Sparse coding for apnea-hipopnea detection

F. E. Lestussi[a] H. L. Rufiner [abc] L. E. Di Persia[ab] D. H. Milone[ab]

[a] *Centro de I+D en Señales, Sistemas e INteligencia Computacional (SINC(i)), Fac. de Ingeniería y Cs. Hídricas, Univ. Nacional del Litoral, Santa Fe (Argentina).*
[b] *Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET), Argentina.*
[c] *Laboratorio de Cibernética, Fac. de Ingeniería, Univ. Nacional de Entre Ríos, Entre Ríos (Argentina).*
Email: {flestussi, lrufiner, ldipersia, dmilone}@fich.unl.edu.ar

Abstract— **The sleep apnea-hipopnea syndrome is a common disorder which could be detected by the decrease of the oxygen saturation. In this work we propose the use of a sparse representation of the SaO_2 signal to help in the detection of this pathology. For train and test we used the SHHS Polysomnography Database which include 1000 polysomnograms. We show that the reconstruction of the SaO_2 signal from an optimal estimated dictionary has errors in the order 10^{-13}. The activation coefficients obtained from this dictionary was used to non-linearly regress the apnea-hipopnea marks by training a multilayered perceptron. Then we estimate the apnea-hipopnea marks from a test dataset with very good results, with a correlation coefficient of 81.74%.**

Keywords— **Sparse representation, Dictionary, Neural network, Sleep apnea-hipopnea syndrome.**

1 INTRODUCTION

The American Academy of Sleep Medicine distinguishes more than 80 different sleep disorders [2]. Problems with falling asleep or daytime sleepiness affect approximately 35 to 40% of the U.S. adult population annually and are a significant cause of morbidity and mortality [5].

One of the pathologies is the obstructive sleep apnea-hipopnea (OSAH) syndrome. The syndrome is characterized by repetitive episodes of airway narrowing or collapse during sleep. The Argentine consensus of respiratory disorder related to sleep shows that the sleep apnea-hipopnea syndrome has a prevalence of 2% to 4% in adults and increase in the elderly [3].

The gold standard diagnostic test for sleep apnea is an overnight polysomnography in a sleep laboratory which is costly in terms of time and money, and the accessibility in some areas is limited. The pulse oximetry is attractive, for the measure of OSAH, because of its availability and ease of application. A decrease of oxygen saturation (SaO_2) detected by pulse oximetry could indicate respiratory events.

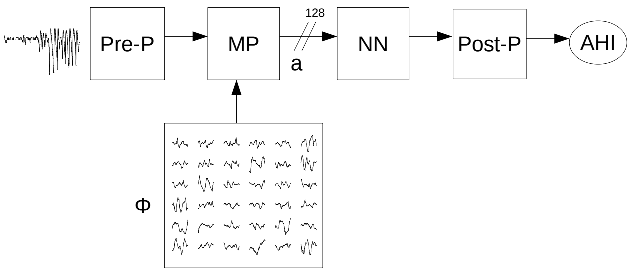

Figure 1: Diagram of the proposed method for apnea-hipopnea detection.

In the last years, several researchers have taken a different approaches to traditional signal processing. These new formulations give rise to techniques based on non-linear systems and higher-order statistics, including independent component analysis (ICA) and methods to obtain sparse representations (SR) of a signal. They provide new ways of phrasing the solution of the problem of signal modeling or representation. One underlying idea is that of representing the signals involved using only a few significant characteristics; that is, as a SR of only a few basic waveforms. There are applications of these techniques to different fields, such as: natural image analysis [16, 15], audio and music signals [17], general biomedical signals [6, 12, 13] and automatic speech recognition [8, 7].

The aim of this work is to estimate a dictionary that produce a SR of the SaO_2 signal, i.e. the coefficients of the representation have few active atoms that preserve the most relevant information. After that, the activation coefficients are used as input of a multilayered perceptron in order to regress the respiratory events from the same dataset used to learn the dictionary. In Fig. 1 we show a block diagram of the proposed system to evaluate the apneas-hipopneas detection. Here, Pre-P is a pre-processing algorithm to denoise and windowing the signal, MP is the maching pursuit algorithm which estimates the activation coefficients of the dictionary of each frame signal, the neural network (NN) previously trained try to estimate the apnea-hipopnea marks by the activation coefficients, and the last block is a post proccesing (Post-P) algorithm which eliminates spureous detections.

