
Research article

Application of Musical Information Retrieval (MIR) Techniques to Seismic Facies Classification. Examples in Hydrocarbon Exploration

Paolo Dell'Aversana^{*}, Gianluca Gabbriellini, Alfonso Iunio Marini, Alfonso Amendola

Eni SpA, San Donato Milanese, Milan, Italy

*** Correspondence:** Email: paolo.dell'avversana@eni.com; Tel: +39-02-52063217; Fax: +39-02-520-63897

Abstract: In this paper, we introduce a novel approach for automatic pattern recognition and classification of geophysical data based on digital music technology. We import and apply in the geophysical domain the same approaches commonly used for Musical Information Retrieval (MIR). After accurate conversion from geophysical formats (example: SEG-Y) to musical formats (example: Musical Instrument Digital Interface, or briefly MIDI), we extract musical features from the converted data. These can be single-valued attributes, such as pitch and sound intensity, or multi-valued attributes, such as pitch histograms, melodic, harmonic and rhythmic paths. Using a real data set, we show that these musical features can be diagnostic for seismic facies classification in a complex exploration area. They can be complementary with respect to “conventional” seismic attributes. Using a supervised machine learning approach based on the k-Nearest Neighbors algorithm and on Automatic Neural Networks, we classify three gas-bearing channels. The good performance of our classification approach is confirmed by borehole data available in the same area.

Keywords: Seismic facies classification; sonification; audio video display; spectral decomposition; pattern recognition; musical information retrieval; MIDI

1. Introduction

Pattern recognition can be defined as “the scientific discipline whose goal is the classification of objects into a number of categories or classes” [1]. Nowadays, modern computer technology allows a wide range of applications of pattern recognition in many scientific and business sectors. These include man-machine communication, speech and image recognition, analysis of market trends, big data mining, automatic categorization, decision-making processes and so forth. Despite the large variability of

algorithms and objectives in the different fields of application, detection and classification by automatic pattern recognition show many common features. In fact, the core of each pattern recognition system is commonly represented by a mathematical formulation for a decision function aimed at categorizing information into several classes.

Nowadays, dealing with big data is the normality in geosciences. For instance the process of finding and producing hydrocarbons generates a huge amount of heterogeneous information. Moreover, volcanic surveillance, earthquake activity monitoring, environmental sciences represent activities that produce big data sets. In these cases, pattern recognition and automatic classification approaches can be useful. Indeed, they have been widely used also for seismic facies classification. In fact, especially when large 3D seismic data sets must be analyzed, it can be useful (or necessary) to support human interpretation with automatic approaches of data mining, recognition, detection and classification [2,3].

Pattern recognition is commonly based on the extraction of several features from the data. A feature is a specific piece of information that can be used for characterizing the data itself. In geophysics these are commonly indicated as attributes. These can be related to amplitude, frequency, geometrical reflector configurations, lithology, geomechanical properties and so forth. The selected attributes are used for classification. This is the process by which each voxel is assigned to one of a finite number of classes (called clusters). The term “clustering” indicates the operation of grouping patterns of signals showing similar values of one or more selected attributes. For instance, seismic reflections showing similar geometrical features and comparable frequency content can be clustered in the same seismic facies. Pattern recognition and data clustering can make use of *unsupervised or supervised learning algorithms*. In the first case, the interpreter provides no prior information other than the selection of attributes and the number of desired clusters. Alternatively, supervised learning classification requires a phase of *machine training*.

A different field where automatic clustering and classification is very useful is digital music. Indeed, interesting applications have been done over the past decade in Musical Information Retrieval (MIR) and musical genre classification. The classification approaches used in digital music are similar to those used in geophysics. In fact, the basic workflow starts from extracting characteristic pieces of information, known as “features”, from music. These musical features are used for clustering the musical pieces into homogeneous classes, commonly called *genres*. Musical features typically fall into one of three main categories: low-level, high-level and cultural features. Spectral or time-domain information extracted directly from audio signals belong to the first category. Instead, melodic contour, chord frequencies and rhythmic properties are typical examples of high-level features. Finally, sociocultural information outside the scope of musical content itself is the main type of cultural feature.

Due to their strict similarity, clustering algorithms used in MIR can be exported, adapted and applied into the geophysical domain. In this paper we discuss a novel methodology of pattern recognition and automatic classification of seismic data applying the same algorithms used for MIR and musical genre classification. In particular, we use clustering algorithms working on MIDI^[1] features. We will show that many benefits are introduced by this cross-disciplinary approach testing it on a real seismic data set.

