AUTOMATIC RECOGNITION OF LOW FREQUENCY RADIO PLANETARY SIGNALS

H. de Lassus* and A. Lecacheux*

Abstract

We explore the problem of automatic detection from spectrograms of mixed low frequency radio astronomy signals with additive plasma noise. We describe and compare four different approaches: Higher Order Statistics clustering by cumulants matching; Wavelet basis nearest neighbor filtering and quantification; Regular grammar of differential operators; Cellular automata grid nested with neural networks. Using real data from Ulysses URAP experiment, we show what can be expected from these various techniques, and give evidence that real time, on board automatic monitoring and tuning of spectrometers is at hand.

1 Introduction

Classification of low frequency radio astronomical signals on the time frequency plane is a task similar to the well known "cocktail party" problem: identify many sources emitting at the same time. Fortunately, the most powerful sources are the planets and the Sun, and the problem can be reduced to the solar system with only a dozen of main sources. The signal under study originates from different parts of the solar system, and the distances, as well as the observation angles, are heterogeneous and moving fast. This leads to unstable patterns on the spectrogram (Figure 3). Monitoring these wide band radio emissions, may give the clue to understand underlying physical phenomena driven by magnetohydrodynamic space plasma laws. A growing need is foreseeable in the near future for artificial, expert readers of spectrograms of low frequency planetary radio emissions. The growth of data output from experiments is reaching such a level that it is already almost impossible for a human expert to extract all the scientific information available from the data. Usually the recognition is done on dynamic spectra. The task of an automatic spectrogram expert reader can be divided in subtasks: detection (separate signals from noise and interferences), estimate the number of current radio sources (component separation), classify the signals (label the signals), finally, evaluate the false alarm risk, accept or reject the classification. We describe and compare four different approaches that

*Groupe de Recherche Instrumental en Radioastronomie, CNRS URA 1757, Laboratoire ARPEGES, Observatoire de Meudon, 5 place Jules Janssen, 92195 Meudon Cedex, FRANCE

realize some of these tasks and enable spectrogram recognition. Higher order statistics clustering by cumulants matching, wavelet basis nearest neighbor filtering and quantification, regular grammar of differential operators, cellular automata grid nested with neural networks. Using real data from Ulysses URAP experiment we show that these techniques are complementary and may be used to design a comprehensive recognition device.

2 The data set

For our comparison we used three data set: one for learning (47000 spectra from Ulysses/URAP of year 1991), one for testing (13000 spectra of year 1991), one to study the stability of the classification process near a major trajectory change of the spacecraft (23400 spectra of year 1992). Each spectrum is sampled over 16 logarithmic spaced radio frequencies. The sampling rate is one spectrum per 144 seconds. Small gaps are filled by bilinear interpolation. Constant background noise is subtracted (i.e. receiver and local plasma noise contributions) and all data below an adaptive detection threshold are neglected. Then normalization is done between $(-1, +1)$ with a 24 hours moving window applied on the data frequency by frequency.

3 Higher order statistics clustering by cumulants matching

The first method that we will mention is particularly concerned with source separation.

3.1 Find templates in the cumulants domain with a learning data base

Following works by Brillinger [Brillinger, 1981], Giannakis proposed to use cumulants matching for signal detection and classification [Giannakis et al., 1990; Giannakis et al., 1992]. Benefiting from the fact that cumulants are insensitive to second order noise, the method suggests to calculate Hamming distances between higher order statistics of the signal and a template using a Viterbi algorithm. We implemented this technique on Ulysses URAP signal using a learning data base and a test data base. On the learning data base we selected samples of each type of emissions and we computed cumulants of order 3 and 4 of signal vectors. Then, we clustered our different types of emissions in subclasses using a clustering method on the cumulants domain. (We tried LBG vector quantification [Gersho et al., 1992] and ART2 on supervised neural networks [Carpenter et al., 1987]). This leads us to a collection of templates of each signal in the cumulants domain. Then, on a test data base, for each sample we computed the Hamming distance to each sub class template we had calculated from the learning data base. We classified the signal by nearest neighbor method.