In the next section we describe the methods to obtain the dictionary and the coefficients of a SR of signals. After that we present preliminary results, where we show the sparsity of the dictionary and classification results.

2 METHODS

2. 1 Sparse representation of signals

Given a signal $\mathbf{s} \in \Re^N$, we consider a representation in terms of a dictionary $\mathbf{\Phi}$ as a decomposition of the form:

$$\mathbf{s} = \sum_{i=1}^{M} \phi_j a_j = \mathbf{\Phi a} \qquad (1)$$

where $\mathbf{a} \in \Re^M$ is the coefficients vector, $\mathbf{\Phi} \in \Re^{N\times M}$ with $M \geq N$, is a collection of waveforms or atoms ϕ_j and both ($\mathbf{a}$ and $\mathbf{\Phi}$) usually unknown.

Some authors use the term "basis" instead of "dictionary"; however, as the set of atoms may not be linearly independent and usually more atoms than the space dimension are used, the last one is preferred.

When there are more waveforms in the dictionary than samples $\mathbf{s}$, i.e., $M > N$ (referred to as an overcomplete dictionary), or when the waveforms do not form a basis, then there will be non-unique representations of the signal. In this situation a suitable criterion is required to select only one of them. In this context, sparseness often refers to the criterion of choosing a representation with "as few non-zero coefficients as possible" (typically using the l_0 norm), although several other criteria have been introduced [4].

The problem of a SR of $\mathbf{s}$ with respect to l_0 could be stated as follows:

$$\min ||\mathbf{a}||_0 \text{ subject to } \mathbf{\Phi a} = \mathbf{s} \qquad (2)$$

It is important to note that in the overcomplete case mentioned above, although Eq. 1 is linear, the coefficients a_γ chosen as the solution correspond to a non-linear function of the data $\mathbf{s} : \mathbf{s} \rightarrow \{a_\gamma\}$.

In order to solve the SR problem, it can be split into two sub-problems: inference and learning. The first one consists in finding the representation coefficients $\mathbf{a}$ which satisfy a given sparsity criterion. The second one involves finding the optimal dictionary $\mathbf{\Phi}$ to represent the data. The last one is usually the most complex of both.

2. 2 Inference of the activation coefficient

A more general framework is assumed, where Eq. 1 is rewritten to include an additive Gaussian noise term ε as follows:

$$\mathbf{s} = \mathbf{\Phi a} + \varepsilon \qquad (3)$$

Following terminology used in ICA, Eq. 3 is referred to as the generative model, to signify that one generates the signal $\mathbf{s} \in \Re^N$ from a set of hidden sources a_j, arranged as a state vector $\mathbf{a}^M$, using a mixing matrix or dictionary $\mathbf{\Phi}$ of size $N \times M$, with $M \geq N$.

The a_j are initially assumed to be statistically independent, with a joint a priori distribution:

$$P(\mathbf{a}) = \prod_{j=1}^{M} P(a_j) \qquad (4)$$

If $\mathbf{\Phi}$ is known and $\mathbf{s}$ is given, the state vector $\mathbf{a}$ can be estimated via the Bayes's rule by considering the posterior distribution:

$$P(\mathbf{a}|\mathbf{\Phi}, \mathbf{s}) = \frac{P(\mathbf{s}|\mathbf{\Phi}, \mathbf{a})P(\mathbf{a})}{P(\mathbf{s}|\mathbf{\Phi})} \qquad (5)$$

A maximum a posterior estimation of $\mathbf{a}$ reads as follows:

$$\hat{\mathbf{a}} = \arg_{\mathbf{a}} \max[\log(P(\mathbf{s}|\mathbf{\Phi}, \mathbf{a})) + \log(P(\mathbf{a}))] \qquad (6)$$

If the posterior is sufficiently smooth, the maximum can be found applying gradient ascent. The solution depends on the form of the distribution chosen for the noise term ε and sources a_j, giving rise to different methods for finding the coefficients. Lewicki and Olshausen [10] proposed an a priori distribution of Laplacian type:

$$P(a_j) = N e^{\rho_j |a_j|} \qquad (7)$$

where ρ_j is given and, if the noise is Gaussian, this leads to the following rule for updating $\mathbf{a}$:

$$\Delta \mathbf{a} = \mathbf{\Phi}^{\mathbf{T}} \Lambda_\varepsilon \varepsilon - \rho^{\mathbf{T}} |\mathbf{a}| \qquad (8)$$

where Λ_ε is the inverse of the noise covariance matrix $E\{\varepsilon^{\mathbf{T}} \varepsilon\}$, with $E\{.\}$ denoting the expected value, and $\rho = \{\rho_j\}$.

