The Thirty-Third International
FLAIRS Conference (FLAIRS-33)

Automatic Identification of Guitar Types from Prerecorded Audio

Lubomir Ivanov
Computer Science Department, Iona College
New Rochelle, NY, 10801 USA
livanov@iona.edu

Jhavagne Dunn
IBM North Castle
1 N. Castle Dr., Armonk, NY, 10504, USA
jhavagne.dunn@ibm.com

Abstract

This paper presents an automatic classification model for the identification of various types/makes of guitars. The classification is carried out by machine learning classifiers trained on mel frequency cepstral coefficients (MFCCs) features extracted from audio recordings. The classification results are analyzed and insights from the experiments are shared.

1. Introduction

Experienced musicians rely on a variety of instruments to help them express specific emotions or to achieve a particular auditory effect. Each type of instrument has its own characteristics, appropriate for certain musical styles and performer preferences and moods – from the soft sounds of a classical guitar to the angry growl of a distorted electric guitar. Many sonic differences are subtle and imperceptible to the casual listener. For example, the much-debated Stradivari vs. Guarneri violin sonic differences are discernible mostly to experienced violinists and select audiophiles. To a trained ear, however, even subtle differences in tone can be quite obvious. To an artist, finding the "perfect sound" is a matter of vision and self-expression. A natural question, therefore, arises: Can a computer identify musical instruments by their sound as readily as a trained musician?

Musical instrument identification is part of the broader topic of automatic music classification, which has been explored extensively (Dieleman et al. 2011, Haggblade et al. 2011). Identifying musical instruments from a mix or solo recordings has been studied for guitar (Johnson et al. 2015), violin (Dalmazzo et al. 2018, Lukasik 2010), piano (Fragoulis et al. 2006), and even drums (Souza et al. 2009). General studies, which attempt to distinguish the sonic characteristics of broad classes of instruments have also been pursued (Hamel et al. 2009, Essid et al. 2006, Agostini et al. 2003). Several software packages for musical feature extraction and classification exist (Moffat et al. 2015).

An ever-popular instrument, the guitar has received its fair share of attention: Researchers have created physical models of guitars and player techniques (Traube and Smith 2001), invented methodologies for automatic recognition of guitar components (Geib et al. 2017) and effects (Schmitt and Schuller 2017) and for extracting guitar-based musical structures (notes, chords, etc.) (Stark and Plumbley 2009, Fragoulis et al. 2006). Some have explored guitar classification (Johnson and Tzanetakis 2015, Freoura et al. 2014).

This paper examines the classification of a set of acoustic and electric guitars in experiments designed to simulate human perception of sound. We present our motivation, methodology and experimental setup, and discuss the results, comparing them to prior work reported in the literature. Directions for further research are outlined at the end.

2. Motivation and Methodology

Our goal was to determine if a computer can distinguish the subtle differences in the sound of different guitars, which a trained human ear can perceive. We wanted to experiment with a variety of guitar models and makes under conditions as close as possible to those a human listener experiences.

2.1 Mel Frequency Cepstral Coefficients

A popular model of human auditory response is based on mel frequency cepstral coefficients (MFCC) used with great success in speech processing. It has been demonstrated that human pitch perception is most accurate in the 100Hz-1000Hz frequency range. Above 1KHz, human hearing accuracy correlates logarithmically with the sound frequency. The mel-scale models this relationship between hearing perception and frequency using units referred to as "mels". Human auditory perception experiments have demonstrated that the human ear acts as a filter for certain sound frequencies. Moreover, these frequencies are non-uniformly spaced

– there are more filters in the low frequency range (100Hz-1KHz) and fewer filters above 1KHz. The MFCCs are the coefficients that make up the mel frequency cepstrum – a mel-scale filter model of the power spectral envelope of a single audio signal frame (typically 20-50ms long).