3.2 Results with four classes on 1,250000 samples

We found that cumulants matching enables source counting but not signal recognition, as the true class could be computed in the mean but not punctually. When five classes

were computed out of four signals, the fifth class embodied all mixed signals. So the problem arose of determining the true number of classes. We introduced a penalty term (Rissannen minimum description length [Rissanen, 1989]) to trade off minimum distance clustering and complexity. This method together with polarization measurements and beam forming is a reasonable way for source counting.

4 Wavelet basis nearest neighbor filtering and quantification

The second method that we explored was wavelet analysis. The idea is similar to the HOS technique except the fact that the matching is done in the wavelet domain. For this technique we used the same learning data base and test data base as for the HOS method. On each signal sample of the learning data base, we computed the wavelet transform module [Daubechies, 1990]. We then normalized each scale. We found by vector quantization N sub classes (LBG or ART2), and we computed the envelopes of the N sub classes which yielded our templates. With a test data base, for each sample we computed the distance to each prototype envelope. and we classified the signal by nearest neighbor method. The results we show were computed from 1,280000 test signal samples. which included 613 solar type III, 88 Jupiter nKOM, 535 Jupiter bKOM, and 45 Jupiter HOM events. A great number of subclasses are needed (10) to get a good result; this is understandable as wavelet filtering is not invariant by translation. Thus, if the samples are in great number and if the signal is broadband, (Jupiter bKOM and solar type III) results of Table 1 show that recognition is effective, and moreover, interference excision should be possible [Maslakovic et al., 1996].

Sub-classes	Recog.	Not recog.	F. alarm
1	20%	31%	49%
2	38%	25%	37%
3	49%	20%	31%
4	61%	12%	27%
10	80%	8%	12%

Table 1: Results of the wavelet domain filtering according to subclasses number.

5 Regular grammar of differential operators

This third method has been designed for solar type III automatic recognition. We define a partial derivative operator of the fifth order in time and of the first order in frequency [de Lassus, 1993]. We convolve the spectrogram with this operator, and apply an adaptive threshold. The resulting image is parsed with a regular grammar automaton to extract the signal. This automaton contains grammar rules fixed according to physical properties of the phenomena. The results we show in Table 2 are computed from a validation data

Day	Recog.	Not recog.	F. alarm
91-09-24	4	0	0
91-09-25	9	1	0
91-09-26	10	0	0
91-11-07	6	0	0
91-11-08	1	1	0
91-11-09	12	0	0
91-11-10	9	1	0
91-12-02	10	2	0
91-12-03	4	0	0
91-12-14	6	1	0
92-02-06	1	0	0

Table 2: Number of type III recognized with differential operator plus grammar automaton.

base of 11 days of Ulysses URAP experiment in autumn 1991. 72 out of 78 solar type III are recognized (92.31%). Only weak emission or mixed signal emission are not recognized. This method yielded no false alarm.

6 Cellular automata grids nested with neural networks

The fourth method appears to be the most promising as it is suited for all types of signals. It uses a particular type of neural network: the Time Delay Neural Network [Lang et al., 1990] and a virtual machine, the inhomogeneous cellular automaton [Aleksic, 1993].

6.1 Inhomogeneous cellular automata

Inhomogeneous cellular automata (CA) are dynamic systems where time and space are discrete. The state of cells on a regular grid is simultaneously updated according to a local interaction rule. Each cell obeys to its own rule, and can reach a finite number of states. At each clock tick all cells calculate their future state with the transition rule watching the state of neighboring cells. With proper rules the system converges to an equilibrium that is the solution to the problem to solve.