2. Links Between Geophysics and Digital Music

In previous works we demonstrated that geophysics and digital music can be linked, with the advantage of sharing technologies and approaches that are commonly applied separately in both domains [4,5]. Recently we discussed the details of our approach based on sonification, for analysis and interpretation of geophysical data [6]. Our workflow includes accurate time–frequency analysis using Stockwell Transform (or S-Transform). The original formulation of this transform applied to a generic time dependent signal $x(t)$, is given by [7]:

$$S(\tau, f) = \frac{|f|}{\sqrt{2\pi}} \int_{-\infty}^{\infty} x(t) e^{-i2\pi ft} e^{-\frac{(t-\tau)^2 f^2}{2}} dt \quad . \quad (1)$$

In the formula (1), τ is the time where the S-Transform is calculated and f is the instantaneous frequency. The exponential function in the integral is frequency dependent, consequently the windowing function scrolling the time series is not constant but depends on the frequency. Thus, this type of mathematical transform is appropriate when instantaneous frequency information changes over the time and must be preserved, such as in seismic signal analysis. Figure 1 shows a visualization of a real seismic trace and the relative Stockwell spectrogram.

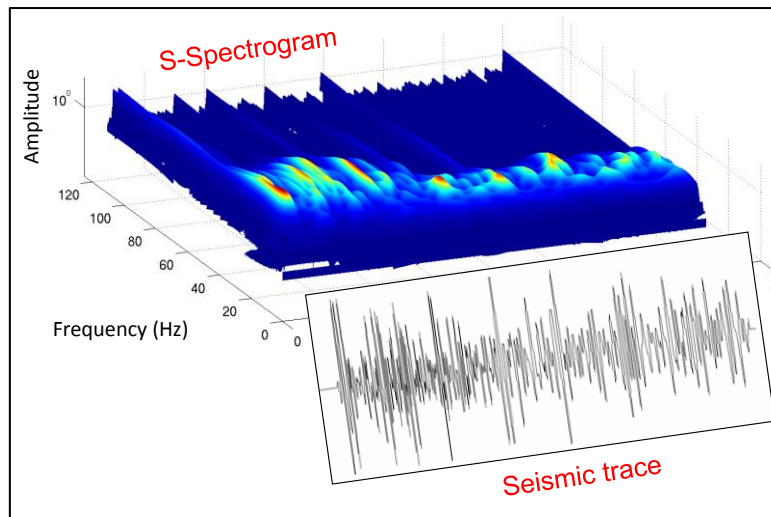


Figure 1. Example of a spectrogram of a seismic trace obtained through application of Stockwell Transform (after Dell’Aversana et al., 2016).

The next step for moving from the geophysical to the musical domain is to transform the physical quantities of a spectrogram (frequency, time, and amplitude) into MIDI attributes (such as pitch, velocity or sound intensity, note duration and so forth). The mathematical relationship between the frequency f and the MIDI note number n is the following:

$$f(n) = 440 \cdot 2^{(n-58)/12} \quad . \quad (2)$$

In equation (2), the symbol n indicates the sequential number of MIDI notes. For instance,
 $n = 1$ corresponds to the note C_0 (16.35 Hz),
 $n = 2$ corresponds to $C\#_0$ (17.32 Hz),
 $n = 108$ corresponds to B_8 (7902 Hz) ...

Once converted in MIDI format, geophysical datasets can be transposed into the audible frequency range, played and interpreted by using modern computer music tools, such as sequencers. Audio analysis is performed simultaneously with interpretation of images, as a complementary tool.

The MIDI standard is well suited to be linked with a time–frequency representation of a signal. Consequently, the conversion to MIDI can bring meaningful information from a physical point of view and can be considered an advanced example of sonification of geophysical data. We applied our approach to real seismic data sets. An interesting e-lecture including examples of audio-video display can be found at the following link: <https://www.youtube.com/watch?v=tGhICX2stTs>.

3. Application of MIR to Seismic Data Classification

In our previous works [5] we introduced the idea that standard seismic data sets converted into MIDI format can be analyzed and classified using musical pattern recognition approaches. In the following paragraphs we will expand that idea, showing applications to real data.

Algorithms capable of performing automatic recognition and classification of musical pieces are becoming increasingly useful over the past decade. The reason is because networked music archives are rapidly growing. Consequently, individual users, music librarians and database administrators need to apply efficient algorithms to explore extended music collections automatically. The science of retrieving information from music is known as Music Information Retrieval (MIR). It finds the main applications in instrument recognition, automatic music transcription, automatic categorization and genre classification. The main goal of MIR algorithms is to recognize occurrences of a musical query pattern within a musical database. This can be a collection of songs or of other types of musical pieces. The same approach can be used for mining every type of data base that includes information codified as sounds. In our approach addressed to geophysical data classification, the database consists of sounds obtained by conversion of geophysical data into musical formats. Our idea is to adapt efficient and established MIR approaches to geophysical purposes, including seismic facies recognition and classification.