2. 3 Learning of the dictionary

Until now, a statistical method has been used to solve the inference problem. In what follows the learning problem is similarly solved. To estimate the atoms of $\mathbf{\Phi}$, the following objective function can be maximized [10]:

$$\hat{\mathbf{\Phi}} = \arg_{\mathbf{\Phi}} \max[L(\mathbf{s}, \mathbf{\Phi})] \qquad (9)$$

where $L = E[\log P(\mathbf{s}, \mathbf{\Phi})]_{P(\mathbf{s})}$ is the likelihood of the data. This likelihood can be found by marginalizing the following product of the conditional distribution of the data, given the dictionary, and the coefficients, together with the coefficients a priori distribution:

$$P(\mathbf{s}|\mathbf{\Phi}) = \int_{\Re^M} P(\mathbf{s}|\mathbf{\Phi}, \mathbf{a}) P(\mathbf{a}) d\mathbf{a} \qquad (10)$$

where the integral is over the M-dimensional state space of $\mathbf{a}$.

The objective function Eq. 9 can be maximized using gradient ascent with the following update rule for the matrix $\mathbf{\Phi}$ [1]:

$$\Delta\mathbf{\Phi} = \eta\Lambda_\varepsilon E[\varepsilon\mathbf{a}^\mathbf{T}]_{P(\mathbf{a}|\mathbf{\Phi},\mathbf{s})} \quad (11)$$

where η is a learning coefficient (between 0 and 1). The problem at this point is how to calculate this update rule, given that it involves solving the following integral:

$$E[\varepsilon\mathbf{a}^\mathbf{T}]_{P(\mathbf{a}|\mathbf{\Phi},\mathbf{s})} = \int_{\Re^M} (\mathbf{s} - \mathbf{\Phi}\mathbf{a})\mathbf{a}^\mathbf{T} P(\mathbf{a}|\mathbf{\Phi},\mathbf{s})d\mathbf{a} \quad (12)$$

As the dimension of $\mathbf{a}$ increases, the previous integral becomes analytically intractable and different authors have proposed approximation methods in order to compute it. Lewicki and Sejnowski [11] used a multivariate Gaussian approximation to the posterior distribution around its maximum $\hat{\mathbf{a}}$:

$$P(\mathbf{a}|\mathbf{\Phi},\mathbf{s}) \approx \sqrt{\frac{|\mathbf{H}|}{2\pi^M}} e^{-1/2(\mathbf{a}-\hat{\mathbf{a}})^\mathbf{T}\mathbf{H}(\mathbf{a}-\hat{\mathbf{a}})} \quad (13)$$

where $\hat{\mathbf{a}}$ is the mean value, and $\mathbf{H}^{-1}$ is the covariance, being $\mathbf{H}$ the Hessian of the log-posterior evaluated in $\hat{\mathbf{a}}$:

$$\mathbf{H} = -\nabla\nabla^T \log P(\mathbf{a}|\mathbf{\Phi},\mathbf{s}) \quad (14)$$

It provides a good approximation for $\mathbf{a}$ close to $\hat{\mathbf{a}}$. This result provides a solution of Eq. 11 given by

$$\Delta\mathbf{\Phi} = \eta\Lambda_\varepsilon(\hat{\varepsilon}\hat{\mathbf{a}} - \mathbf{\Phi}\mathbf{H}^{-1}) \quad (15)$$

where $\varepsilon = \mathbf{s} - \mathbf{\Phi}\hat{\mathbf{a}}$

In order to obtain the dictionary and the coefficients (Eqs. (8) and (15)), in this paper we use the implementation proposed by Lewicki and Olshausen [10] (using Lewicki's noise overcomplete ICA code or NOCICA), at the training of the dictionary.

