and reached 82% recognition rate with the best combination of parameters and training material. Martin [3] has classified sets of 6, 7 and 8 instruments, reaching 82.3% (violin, viola, cello, trumpet, clarinet, and flute), 77.9% and 73% recognition rates respectively. He

used up to 3 different recordings from each instrument; in each experiment one recording was classified while the rest were learned. The feature set was relatively large and consisted of 31 one-dimensional features. For a comprehensive review of instrument recognition, see [4].

The work on solo recognition is not yet exhausted. Although it seems that there are not many applications which actually require solo recognition, yet as we shall demonstrate at the end of this paper, knowledge of how to deal well with solos can also help in recognition of multi-instrumental music (where several instruments play concurrently). The subject of musical instrument recognition in multi-instrumental music is difficult and is just beginning to get explored (e.g. [5]).

We begin the paper by presenting a process for recognition of a set of instruments (bassoon, clarinet, flute, guitar, piano, cello and violin) which yields a high average recognition rate - 88.13% when classifying 1-second pieces of real recordings.

A large and very diverse solo database is used for learning and evaluating the recognition process. It contains 108 solo performances, all by different musicians, and apparently supplies a good generalization of the different sound possibilities of each instrument in various recording conditions, playing techniques, etc., thus providing a good generalization of the sounds each instrument is capable of producing in different recordings - what we call the "concept instrument". In order to evaluate the generalization ability of the classifier, the same solos are never used both in the learning and test sets; we have proved that a classification evaluation process in which the training and learning sets both contain samples recorded in very similar conditions is likely to produce misleading results [6].

We use a very large collection of features for solo recognition - 62 different feature types [7] which were developed and used in the Cuidado project. Using our GDE feature selection algorithm, we select a smaller feature set best suited for solo recognition in real-time (of our 7 instruments), with only a small reduction in recognition rate (85.24%) compared to the complete feature set. We present the features of this real-time feature set, which was actually implemented in a real-time solo recognition program.

We end the paper by demonstrating that the same features and techniques we used for real-time solo recognition can also help to perform instrument recognition in duet performances.

## 2. SOLO DATABASE

Our sound database consists of 108 different ‘real-world’ solo performances (by “solo” we mean that a single instrument is playing, in monophony or polyphony) of 7 instruments: bassoon, clarinet, flute, classical guitar, piano, cello and violin. These performances, which include classical, modern and ethnic music, were gathered from commercial CD’s (containing new or old recordings) and MP3 files played and recorded by professionals and amateurs.

Each solo was performed by a different musician and there are no solos taken from the same concert. During the evaluation process we never use the same solo, neither fully or partly, in both the learning set and the test set. The reason for these limitations is that we need the evaluation process to reflect the system’s ability to generalize – i.e. classify new musical phrases which were not learned, and were recorded in different recording conditions, different instruments and played by different performers than the learning set. We have proved [6] that the evaluation results of a classification system which does learn and classify sounds performed on the same instrument and recorded in the same recording conditions, even if the actual notes are of a different pitch, are much higher than when classifying sounds recorded in different recording conditions. The reason is that such an evaluation process actually shows the system’s ability to learn and then recognize specific characteristics of specific recordings and not its ability to generalize and recognize the “concept instrument”.

### 2.1. Preprocessing

All solos were downsampled to 11Khz, 16bit. Only the left channel was taken out of stereo recordings<sup>1</sup>. A 2-minute piece was taken from each solo recording and cut into 1-second cuts with a 50% overlap – a total of 240 cuts out of each solo.

## 3. FEATURE DESCRIPTORS

The computation routines for the features we use in the classification process were written by Geoffroy Peeters as part of the Cuidado project. Full details on all the features, can be found in [7].

The features are computed on each 1-second solo-cut separately. Besides several features<sup>2</sup> which were computed using the whole signal of the 1-second cut, most of the features were computed using a sliding frame of 60 ms with a 66% overlap. For each solo-cut of 1 second, the average and standard deviation of these frames were used by the classifier.

---

<sup>1</sup> It could be argued that it is preferable to use a mix of both channels. Which method is actually better depends on the specific recording settings of the musical pieces.

<sup>2</sup> Some features contain more than a single value, e.g. the MFCC’s; we use the term “features” regardless of their number of values.

Initially, we used a very large feature collection – 62 different features of the following types [8]:

### 3.1.1. Temporal Features.

Features computed on the signal as a whole (without division into frames), e.g. log attack time, temporal decrease, effective duration.

### 3.1.2. Energy Features.

Features referring to various energy content of the signal, e.g. total energy, harmonic energy, noise part energy.