Clustering and classification can be performed extracting from the data both audio and “symbolic” features (also called audio and symbolic attributes, respectively). In fact, musical data can be stored digitally as either audio data (e.g. wav, aiff or MP3) or symbolic data (e.g. MIDI, GUIDO, MusicXML or Humdrum). Audio formats represent actual sound signals because analog waves are entirely encoded into digital samples. Instead, symbolic formats allow storing just instructions for reproducing the sounds rather than actual waves.

Examples of audio features are the pitch (related to the frequency) the timbre, and the energy (a measure of how much signal there is at any one time). Other useful audio features are the cepstral coefficients. These are obtained by applying Fourier transform to the log-magnitude Fourier spectrum and have been frequently used for speech recognition tasks.

MIDI protocol includes basic attributes like pitch and sound intensity (called “velocity”). Furthermore, there are “high-level” MIDI features based on musical texture, rhythm, density of musical notes (average number of notes per second), melodic trends and note duration, pitch statistics and histograms of sound intensity, and so forth. Both audio and symbolic formats show benefits and limitations.

Audio formats allow preserving the full waveform information, but at expenses of high memory requirement. On the other side, MIDI recordings contain only instructions to send to a digital

synthesizer for reproducing the sounds. Thus, the original geophysical information can be preserved only partially. Information loss can be negligible only if the transformation from geophysical signals to MIDI instructions is extremely accurate, like in our approach [6]. However, MIDI files are easier to store and much faster to process than audio data. Consequently, MIDI format is suited for classification workflows applied to big data sets. MIR algorithms based on high-level features work efficiently on MIDI attributes.

A crucial point is that, many new attributes can be introduced into the geophysical domain after transforming geophysical signals into MIDI format. Some of these features have no equivalent attributes in the seismic domain. In fact, differently from seismic attributes, symbolic representations like MIDI instructions can store high-level musical information. The concept of “high-level features” refer to pieces of information that are based on musical abstractions, such as rhythmic, melodic and harmonic patterns, chord sequences and other significant structures. Instead, low-level features refer to signal-processing basic characteristics derived from signals, like instantaneous frequency or phase. It is reasonable to expect that clustering algorithms will work with improved efficiency when high-level features are included in the classification process.

More than hundred features can be extracted from the geophysical data converted into MIDI files. Although not all of these features are necessarily useful for classification purposes, many different combinations of these attributes can be used for clustering the data in homogeneous classes. A typical application is to cluster/classify a seismic data set in different categories addressed to seismic facies identification. The data set is scanned trace by trace, with the objective to assign different data segments to distinct classes.

An additional benefit derived from MIR algorithms is that they are commonly designed and applied for extracting information in very noisy environmental conditions. Indeed, they must be able to recognize a song in busy and buzzing environments. For instance, SHAZAM, one of the most popular apps in the world for Macs, PCs and smartphones, has acknowledged music identification capabilities, actually expanded to integrations with cinema, advertising, TV and retail environments [8]. It has been used to identify 15 billion songs. Thus, we can argue that similar algorithms (and/or analogous approaches working on MIDI files) can be suited for detecting and classifying “geophysical sounds” extracted from noisy data bases. Indeed, our tests, discussed in the following sections, seem to support that possibility.

4. Classification Approach

There are many possible classification methods [2,9]. In our MIR approach applied to geophysical data, we tested both unsupervised and supervised classification methods based on statistical pattern recognition and machine learning [10]. Finally, we decided to apply a supervised approach to our real data, because it allowed us controlling better the classification workflow, taking in account for prior information. In fact, in the case of unsupervised methods, the interpreter provides no prior data other than the selection of features and the number/type of desired classes (taxonomy). Instead, supervised learning classification requires a phase of *training*, where the interpreter uses a *training data set* that helps the pattern recognition system to decide what is typical of a particular class.

A simple supervised method is the k-Nearest Neighbors algorithm (or k-NN for short). It is used for both classification and regression. In classification problems, an object is classified taking in account for the properties of its neighbors (commonly forming the training set). Training is

performed by storing the coordinates of each training sample in a multidimensional feature space. Then, a sample is classified by examining the selected features in an ensemble of the nearest k training points surrounding the sample point in the feature space. The “probability” that a sample belongs to a given category is estimated by the number of training points captured divided by k . That estimation represents a “classification score” assigned to each sample. The value k is frequently defined as a function of the number of the training samples (commonly the square root).

K-NN classifiers offer advantages in training speed; however, they are limited in modeling relationships between features. In fact, their classification approach is based entirely on some type of *distance* that expresses the similitude of records estimated on one or more features. That distance cannot take in account for conditional relationships between the features. A different category of classifiers that takes in account for these relationships is a set of techniques falling under the label of Artificial Neural Networks (ANNs). These consist of units (emulating human neurons) connected by links, each with an associated weight. Learning in ANN's takes place by modifying the values of the weights through an iterative optimization process.