Typical CA tasks are to estimate probability densities from a cluster of neural networks, recover signal shape from an energy peak, and estimate cyclo-stationarities from signals. They can also filter neural network outputs, according to physical laws, filter RFIs, or implement any human expertise in a decision process. Figure 1 through figure 2 show a CA classifier converging to the color which has the greatest density in the original image (white).

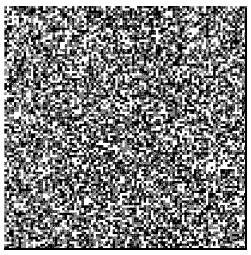
(a)

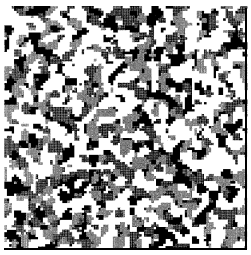
(b)

Figure 1: 100×100 *4 classes cellular automata classifier (a) time 1, (b) time 2.*

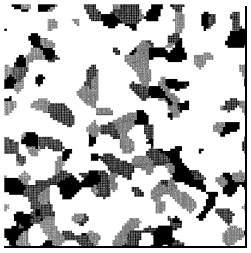
(a)

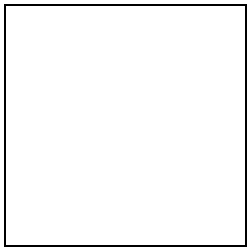
(b)

Figure 2: 100×100 *4 classes cellular automata classifier (a) time 8, (b) time 12.*

6.2 The Time Delay Neural Network

Conventional classifiers estimate the likelihood probabilities by assuming they can be well modeled by specific parametric distributions such as Gaussian mixture distributions. Training involves estimating the parameters of the likelihood distributions and estimating the a priori class probabilities from training data. In contrast, neural networks do not estimate Bayesian probabilities in this indirect way. Instead, when the desired outputs are 1 of K and an appropriate cost function is used, Bayesian probabilities are estimated directly. Neural network outputs can then be used as Bayesian probabilities if three conditions hold: sufficient training data are available, the network is complex enough, and the classes are sampled with the correct a priori class probabilities [Richard et al., 1991]. The difficulty with usual static neural networks such as multi layer perceptrons, is that time frequency plane is not taken into account specifically. This is why Lang & al. introduced the Time Delay Neural Network (TDNN) [Lang et al., 1990] as a special multi layer perceptron for the dynamic approximation of speech contrast functions on spectrograms. The architecture of the TDNN tries to minimize the parameter number of the contrast function. Using a "shared weight" principle, it slides a small perceptron along a windowed spectrogram, transmitting a sum of local decisions to upper layers of the network. Results by Miniere [Miniere, 1994] indicate that this method may be the right choice for time frequency detection of transient signals.

6.3 Learning from the data

The different data sets from Ulysses/URAP are manually labeled sample by sample and frequency by frequency. A single spectrum may include up to four different classes of signals (TIII, bKOM, nKOM, HOM). Two human experts are needed to confront uncertain recognition, and human errors are commonly rated at $5-10\%$, while 10% of the data are

impossible to classify and 20% doubtful. This leads to 60% of clearly identified signals. The distribution of useful events is 40% TIII, 45% bKOM, 10% HOM,5% nKOM. The technique chosen implies that only local maxima of energy are selected for the first classifier. Each frequency is swept and when a local maximum of energy is found, a window of 10×6 samples, chosen with the zero crossings of the average autocorrelation function of the signal, is taken centered around the maximum. This matrix of 60 samples is the input of a Time Delay Neural Network detector. As the statistical distribution of the samples among classes is uneven, we chose to incorporate neighboring samples around the maxima for those classes that were under–represented.