The MFCC extraction process consists of the following steps: An audio signal is initially broken into short frames (25ms in our case). After applying a Fast Fourier Transform, the obtained frequency components are mapped to the mel scale by using 26 triangular filters. The logarithm of all filterbank values is taken and the Discrete Cosine Transform is applied to decorrelate the frequency components. In the end, only the first 13 MFCC coefficients are kept, since the remaining ones represent higher frequencies to which the human ear is less sensitive. Thus, an audio signal is represented as a sequence of 13-dimensional MFCC vectors.

2.2 Classification Methodology

To carry out guitar identification, we used support vector machines with sequential minimal optimization (SMO) and multilayer perceptrons (MLP) implemented in the WEKA machine learning software (Hall et al. 2009) We also conducted experiments with random forest classifiers, but the results were weaker than those using SMOs and MLPs.

3. Experimental Setup

The guitars used in our experiments are described in Table 1. They were selected to have a broad mix of physical features – different body/neck/fingerboard woods, pickups, bridges, string materials and gauges. We included guitars, which sound quite distinctly as well as guitars, which sound similarly. The electric guitars were recorded using an Orange Crush 20ldx solid state amplifier. No effects were used except for a 2006 Lovepedal Eternity E6 overdrive pedal for the distorted guitar signals. The acoustic guitars were not amplified. All guitars used the standard E-A-D-G-B-E tuning. All audio was recorded with an iPhone 5 placed 30 inches away from the amplifier. We purposefully did not use professional studio equipment to recreate the real-life experience of a typical listener. In that respect, our experiments are different from all other research we have come across. The recorded audio consisted of 4-second m4a clips, containing one of the following:

- A single clean or overdriven (distorted) note (A through G) in the lower octave played on the A string.
- A single clean or overdriven note (A through G) in the higher octave played on the high-E string.
- An open major or minor chord (A through G), except for the barred Cm, Gm, and F chords. The chords were recorded clean only (no overdrive).
- An overdriven (distorted) three-note power chord (root note, 5^{th} note, and root note in the next octave).

The audio clips were recorded for each of the seven guitars except no overdriven notes or power chords were recorded for the classical and the 12-string acoustic guitars. The Audacity sound processing software (Audacity Team 2019) was used to cut the original m4a clips into 25ms frames and store them as wav files for feature extraction. The MFCC vectors were extracted using Python programs we wrote based on the methodology described in section 2.2. The extracted MFCC vectors were recorded into WEKA ARFF file used for training the machine learning classifiers.

Table 1. Guitars used in the experiments

Guitar	Type	Model
Guitar A	Electric	Sterling by Music Man Axis AX40, maple body and neck, dual humbucking pickups, Floyd-Rose tremolo bridge, Fender 250LR .09-.46 nickel-plated steel strings.
Guitar B	Electric	1996 Fender Stratocaster (MIJ), Blackmore model, alder body, maple neck with scalloped rosewood fingerboard, three Lace sensor gold single-coil noiseless pickups, tremolo bridge, Fender 250LR .09-.46 nickel-plated strings.
Guitar C	Acoustic	Cordoba Dolce 7/8 Classical Guitar, cedar top, mahogany back and sides, mahogany neck with rosewood fingerboard, Savarez 500CJ high-tension nylon strings.
Guitar D	Electric	2007 Fender Deluxe Stratocaster (USA), alder body, maple neck with rosewood fingerboard, three samarium cobalt single-coil noiseless pickups, .10-guage steel strings.
Guitar E	Electric	2016 Gibson Les Paul Traditional, mahogany body with maple top, mahogany neck, rosewood fingerboard, dual '57 classic humbucking pickups, Nashville tune-0-matic bridge and stopbar, Gibson light .10-.46 strings.
Guitar F	Electric	1980s Asama (a Japanese Fender Stratocaster clone), basswood body, maple neck and fingerboard, DiMarzio Area '58 (neck), Area '67 (middle), and Area '61 (bridge) pickups, Fender 250LR .09-.46 nickel-plated strings.
Guitar G	Acoustic	Seagull Coastline Series S-12 dreadnought guitar, cedar neck, cherry back and sides, silver leaf maple neck, steel 12-string set.