ANNs have been already used in geophysics for many different purposes, such as lithofacies recognition, formation evaluation from well-logs data, automated picking of seismic events (reflections and first arrivals), identification of causative source in gravity and magnetic applications, pattern recognition and classification in geomorphology, target detection and big data mining in hydrocarbon exploration [11]. However, our approach is new in the geosciences domain, because ANNs have been never applied to classify geophysical data transformed into MIDI files. Using MIDI attributes, we verified that both K-NN and ANNs classifiers work very quickly and efficiently.

An Artificial Neural Network is commonly defined by three types of parameters: a) the interconnection pattern linking the different layers of units (neurons); b) the learning process for updating the weights of the interconnections; c) the activation function that converts a neuron's weighted input to its output activation. Multilayer feed-forward ANNs represent the type most frequently used for supervised learning for the purpose of classification. Inspired by neurobiological processes in the human brain, these networks are composed of multiple layers of units. The individual units deployed on each layer are connected by links and their associated weights. In supervised training, the user provides both the inputs and the outputs. In our application, we derived the training sub-set using seismic traces calibrated by borehole data and transformed into MIDI format. These represent “labeled” records from which a selected set of MIDI features are extracted. The network then processes the inputs and compares its resulting outputs (classification results) against the desired outputs. That difference allows calculating an *error function*. This is then propagated back through the system, with the final goal to adjust iteratively the weights until the output values (output classification result) will be closer to the “correct” values (correct classification). In other words, the weights are iteratively adjusted in order to make the ANN output closer to the model (labeled records) through some type of optimization process (such as a gradient descent algorithm). In such a way the artificial network can learn how to classify the rest of the data set. A classification score is commonly estimated based of the error function at the end of the process.

5. Example

We tested our automatic classification approach on a seismic data set calibrated by several wells drilled in correspondence of gas bearing formations. These form a multi-layer reservoir consisting of

Oligocene, Eocene and Paleocene sand channels with high gas saturation at depth between 1.5 km and 2.0 km below the sea floor. The sedimentary sequence above the reservoir includes Miocene and Pliocene turbidites. Below the reservoir, the sequence continues with carbonate formations.

5.1. Taxonomy

The first step of our classification workflow was to set a reasonable “taxonomy”. It means that we fixed in advance the number and the type of classes in which our data set should be clustered. That choice is subjective, however it can be properly based on robust geological criteria. In our case, several wells have been drilled in the area of the test, thus our taxonomy was driven by the analysis of both seismic and borehole data. We assumed that the data set could be clustered in four distinct classes related to corresponding gas-bearing paleo channels drilled by the wells. Furthermore, we introduced an additional class including the seismic facies outside the paleo channels. Figure 2 shows an example of seismic data where the different gas-bearing channels have been drilled. They are here labeled, respectively, as Channel “A”, “C” and “D”. There is another channel labelled “Channel B” not included in this figure. Our taxonomy represents just an initial hypothesis that can be improved after the first classification trial.

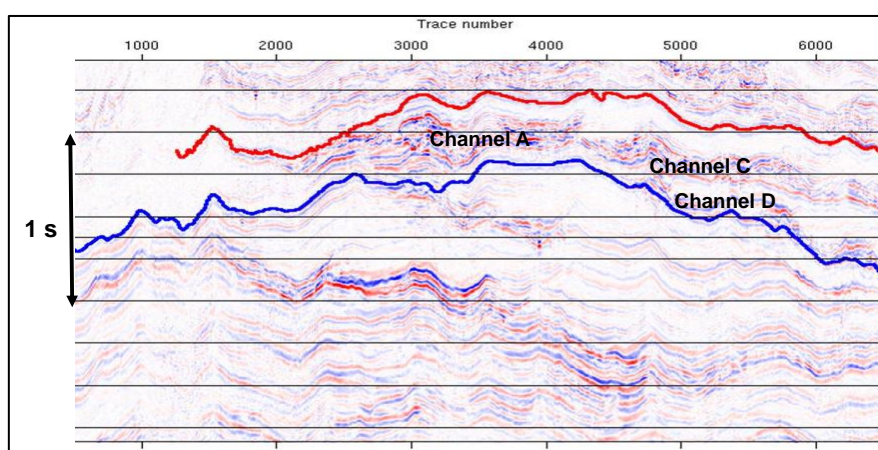


Figure 2. Portion of the seismic data used in the classification test. The red and blue lines shows, respectively, the top and bottom of the seismic volume used in the test. Horizontal axis: trace number. Inter trace distance is 12.5m. Vertical axis: two-way times (s).

5.2. Feature Extraction

After defining an initial taxonomy, we performed the phase of feature extraction. We extracted MIDI features related to pitch, sound intensity, note duration, melodic, harmonic and rhythmic patterns. We started with classification tests based on basic features commonly used in “standard” seismic attribute analysis. These are, for instance, features related to frequency and signal intensity. Then we progressively increased the number of the features, including also multi-valued features that have no equivalent “conventional” seismic attributes. For instance we included new features related to melodic, harmonic and rhythmic patterns. This approach was useful also for comparing the classification effectiveness of standard attributes vs. high-level MIDI features.