6.3.1 Overlapping windows

We could not broaden too much the size of our observation window in order to keep a reasonable number of weights in the Time Delay Neural Network. It is well established that a minimum of 10 independent learning examples per weight are required. This constraint was such that many events could not fit in the window 10×6. This problem has a solution with overlapping time frequency windows. Incorporating an overlapping window scheme of three Time Delay Neural Networks for each frequency detector, we succeeded in monitoring large events, as well as small ones. We decided to separate the signals in three classes on a first neural step, leaving the separation between bKOM and HOM classes to a subsequent cellular automaton device. The statistical properties of these two signals are very similar, and to separate them we have to apply physical observation laws, such as time occurrence or observing geometry, leading to pre-determined occurrence rules in the time frequency plane. This could not be implemented easily on the static window scheme we chose for the Time Delay Neural Network, whereas recursive neural nets could not be implemented successfully because of telemetry gaps. Cellular automata seemed the appropriate tool for this task.

6.3.2 The neural network pre–classifies the signal in three classes

The architecture of the classifier is composed of one neural system for each frequency channel. A system is composed of three Time Delay Neural Networks, each watching different windows, but all windows include the frequency where the local maximum lies. The Time Delay Neural Networks have four layers of respectively 60, 6, 5, 3 nodes. The "position independent feature extractor" [Lang et al., 1990] of the neural network has a width of three samples. At the top of this construction, a multi-layer perceptron receives the 3×3 outputs of the Time Delay Neural Networks as its inputs. This perceptron has respectively 9, 5, 4, 3 nodes on each of its four layers. The complete neural classifier includes 27 Time Delay Neural Networks and 7 multi-layer perceptrons. This architecture has been selected after intensive trials, concerning information merging at the output of adjacent detectors. Learning is performed with the 60 samples windows centered on an energy peak with back-propagation gradient descent and the labels. The data presentation order is shuffled between each iteration. The Time Delay Neural Networks perform their learning independently, then their weights are frozen and the multi-layer perceptron is

added. Learning is continued with frozen weights until generalization is degrading on the test set. Then learning is resumed for the complete system for a few iterations. The neural classifier yields an output mapping the input space $(3 \times 6 \times 10)$ $(-1; +1)$ in a Bayesian probability approximation contrast space $(3 \times 1)(0; +1)$ which gives only a skeleton of the different signals. This output has now to be processed by the cellular automata device with two goals: label the bKOM and HOM classes, and recover the signal around the skeleton for all classes.

6.4 Cellular automata detection

We propose to apply the cellular automaton paradigm to the spectrogram segmentation, incorporating in the transition rules physical observation laws. In our experiment we chose the 9×2 cells neighborhood with five states (Jovian bKOM and HOM, solar type III, Jovian nKOM, and noise) on two regular grids of 72000 cells (48 hours moving grid). Each frequency channel has its own local rules which enable: energy gradient estimation, the diffusion of labels along iso-energy lines, labels density estimation, lagged correlations of labels, and winner takes all decision. During initialization, only the highest frequencies are labeled HOM and the lowest bKOM. All other frequencies have an undetermined bKOM-HOM label. Then the system is let to evolve freely. After a bounded number of iterations depending upon the rules implemented, (typically 15 iterations) the system reaches stability. bKOM and HOM are labeled, other signals have recovered their size around the skeleton given by the neural networks.

7 Results and discussion

Examining the results on the test set (Table 3), we saw that all errors were concentrated on ambiguous samples. 10% of errors were due to an obvious human error in labeling. 80% were doubtful samples. As human experts had balanced between two labels, if the network took the second choice it could not be considered false. Real errors (10%) occurred with interferences on the channel, proximity of a telemetry gap, or anomalous signals looking much like an other class. nKOMs suffer lesser results, because their number is dramatically low in the learning set. We have tried an independent test to evaluate who was right from the human expert and the network in doubtful cases. A measure of the polarization of the received radio signal is available in URAP data from time to time. It enables the discrimination between unpolarized solar signals and highly polarized Jovian signals taken as a whole. This test shows that the neural network took the right choice for the doubtful samples much more often than the human expert (90% of the time). So we reassessed the learning data base, modifying detected human labeling errors, and we settled many doubtful cases with this polarization test. After a second learning, results were above 90% of success. Table 3 shows that recognition rate depends upon the frequency channel. There is no recognition for channels 1 to 4, as there was no scientific interest for this. Cellular automata segmentation gives 99% of good labeling between bKOM and HOM on signals originally labeled as bKOM-HOM by the networks.