2.4 Maching pursuit

The solution of the problem given by Eq. 2 is in terms of l_0 norm, which turns it computationally intractable. Eq. 8 solves a near problem that generally shares the same global minimun. This solution is exact in terms of l_1 norm but is computationally very expensive. For the dictionary training stage, we need the best inference of $\mathbf{a}$ to obtain an optimal estimation, so we use that last one. But once the dictionary was estimated, we could even approximate the inference solution to obtain a faster algorithm.

Mallat and Zhang [14] proposed a general method to approximate the solution of Eq. 2. Sparsity is directly included by choosing an appropriate number of terms. Given an initial approximation $\mathbf{s}^{(0)} = \mathbf{0}$ and an initial residual $\mathbf{R}^0 = \mathbf{s}$, a sequence of approximations is iteratively constructed. At step k the parameter $\gamma = \hat{\gamma}$ is selected, such that the atom $\phi_{\hat{\gamma}}^{(k)}$ best correlates with the residual $\mathbf{R}^k$, and a scalar multiple of this atom is added to the approximation at step $k-1$, obtaining:

$$\mathbf{s}^{(\mathbf{k})} = \mathbf{s}^{(\mathbf{k-1})} + a_{\hat{\gamma}}^{(k)}\phi_{\hat{\gamma}}^{(k)} \quad (16)$$

where $a_{\hat{\gamma}}^{(k)} = \langle \mathbf{R}^{k-1}, \phi_{\hat{\gamma}}^{(k)} \rangle$ and $\mathbf{R}^k = \mathbf{s} - \mathbf{s}^{(k)}$. After m steps an approximation to Eq 1 is obtained, with residue $\mathbf{R} = \mathbf{R}^{(m)}$.

2.5 Classifier

After training the dictionary, a multilayered perceptron NN with 128 inputs is used to classify frames and detect apnea events. We tried different structures, i.e., we used differents numbers of neurons (5-20) in the hidden layer, always with sigmoidal activation functions; also for the output, we used 3 outputs neurons, to obtain in each neuron the normal, apnea or hipopnea respiration with sigmoid activation functions, or in other experiment with one output neuron with a linear activation function which resolved in three levels all types of respiration. From now, when we said the NN we talk about a multilayered perceptron with 128 inputs, 20 neurons in hidden layer and one output neuron which regress the apnea-hipopnea marks, this structure has the best results. The activation function for all neurons is:

$$tansig(x) = \frac{2}{1+e^{-2x}} - 1. \quad (17)$$

The Levenberg-Marquardt algorithm was used to train network with the activation coefficients as input and the apnea-hipopnea marks as target where "1" is a apnea-hipopnea event and "-1" is a normal respiration. The same train dataset used to learn the dictionary was used as training set.

After the training stage, the dictionary and the NN parameters are fixed. The test set is analyzed using frames of 128 samples. With the trained dictionary, we found the coefficients $\mathbf{a}$ by the maching pursuit algorithm. After that, we use that activation vector as input to the trained multilayered perceptron and the output give us "-1" if the frame is classifield as normal or "1" if it represents an apnea-hipopnea. At last, we sum all the apnea-hipopneas and divide by the sleep time (in hours) and obtain the oxigen desaturation index (ODI), which is used as a predictor of the apnea-hipopnea index (AHI) which is a parameter of the degree of illness.

3 MATERIALS AND PRELIMINARY RESULTS

For this preliminary work we used the Sleep Heart Health Study Polysomnography Database. This database include 1000 overnight polysomnograms to study the relationship of sleep disordered breathing and cardiovascular disease. Its has several biomedical signals for each study: EEG, EOG, EMG, ECG, nasal airflow and respiratory effort signals, SaO_2 and

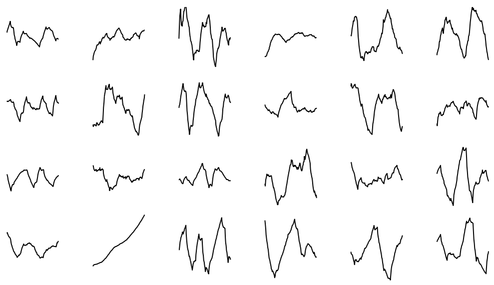

Figure 2: Some atoms of the dictionary estimated by NOCICA.

heart rate, and annotations of: sleep stages, arousals and respiratory events. The respiratory events of our interest are defined as,

- apneas: if the amplitude of the airflow signal decreases below at least 25% of the amplitude of "baseline" breathing (identified during a period of regular breathing with stable oxygen levels), and is longer than 10 seconds.
- hypopneas: if the amplitude of any respiratory signal decreases below 70% of the amplitude of "baseline", and is longer than 10 seconds and for >2 breaths. These require at least a 2% desaturation [18].