4. Results and Discussion

We conducted a large number of machine learning experiments using the MFCC vector sequences extracted from the guitar sound recordings described in section 3. Initially, we experimented with pairwise identification of guitars, testing each guitar against all others. Next, we experimented with sets of three guitars, focusing on comparing humbucker pickup guitars to single-coil pickup guitars as well as comparing similarly sounding guitars to each other. Finally, we attempted a concurrent identification of all electric guitars and then the full set of guitars. We used SMO and MLP classifiers with leave-one-out (L1O) validation.

Overall, the results (Table 2) exceeded our expectations. In the head-to-head comparisons, most results were in the upper 90th accuracy percentile for both clean and distorted signals. In many cases the classifiers identified distorted guitars slightly better than the same guitars played clean. As expected, the acoustic guitars, especially the classical Cordoba nylon string guitar, were readily identifiable. The average SMO and MLP accuracies for the Cordoba guitar in the pairwise clean guitar comparisons were 95.17% and 96.83% respectively. For the 12-string Seagull guitar, the average SMO and MLP pairwise clean comparison accuracies were 93% and 98.17%. In the experiments with distorted guitar signals, the acoustic guitars were identifiable even more readily: The average SMO/MLP accuracies for the nylon string guitar were 95.33% and 97.83% respectively, and for the 12-string guitar - 93% and 97.5%.

The next interesting question was the distinction between electric guitars with humbucking vs. single-coil pickups. To a trained human ear, the humbucking sound has a richer frequency spectrum and timbre. The single-coil pickups tend to sound brighter (usually described as a "bell-type" sound). In the clean signal humbucking-vs-single-coil experiments the SMO/MLP classifiers obtained 90.5% and 93.75% average accuracies respectively. The average accuracies in the distorted signal experiments were lower – 83.75% for SMO and 87.25% for MLP. This is not surprising since the distorted signals have many more overlapping overtones, which makes the identification difficult even for a trained musician. Table 2 reveals that the average accuracy is lower when comparing the American Fender Stratocaster to the Gibson and Axis guitars. For the other two Stratocasters, the identification produces high accuracies. The most likely explanation is that the samarium cobalt single-coil pickups of the American Stratocaster have richer dynamics than the traditional single-coil pickups, and resemble the frequency response of humbucking pickups.

Next, we compared the three Stratocaster models – first two-by-two, then all three simultaneously. The average accuracies in the pairwise tests were 69.33% (SMO, clean), 80% (MLP, clean), 91.67% (SMO, distorted), and 92.33% (MLP, distorted). In the three-way Stratocaster test, the accuracies were 57% (SMO, clean), 74% (MLP, clean) 83% (SMO, distorted), and 84% (MLP, distorted). As expected, the accuracies were lower overall, though, interestingly, the distorted signal experiments produced stronger results than the clean signal experiments. This is likely due to the interaction between the guitar pickups and the overdrive pedal, which alters the pickups' tonal characteristics when engaged. By comparison, the pairwise test of the two humbucking guitars produced strong identification results 95% (SMO, and MLP, clean), 83% (SMO, distorted), and 86% (MLP, distorted). Apparently, the tonal characteristics of Axis and Gibson humbuckers are sufficiently different to be easily identifiable by the classifiers.

Table 2: Classification accuracies.