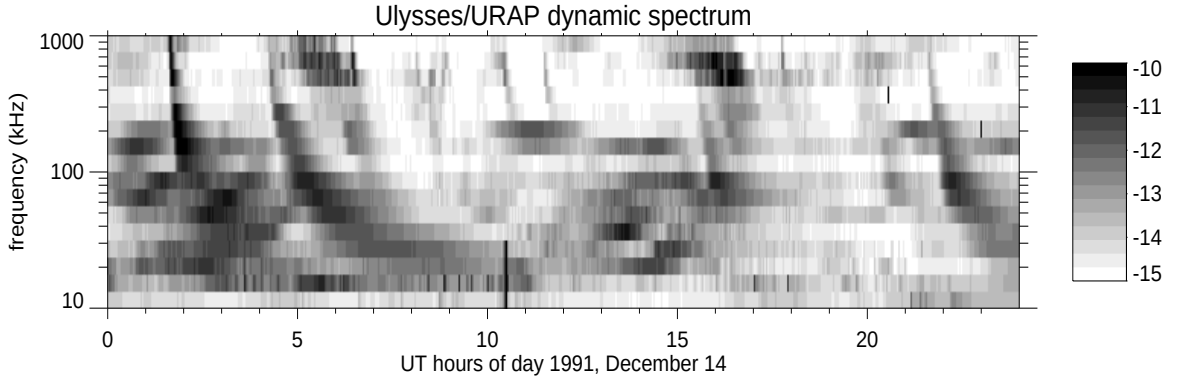

Figure 3: Dynamic spectrum from Ulysses URAP 12-14-91.

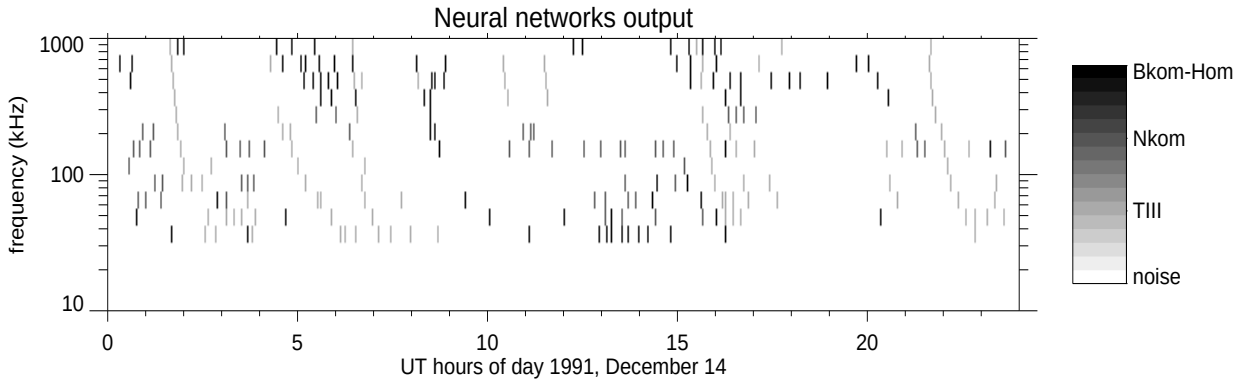

Figure 4: The output of the Time Delay Neural Networks.