### 3.1.3. Spectral Features.

Features computed from the Short Time Fourier Transform (STFT) of the signal, e.g. spectral centroid, spectral spread, spectral skewness.

### 3.1.4. Harmonic Features.

Features computed from the Sinusoidal Harmonic modelling of the signal, e.g. fundamental frequency, inharmonicity, odd to even ratio.

### 3.1.5. Perceptual Features.

Features computed using a model of the human hearing process, e.g. mel frequency cepstral coefficients, loudness, sharpness.

Later in the paper we shall use our GDE feature selection algorithm to reduce the number of features in order to perform instrument recognition in real-time.

## 4. “MINUS-1 SOLO” EVALUATION METHOD

After the features are computed, they are normalized using “min-max” normalization method (to the range of 0 – 1). For every solo in its turn, its 1-second solo-cuts are removed from the database and classified by the rest of the solos. This process is repeated for all solos, and the average recognition rate for each instrument is reported along with the average recognition rate among all instruments. These results are more informative than the average recognition rate per solo, as the number of solos performed by each instrument might be different.

The classification is done by first performing Linear Discriminant Analysis (LDA) [9];[10] on the learning set, multiplying the test set with the resulting coefficient matrix and then classifying using the K Nearest Neighbours (KNN) algorithm. For the KNN we use the “best” K from a range of 1 - 80 which is estimated using the leave-one-out method on the learning set<sup>3</sup> [11].

---

<sup>3</sup> The “best K” for our database was estimated as 33 for the full feature set and 39 for the real-time set. Experiments with solo-cuts using an overlap of 75% instead of 50% (resulting in 480 solo-cuts per solo instead of 240), reported a “best K” of 78 for the full feature set and 79 for the “real-time” set.

## 5. FEATURE SELECTION

After computing the recognition rate using the full feature set, we use our Gradual Descriptor Elimination (GDE) feature selection method [11] in order to find the most important features. GDE uses LDA repeatedly to find the descriptor which is the least significant and remove it. This process is repeated until no descriptors are left. At each stage of the GDE the system recognition rate is estimated.

In this section we have set the goal to achieve a smaller feature set which will be quick to compute - allowing us to perform solo recognition in real-time, and will compromise the recognition rate as little as possible, compared with the results obtained by using the complete feature set. By “real-time” we mean here that while the solo is recorded or played the features of each 1-second fraction of the music are computed and classified immediately after it was performed, before the following 1-second has finished playing/recording<sup>4</sup>.

We removed the most time-consuming features and used GDE to reduce the feature-data until the number of features went down from 62 to 20. Using these features we have actually implemented a real-time solo phrase recognition program which works on a regular Intel Processor and is written in plain Matlab code (without compilation or integration with machine language boost routines).

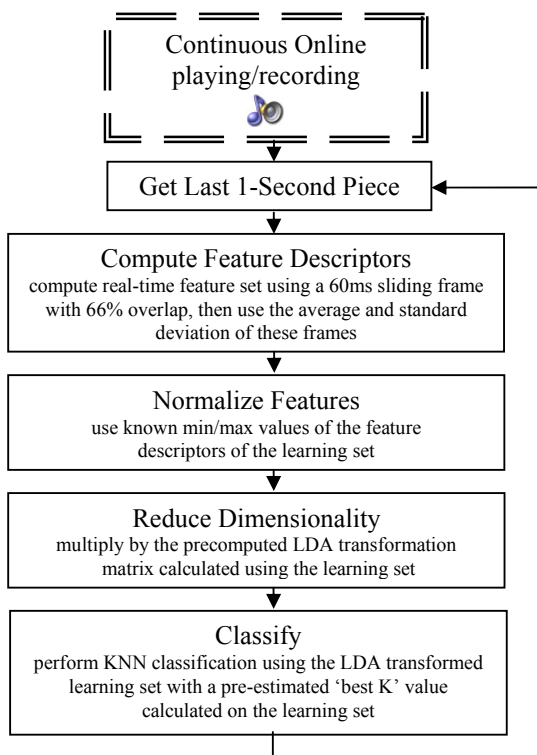


Figure 1: Real-time solo recognition process.

<sup>4</sup> Because the classified 1-second solo pieces can partially overlap, the theoretical upper limit for the recognition resolution is 1 sample.

Naturally, this program uses a precomputed LDA matrix and pre-estimated ‘best K’ for the KNN classification, as the learning set remains constant and should not depend on the solo input.