Guitars	**SMO Accuracy** (Clean/Distorted)	**MLP Accuracy** (Clean/Distorted)
12-String – Axis	100% / 100%	100% / 100%
12-String – Blackmore	88% / 91%	98% / 98%
12-String – Cordoba	94%	95%
12-String – Fender (USA)	91% / 87%	99% / 98%
12-String – Asama	96% / 97%	99% / 98%
12-String – Gibson	89% / 89%	98% / 96%
Axis – Blackmore	92% / 93%	98% / 96%
Axis – Cordoba	100% / 100%	100% / 100%
Axis – Fender (USA)	95% / 85%	98% / 73%
Axis – Gibson	95% / 83%	95% / 86%
Axis – Asama	94% / 95%	94% / 94%
Blackmore – Asama	57% / 83%	74% / 83%
Blackmore – Cordoba	96% / 96%	98% / 99%
Blackmore – Fender (USA)	82% / 94%	91% / 97%
Cordoba – Fender (USA)	86% / 87%	94% / 96%
Cordoba – Gibson	97% / 96%	96% / 98%
Cordoba – Asama	98% / 99%	98% / 99%
Gibson – Asama	91% / 94%	96% / 100%
Fender (USA) – Asama	69% / 98%	75% / 97%
Fender (USA) – Gibson	82% / 61%	87% / 82%
Fender (USA) – Blackmore – Asama	57% / 83%	74% / 84%
Axis – Fender (USA) – Gibson	81% / 70%	86% / 75%
Axis – Cordoba – Fender (USA) – Gibson	95%	91%
Axis – Cordoba – Fender (USA) – Gibson	87%	79%
All electrics guitars	61% / 51%	74% / 64%
All seven guitars	59% / 49%	76% / 70%

The all-electric-guitars and the all-guitars tests produced lower accuracies. An examination of the confusion matrices reveals that this is due primarily to the misclassifications between the different Stratocaster models and between the Fender (USA) and the humbucking Gibson and Axis guitars.

As can be readily seen from Table 2, in almost all experiments the MLP classifier outperformed the SMO classifier, in some cases by a wide margin. However, the MLPs took significantly longer to train compared to the SMOs.

The results reported in this study are fairly consistent with other reported results in the literature. Closest to our work is that of Johnson and Tzanetakis (2015), who presented a study of the timbral properties of fourteen acoustic and three electric 6-string guitars. All guitars in that study used steel strings and were recorded clean, with no effects. Johnson and Tzanetakis used support vector machines (SVM) and k-nearest neighbor (kNN) classifiers trained on linear predictive cepstral coefficients (LPCC) and MFCC. The authors reported results in the 50% accuracy range, with one experiment producing an accuracy in the upper 70% range. Their best results were obtained using an SVM classifier.

The primary difference with our work is our emphasis on emulating the human experience of identifying guitars. To that end, we used a more diverse (though smaller) set of guitars including a classical nylon string guitar and a 12-string dreadnought, an overdrive pedal, and no professional studio recording equipment. To further recreate the human experience, we focused exclusively on MFCCs. Instead of kNN, we used SMO and MLP classifiers that produced stronger results, especially in the pairwise guitar comparisons.

4. Conclusion and Future Work

In this paper we presented an empirical study of guitar identification from recorded single-instrument audio. The results of the experiments confirm that guitar identification can be successfully carried out by a computer and provide insights into the timbral characteristics of different types of guitars. The study raised numerous additional question:

- How much effect do the structural materials (wood, plastic, metal parts) have the tonal characteristics of guitars?
- To what extent does the choice of strings materials or gauge affect guitar identification?
- Can better results be obtained from recordings of alternate tunings or more advanced chords (9's, 11's, sus-2, etc.)?
- To what extent do effects alter the tonal characteristics of the instruments and how is identification impacted?
- Is guitar identification affected by the choice of amplifier?
- Will an ensemble classifier produce stronger identification results even for large guitar collections?

In addition to addressing the questions above, we intend to expand our experimentation to include additional string instruments (violins, cellos, ukuleles, banjos, etc.). We are also working on a methodology for extracting single-instrument sound signals from a polyphonic mix, to allow us to extract large collections of sounds from numerous instrument makes and models without having to carry out individual instrument recording. This will not only provide more data for accurate training of the machine learning classifiers but will also make the process of identifying instruments closer to the experience of a human listener.