In order to clarify the meaning of some MIDI features and their diagnostic power, Figure 3 shows an example of MIDI display of a seismic trace. Here, two among the most relevant MIDI features are showed. These are pitch and “MIDI velocity”, related to instantaneous frequency and sound intensity, respectively. That type of display is known as “MIDI piano roll”. In the upper panel, the vertical axis represents the pitch. This is extracted from the spectrogram obtained by applying Stockwell transform to the seismic traces. I remark that the original frequency content of the seismic data has been transposed into the audible frequency range, so that we can listen to the derived MIDI file.

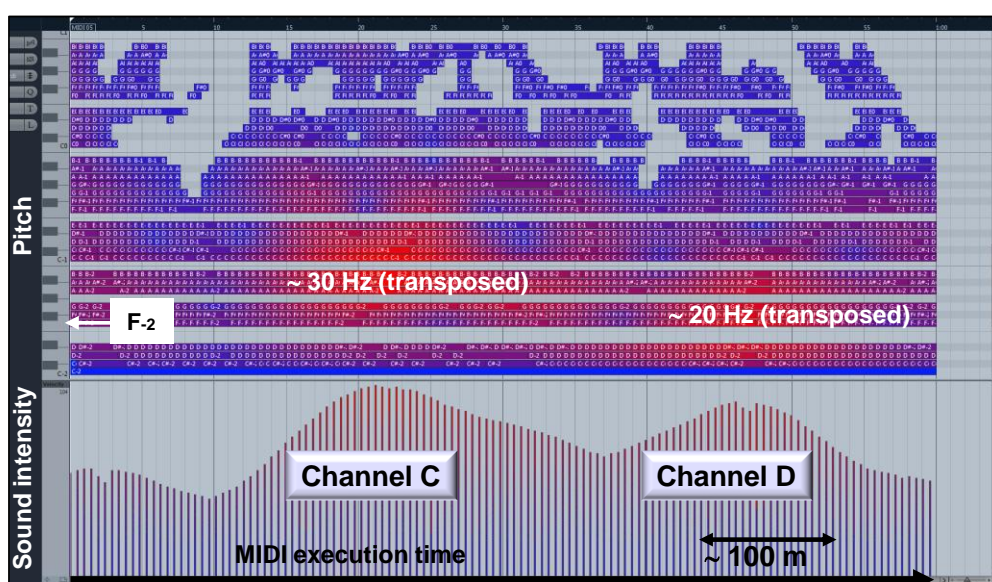


Figure 3. MIDI Piano roll display of a seismic trace crossing channels C and D (see text for details). Depth increases from left to right, consistently with the MIDI execution time.

The pitch is represented using a virtual keyboard shown on the left boundary of the figure. The different colors indicate sound intensity associated to the musical notes: red corresponds to high values and blue represents low values. This type of display is a sort of discrete spectrogram where each pixel represents a musical note with a certain sound intensity. The horizontal axis represents the “MIDI execution time (settable by the user). The direction from left to right corresponds to increasing travel time or, equivalently, increasing depth. Furthermore, the histogram of the sound intensity vs. time, is shown in colors on the bottom of the figure.

This particular seismic trace is interesting because it crosses two distinct channels, “C” and “D”, showing different spectral contents. In the figure, we can see that the deepest gas-bearing channel shows a frequency content lower than the other. Its dominant frequency is around 20 Hz, corresponding to the musical note F_2 in the piano roll (after proper pitch transposition). Instead the dominant frequency of the upper channel is about 30 Hz. That difference can be interpreted as a different attenuation effect caused by different gas saturation in the two channels. That interpretation was driven by analysis of borehole data used for calibration purposes. The ensembles of musical notes in the two gas-bearing channels create musical patterns and represent an important class of “polyphonic MIDI feature”.

This type of multivalued spectral feature (MIDI pattern) can be further analyzed using a different type of display, as showed in Figure 4. This is an example of four adjacent seismic traces

converted into MIDI format and displayed as *pitch histograms*. The vertical axis represents depth, in meters. A vertical segment of about 500 m has been selected for each trace, crossing the same two gas-bearing channels (again labeled as “Channel C” and “Channel D”). Now, colors represent the different pitches. The scale on the left shows that a different color is assigned to each musical note in every octave. For instance, red is assigned to the musical note *C*, yellow is assigned to the musical note *F*, and so forth. The bar length in the histograms is proportional to sound intensity.