We can see in Figure 1 that the classification process uses at each round the last 1-second of the recording, which makes the recognition ‘resolution’ increase in direct relation to the hardware speed and efficiency of the sub-algorithms being used.

## 6. RESULTS

	“Real-Time” 20 features	“Complete Set” 62 features
Bassoon	86.25 %	90.24 %
Clarinet	79.29 %	86.93 %
Flute	83.33 %	80.87 %
Guitar	86.34 %	87.78 %
Piano	91.00 %	93.88 %
Cello	82.18 %	88.72 %
Violin	88.27 %	88.47 %
<b>Average</b>	<b>85.24 %</b>	<b>88.13 %</b>

Table 1: Minus-1 Solo recognition results

We can see in Table 1 that the “Real-Time” average recognition rate is indeed rather close to the “Complete Set”. It is interesting to note that while reducing the feature set we have actually improved the recognition rate of the flute; LDA does not always eliminate confusion caused by interfering features.

### 6.1. The Real-Time Feature Set

We bring in Table 2 the resulting 20 feature list for real-time classification of solos, sorted by importance, from the most important feature to the least.

1. Perceptual Spectral Slope	2. Perceptual Spectral Centroid
3. Spectral Slope	4. Spectral Spread
5. Spectral Centroid	6. Perceptual Spectral Skewness
7. Perceptual Spectral Spread	8. Perceptual Spectral Kurtosis
9. Spectral Skewness	10. Spectral Kurtosis
11. Spread	12. Perceptual Deviation
13. Perceptual Tristimulus	14. MFCC
15. Loudness	16. Auto-correlation
17. Relative Specific Loudness	18. Sharpness
19. Perceptual Spectral rolloff	20. Spectral rolloff

Table 2: A sorted list of the most important features for real-time solo classification (of our 7 musical instruments).

We can see in Table 2 that the 10 most important features are the first 4 Moments and the Spectral Slope, computed in both the perceptual and spectral models. See [7] for a full explanation of each feature.

	Bassoon	Flute	Clarinet	Guitar	Piano	Violin	Cello	Total Correct
<b>Castelnuovo: Sonatina</b>	16.2 %				83.8 %			100.0 %
<b>Stockhausen: Tierkreis</b>		50.0 %	50.0 %					100.0 %
<b>Scelsi: Suite</b>		71.1 %	28.9 %					100.0 %
<b>Carter: Esprit rude</b>		45.0 %	52.5 %				2.5 %	97.5 %
<b>Kirchner: Triptych</b>				2.5 %		37.5 %	60.0 %	97.5 %
<b>Ravel: Sonata</b>				2.6 %		38.5 %	59.0 %	97.5 %
<b>Martinu: Duo</b>				2.7 %		27.0 %	70.3 %	97.3 %
<b>Pachelbel: Canon in D</b>		17.6 %		5.0 %			77.5 %	95.1 %
<b>Procaccini: Trois pieces</b>	44.4 %	5.6 %			50.0 %			94.4 %
<b>Bach: Cantata BWV</b>		45.2 %	9.6 %				45.2 %	90.4 %
<b>Sculptured: Fulfillment</b>		25.0 %	11.1 %				63.9 %	88.9 %
<b>Ohana: Flute duo</b>		86.8 %	13.2 %					86.8 %
<b>Bach: Cantata BWV</b>				13.6 %		6.8 %	79.5 %	86.3 %
<b>Pachelbel: Canon in D</b>				15.4 %		0.0 %	84.6 %	84.6 %
<b>Idrs: Aria</b>	43.2 %	16.2 %	2.7 %		37.8 %			59.4 %
<b>Feidman: Klezmer</b>			6.7 %	40.0 %	50.7 %		2.7 %	46.7 %
<b>Copland: Sonata</b>			0.0 %	29.5 %	45.5 %		25.0 %	45.5 %
<b>Guiliani: Iglou</b>		8.1 %		32.4 %	10.8 %		48.6 %	40.5 %

Table 3: Duet Classification using our real-time solo recognition program.

## 7. MULTI-INSTRUMENTAL EXAMPLES

In Table 3 we bring some examples for instrument recognition in real performance duets (where 2 instruments are playing concurrently) using our solo-recognition process with the real-time feature set. Obviously, this section is not pretending to be an extensive research of multi-instrumental classification, but rather comes to demonstrate that successful solo recognition might actually be useful for instrument recognition in multi-instrumental music.

From each real performance duet, a 1-minute section was selected in which both instruments are playing together and each second of this section was classified by our real-time solo recognition program.