References

Agostini, G., Longari, M., Pollastri, E. 2003. "Musical instrument timbres classification with spectral features," *EURASIP J. Appl. Signal Process.*, vol. 2003, pp. 5–14, Jan. 2003. [Online]. Available: http://dx.doi.org/10.1155/S1110865703210118

Dalmazzo, D., Tassani, S., Ramirez, R., 2018, A Machine Learning Approach to Violin Bow Technique Classification: A Comparison Between IMU and MOCAP systems. 1-8.10.1145/3266157. 3266216.

Dieleman, S., Brakel P., Schrauwen, B. 2011. Audio-based music classification with a pretrained convolutional network. *12th International Society for Music Information Retrieval Conference (ISMIR-2011)*. University of Miami, Florida, USA.

Essid, S., Richard, G., David, B.2006. Musical instrument recognition by pairwise classification strategies.. *on Audio, Speech, and Language Processing,* , vol. 14, no. 4, pp. 1401–1412.

Fragoulis, D. et al. 2006. Automated classification of piano-guitar notes. In *Audio, Speech, and Language Processing, IEEE Transactions on*, vol. 14, no. 3, pp. 1040–1050.

Freoura V. et al, 2014. Evaluation and classification of steel string guitars using bridge admittances, ISMA'14, Le Mans, France.

Geib, T., Schmitt, M., BjörnSchuller, B. 2017. Automatic Guitar String Detection by String-Inverse Frequency Estimation, In *Lecture Notes in Informatics,* Gesellschaft für Informatik, Bonn,p.127.

Haggblade, M., Hong, Y., Kao, K. 2011. "Music genre classification." *Department of Computer Science, Stanford University.*

Hall M., Frank E., Holmes G., Pfahringer B., Reutemann P., Witten I. 2009. The WEKA Data Mining Software: An Update; *SIGKDD Explorations*, Volume 11, Issue 1

Hamel, P., Wood S., Eck, D. 2009. Automatic Identification of Instrument Classes in Polyphonic and Poly-Instrument Audio. *ISMIR*

Johnson, D., Tzanetakis, G. 2015. Guitar Model Recognition from Single Instrument Audio Recordings, @article {Johnson2015 GuitarMR, In *IEEE Pacific Rim Conference on Communications, Computers and Signal Processing (PACRIM)*, pp. 370-375.

Lukasik, E. 2010. Long Term Cepstral Coefficients for Violin Identification. *AES Convention*, Volume: 128

Moffat, D., Ronan, D., Reiss, J. 2015. An evaluation of audio feature extraction toolboxes. In *Proceedings of the 18th Int. Conference on Digital Audio Effects (DAFx-15), Trondheim, Norway.*

Schmitt, M., Schuller, B., 2017. Recognizing Guitar Effects – Which Acoustic Features Really Matter?. In *Lecture Notes in Informatics (LNI)*, Gesellschaft für Informatik, Bonn 2017, 177.

Souza, V., Batista, G., Souza-Filho, N. 2015. "Automatic classification of drum sounds with indefinite pitch." *Neural Networks (IJCNN), 2015 International Joint Conference on*. IEEE, 2015.

Stark, A., Plumbley, M. 2009. Real-Time Chord Recognition for Live Performance. In *International Computer Music Conference*.

Audacity Team (2019). Audacity(R): Free Audio Editor and Recorder [Computer application]. Version 2.3.2 retrieved May 20th 2019 from https://audacityteam.org/

Traube, C., Smith, J. O. 2001. Extracting the fingering and the plucking points on a guitar string from a recording, *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*.

Van Steelant, D et al. 2005. Support vector machines for bass and snare drum recognition. In *Classification—the Ubiquitous Challenge*. Springer, Berlin, Heidelberg, 2005. 616-623.