Pitch histogram is a useful MIDI “multi-valued, or polyphonic feature” that provides a discrete representation of the frequency content of seismic trace segments. It can be diagnostic for distinguishing different seismic facies. For instance, looking at Figure 4, we can see that each individual channel shows its peculiar pitch histogram. Consequently, it is reasonable to expect that the pitch histogram can contribute to distinguish/classify different channels.

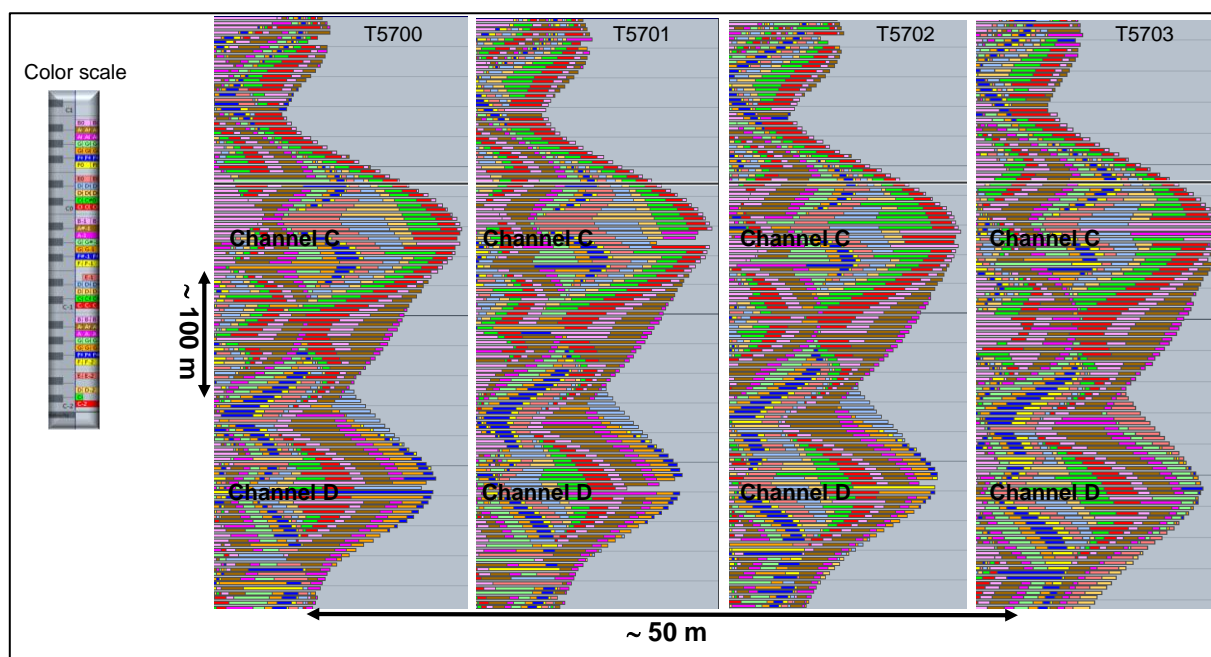


Figure 4. Pitch histograms for 4 different seismic traces crossing channels C and D.

Of course, figures 3 and 4 are just different modalities to represent seismic information in the frequency-time domain (spectrograms). However, they can add value for data analysis and interpretation for several reasons. First, they represent MIDI files, thus they can be played using a sequencer, adding a sound dimension commonly not included in any standard seismic data analysis. An interesting example of audio-visual display of a seismic trace transformed into MIDI spectrogram and Pitch histogram can be found and listened at the following link: <https://www.youtube.com/watch?v=BmyALYtD0WA>.

An additional benefit is that MIDI spectrograms and Pitch histograms can be analyzed, processed, filtered, combined, mixed, stacked and quickly integrated with other information, using the advanced technology commonly applied in the industry of digital music. Many analysis tools are included in the software packages available on the market at relatively low cost. Of course, these cannot substitute the software platforms commonly used for seismic data analysis. However, they can represent a useful complementary tool kit.

5.3. Training

A significant percentage of the seismic traces extracted at locations near the wells have been used as a training data set. We performed several trials by testing different percentages of data to be used for training purposes. We selected empirically a subset of data sufficient for obtaining satisfactory training results. We used all the selected features for training the classification algorithm, starting from the simplest attributes and moving progressively to polyphonic features. In order to decide the final set of features to use in the classification, we conducted many tests to verify the relative performance of each ensemble of attributes. We performed two types of feature selection: basic on/off feature selection and feature weighting. We verified that single-value features (referred to here as one-dimensional features, such as “velocity”) work efficiently for a preliminary classification approach. They allow performing quick training on a limited subset of data. However, in order to improve the classification, polyphonic features consisting of an array of values (referred to here as multi-dimensional or multi-value features, such as pitch histograms, patterns of notes, or patterns of chords) must be used. We remark that most of these MIDI features have no equivalent attribute in the “traditional” seismic domain. They can effectively provide a complementary contribution to classify seismic facies.