The first column in Table 3 contains the partial name of the musical piece. Columns 2 to 7 contain the percentage of solo-cuts in the musical piece which were classified as the corresponding instrument. The white cells indicate correct classifications – recognition of instruments which actually played in the corresponding solo-cuts, while the black cells indicate misclassified solo-cuts. The last column is the total percentage of solo-cuts that were correctly classified as one of the playing instruments.

We can see that there is a considerable number of examples where the classification was correct although, as we know, the classifier is very naïve and does not use  $f_0$  nor attempts to perform any source separation. We shall study in future work the reasons why specific instrument combinations produce more recognition errors and how to improve the recognition of these combinations, e.g. we can see in Table 3 that the guitar was a most common misclassification and that we probably need extra features to discriminate it better.

## 8. SUMMARY

We presented a process for continuous recognition of musical instruments in solo recordings which yields a high recognition rate. Our results are based on evaluation with a large and very diverse solo database which allowed us a wide generalization of the classification and evaluation processes using diverse sound possibilities of each instrument, recording conditions and playing techniques.

We used our GDE feature selection algorithm with a big feature set and considerably reduced the number of features, down to a feature set which allowed us to perform real-time instrument recognition in solo performances. This smaller feature set delivers a recognition rate which is close to that of the complete feature set.

Lastly, we have shown that our recognition process and real-time feature set, without any modifications, can also be useful for instrument recognition in duet music. This exemplifies our initial claim that learning to achieve high recognition rates in solos could also be useful for instrument recognition in multi-instrumental performances.

## 9. FUTURE WORK

We shall continue researching instrument recognition in multi-instrumental music. We intend to study the reasons for correct recognition in some duets and incorrect recognition in others by our solo classifier. We have started working on a multi-instrument recognition process where each solo-cut can be classified as more than one instrument. This process also provides a confidence level for every classification.

We will work on partial “source reduction”, where we shall not attempt to actually separate the instruments but rather to weaken the influence of some of the tones and then use a modified “solo” classifier.

New features will be developed and used in the feature selection process; some of them especially designed with multi-instrumental recognition in mind.

## 10. ACKNOWLEDGMENTS

Thanks to Geoffroy Peeters for using his feature computation routines and sharing his knowledge and experience.

Thanks to Emmanuel Vincent for sharing his solo database.

## 11. REFERENCES

- [1] J. Marques and P. J. Moreno, “A study of musical instrument classification using Gaussian mixture models and support vector machines,” *Cambridge Research Laboratory Technical Report Series*, CRL/4, 1999.
- [2] J. C. Brown, O. Houix and S. McAdams, “Feature dependence in the automatic identification of musical woodwind instruments,” *Journal of the Acoustical Society of America*, Vol. 109, No. 3, pp 1064-1072, 2001.
- [3] K. Martin, “Sound-source recognition: A theory and computational model,” *PhD Thesis*, MIT, 1999.
- [4] P. Herrera, G. Peeters and S. Dubnov, “Automatic Classification of Musical Sounds,” *Journal of New Musical Research*, Vol. 32, No. 1, pp 3-21, 2003.
- [5] J. Eggink and G. J. Brown, “Instrument recognition in accompanied sonatas and concertos,” To appear in *Proc. International Conference on Acoustics, Speech, and Signal Processing (ICASSP '04)*, 2004.
- [6] A. Livshin and X. Rodet, “The Importance of Cross Database Evaluation in Musical Instrument Sound Classification,” In *Proc. International Symposium on Music Information Retrieval (ISMIR '03)*, 2003.
- [7] G. Peeters, “A large set of audio features for sound description (similarity and classification) in the CUIDADO project,” 2003. URL: [http://www.ircam.fr/anasy/peeters/ARTICLES/Peeters\\_2003\\_cuidadoaudiofeatures.pdf](http://www.ircam.fr/anasy/peeters/ARTICLES/Peeters_2003_cuidadoaudiofeatures.pdf)
- [8] G. Peeters and X. Rodet, “Automatically selecting signal descriptors for Sound Classification,” in *Proc. International Computer Music Conference (ICMC'02)*, 2002.
- [9] G. J. McLachlan, *Discriminant Analysis and Statistical Pattern Recognition*. New York, NY: Wiley Interscience, 1992.
- [10] K. Martin and Y. Kim, “Musical instrument identification: a pattern-recognition approach,” In *Proc. 136th Meeting of the Acoustical Society of America*, 1998.
- [11] A. Livshin, G. Peeters and X. Rodet, “Studies and Improvements in Automatic Classification of Musical Sound Samples,” In *Proceedings of the International Computer Music Conference (ICMC'03)*, 2003.