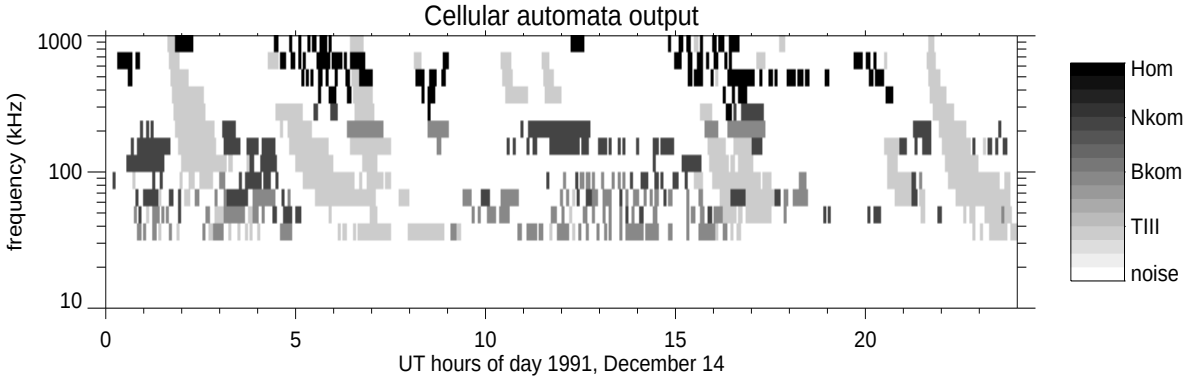

Figure 5: The output of the Cellular Automata after neural network pre-treatment.

channel	freq.	Type III	bKOM	nKOM
16	940 kHz	95%	98%	
15	740 kHz	96%	96%	
14	540 kHz	94%	94%	
13	387 kHz	92%	96%	96%
12	272 kHz	91%	92%	95%
11	196 kHz	94%	95%	91%
10	148 kHz	97%	98%	87%
9 (RFI)	120 kHz	78%	81%	65%
8	100 kHz	93%	97%	80%
7	81 kHz	89%	84%	73%
6	63 kHz	86%	94%	
5	52 kHz	93%	94%	

Table 3: Test set recognition rates of neural networks plus cellular automata.

The recovering of signals around the skeleton given by the neural networks for all classes is almost perfect. In the extrapolation test from January 1, 1992, to Jupiter swing by on February 8, 1992, results keep the same high standard of quality.

The Ulysses/URAP dynamic spectrum, as shown in Figure 3, is treated with the Time Delay Neural Networks (Figure 4), the output of the Cellular Automata after neural network pre–treatment is demonstrated in Figure 5.

7.1 Cellular automata grids nested with neural networks: an assessment

Combined together, the neural network approach and the cellular automata paradigm is a nice method. Cellular automata separate signals with same texture, they incorporate easily physical laws (e.g. cyclo-stationarities) in the recognition process. Neural networks separate signals with different texture [de Lassus et al., 1996a], and are able to classify

patterns beyond human expertise (e.g. low energy overlapping signals), they can detect abrupt changes and reduce false alarms [de Lassus et al., 1996b].

8 Conclusion

We have compared four different methods to recognize low frequency planetary radio emissions. Differential operators work well on signal with a "shape" like solar type III, they suffer no false alarm. Cumulants clustering combined with direction of arrival techniques allows source counting. Wavelets do not see well narrow band signals, thus, they may be efficient at filtering narrow band interferences. Neural network is the right solution for recognition and change detection, while cellular automata give the glue to integrate real time recognition systems. Combining these various techiques, enables radio planetary low frequency signal recognition and tracking. Typical applications could be, real time recognition on board a spacecraft, or high resolution spectrogram scanning. (1024 or 2048 points FFTs over 20MHz [Kleewein et.al.,this volume]). New challenges are now ahead: massive data mining for comparative astronomy, real time spectrometer tuning, detection of transients, excision of RFIs.

Acknowledgments: We acknowledge fruitful collaboration from S. Thiria and F. Badran from Université Paris VI and Conservatoire National des Arts et Métiers Paris. This work has been partly supported by the French Space Agency (CNES), the French Programme National de Planétologie (CNRS-INSU) and Paris-Meudon Observatory.