5.4. Validation Tests and Classification

We classified the whole seismic data using both single-value and multi-valued features. We used a combination of classifiers: K-NN for single-value features and ANN for multi-valued features.

Figure 5 shows some indicative results of the validation test. The arrows indicate the position of some areas where the channels are calibrated by wells. The numbers indicate the correspondent classification scores. In this particular example, the scores represent the percentage of recordings that were assigned the correct category in the areas of the wells.

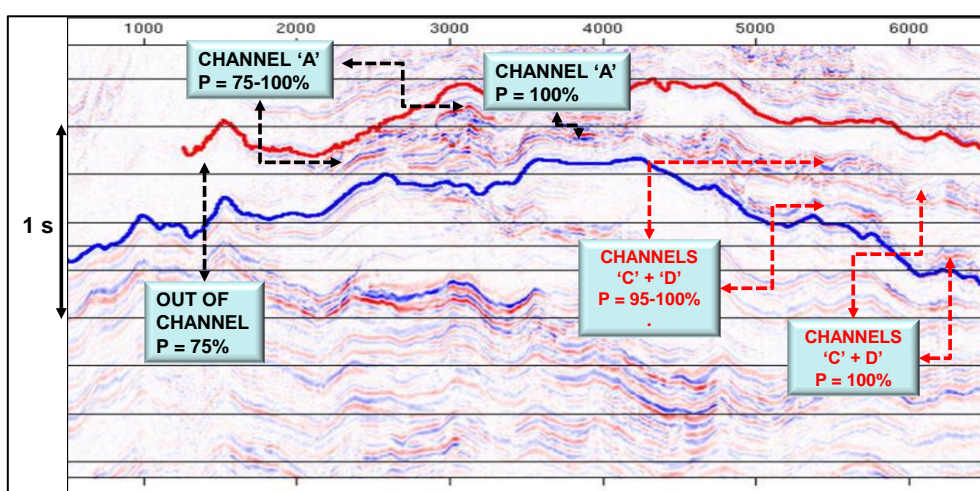


Figure 5. Validation tests using a hybrid K-NN and ANN approach. The arrows indicate approximately the locations of the traces in the areas of the wells used for the verification tests. Numbers indicate the percentage of correct classification in the areas of the wells. Horizontal axis: trace number. Inter trace distance is 12.5m. Vertical axis: two-way times (s).

Figure 5 shows just locally the scores of the validation tests. Instead, Figure 6 provides an extended view of the classification results co-rendered with the seismic section. The different channels are clearly detected in separate classes. In the Figure, the classification scores trace by trace are represented in colors and superimposed to the seismic section. Yellow area indicates that the channel “A” is classified with a variable score ranging from 75 to 100: orange indicates a less ambiguous area, where the same channel “A” is classified with higher reliability (given by a classification score close or equal to 100). The green and violet areas indicate, respectively, two different channels, “C” and “D”, classified with scores of almost 100.

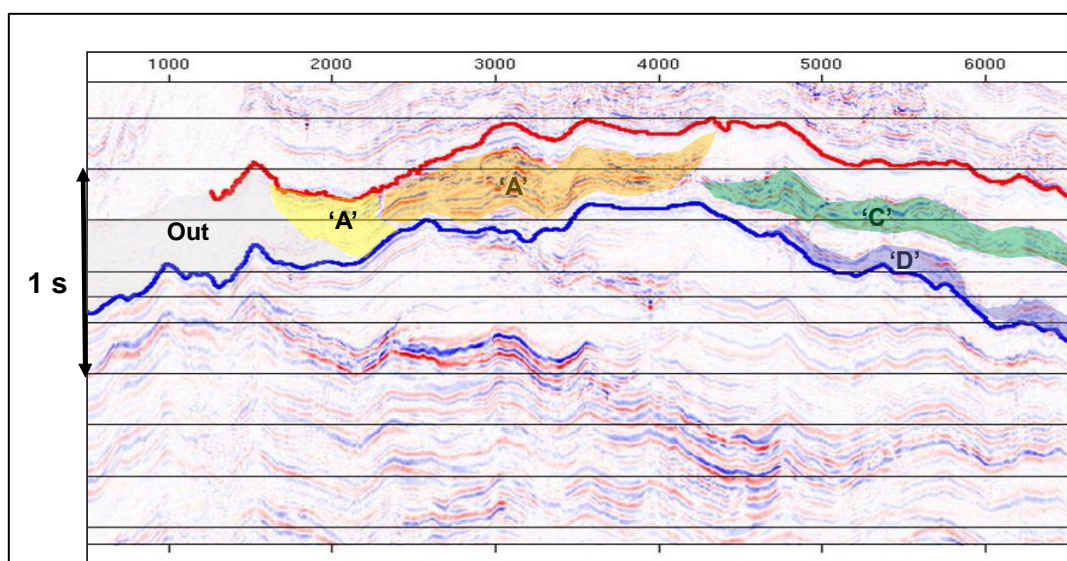


Figure 6: Overall classification results co-rendered with the seismic section. Yellow: classification score for channel “A” ranging from 75 to 100. Orange: classification score for channel “A” close or equal to 100. Green and violet areas indicate, respectively, the channels “C” and “D” classified with scores of almost 100. Grey indicates that no MIDI segment has been clustered in anyone of the channel classes. Horizontal axis: trace number. Inter trace distance is 12.5m. Vertical axis: two-way times (s).

5.5. Test Discussion

The demonstrative example discussed in this paper is aimed at providing a proof of concept of our approach. Thus, we used a good quality data set consisting of a limited number of seismic traces (several thousands). At this initial stage of our research, our first goal is to verify that our MIDI-based approach is able to classify seismic data at least with comparable performance with respect to other consolidated approaches. First, we compared our classification results vs. analogous results obtained with “standard” seismic attributes, such as those related to various functions of frequency and amplitude. Our main conclusion is that when also “high-level” MIDI features are included (such as rhythmic, melodic and harmonic multi-valued features), the classification performance improves in terms of cluster separation. Moreover, the MIDI-based approach is very fast, due to the simple structure of MIDI files. Indeed, multiple MIDI messages related to pitch, amplitude, note length, rhythmic, melodic and harmonic patterns extracted from a seismic trace can be codified in a file of few

Kbytes for each trace. When classification must be performed on big data sets, our approach can be advantageous because of the speed with which MIDI traces can be processed. Finally, several MIDI features have no equivalent seismic attributes. These additional MIDI features can be used jointly with the “conventional” seismic attributes, improving the performance of the classification process. In summary, we think that our method can be complementary with respect to other classification approaches based on well consolidated seismic features.

6. Final Remarks

Digital music technology can be effectively used for analyzing and interpreting geophysical data. A double workflow can be applied: it consists of interactive audio-video display on selected subsets of the data, combined with automatic pattern recognition and classification approaches. Previous works demonstrated that double sensory analysis of seismic data can effectively improve the perception of anomalies and patterns related to hydrocarbon filled layers. In this paper, we show that efficient techniques of pattern recognition and clustering commonly applied in the domain of musical information retrieval (MIR) can be applied to geophysical data. These must be accurately converted into MIDI files (or into other symbolic formats). The fact that this type of format does not require large memory resources, makes this automatic approach fast and accurate at the same time.

We performed a proof of concept test on real data using a MIR approach to cluster and classify automatically a seismic data set. Our test shows that gas-bearing layers can be detected and properly classified using a hybrid approach combining k-Nearest Neighbors and Automatic Neural Network supervised algorithms. Many MIDI features have been extracted from the data, including single-valued and polyphonic features. Some of these do not have any correspondence with the standard seismic attributes. They represent a new class of diagnostic features that can be used as complementary classification attributes.

Conflict of Interest

All authors declare no conflict of interest in this paper.

Note

MIDI: short for Musical Instrument Digital Interface. It is a technical standard that describes a protocol, digital interface and connectors for linking and connecting a wide variety of electronic musical instruments, computers and other related devices.

References

1. Theodoridis and Koutroumbas (1998) Pattern Recognition, Academic Press, London.
2. Zhao T, Jayaram V, Roy A, et al (2015) A comparison of classification techniques for seismic facies recognition. *Interpret* 3: 29-58.
3. Duda RO, Hart PE, Stork DG (2001) Pattern classification. 2nd ed: John Wiley & Sons.
4. Dell'Aversana P (2013) Listening to geophysics: Audio processing tools for geophysical data analysis and interpretation. *The Leading Edge* 32: 980-987.

5. Dell'Aversana P (2014) A bridge between geophysics and digital music. Applications to hydrocarbon exploration. *First Break* 32: 51-56.
6. Dell'Aversana P, Gabbriellini G, Amendola A (2016) Sonification of geophysical data through time-frequency transforms. *Geophys Prospect*.
7. Stockwell RG, Mansinha L, Lowe RP (1996) Localization of the complex spectrum: the S Transform. *IEEE Trans Signal Process* 44: 998-1001.
8. Wang A (2003) An industrial-strength Audio Search Algorithm, ISMIR, London: Shazam Entertainment Ltd.
9. Kiang MY (2003) A comparative assessment of classification methods. *Decis Support Syst* 35: 441-454.
10. McKay C (2004) Automatic Genre Classification of MIDI Recordings, Music Technology Area Department of Theory, Faculty of Music, McGill University, Montreal.
11. Aminzadeh F and de Groot P (2006) Neural Networks and Other Soft Computing Techniques with Applications in the Oil Industry, EAGE Publications.



AIMS Press

© 2016 Paolo Dell'Aversana, et al., licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0>)