Both events are marked at the nasal airflow signal. This signal has a different dynamic than the SaO_2 signal, usually a desaturation occurs several seconds after the apnea event. Moreover, the desaturation and the flow reduction have different duration. In this work we try to estimate the apnea-hipopnea event using the SaO_2 signal, so we measure the ODI to predict the AHI.

We used segments of 128 samples of the SaO_2 signal of 84 studies to train the dictionary. The data set has studies with different grades of illness, a balanced train set was created by 21 studies with AHI <5, 21 with 5<AHI<10, 21 with 10<AHI<15, and 21 with AHI>15. The rest of the data set is an unbalanced test set with approximately 916 studies with different degrees of illness.

We use the wavelet processing technique proposed in [9] to denoise the SaO_2 signal. The signal, sampled at 1Hz, was denoised by zeroing the approximation coefficients, at level 8, of the dyadic discrete wavelet transform with mother wavelet Daubechies 2. This procedure has the effect of a highpass filter to eliminate the baseline wander and low frequency noise. From now, when we talk about SaO_2 signal is the denoised one.

After that, we use all apnea-hipopnea segments and randomize up to 200 frames of 128 samples of each studies; so we train the dictionary with 16800 frames of 2 minutes each one. In Fig. 2 we show an extract of a complete 128×128 dictionary generated by NOCICA

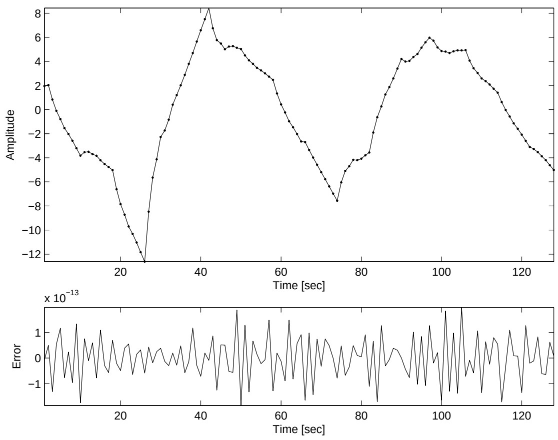

Figure 3: Top: in solid line a segment of SaO_2 signal which has marks of apneas, in dots the reconstruction by all the atoms in the dictionary. Down: the instantaneous reconstruction error.

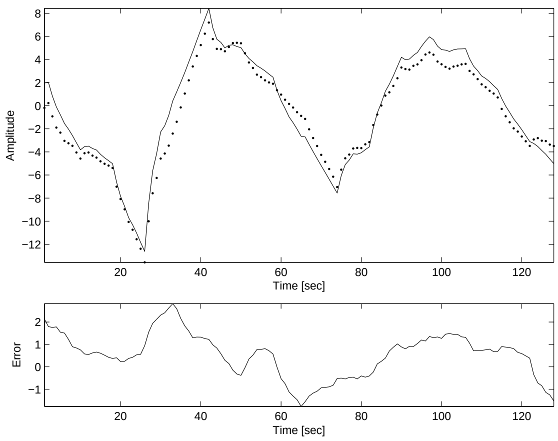

Figure 4: Top: in solid line a segment of SaO_2 signal which has marks of apneas, in dots the reconstruction by the 32 (25%) most important atoms of the dictionary. Down: the instantaneous reconstruction error.

method. Observe that the dictionary capture different specific waveforms, for example: the wave in the first row and the fourth column takes low frequencies information, the one in the fourth row and the second column looks like a linear function, and all over the dictionary there are sawtooth patterns with different numbers of peaks which are similar to the desaturation patterns that can be observed in a SaO_2 signal.

We show in Fig. 3 a frame of the SaO_2 signal of the test set with some desaturations, the reconstruction of that segment and the instantaneous errors; the dictionary represents the desaturation with a little error (in order of 10^{-13}). The same segment is shown in Fig. 4 but the reconstruction was performed with the 32 most important atoms, measured by their absolute activation value (this represent the 25% of the dictionary).

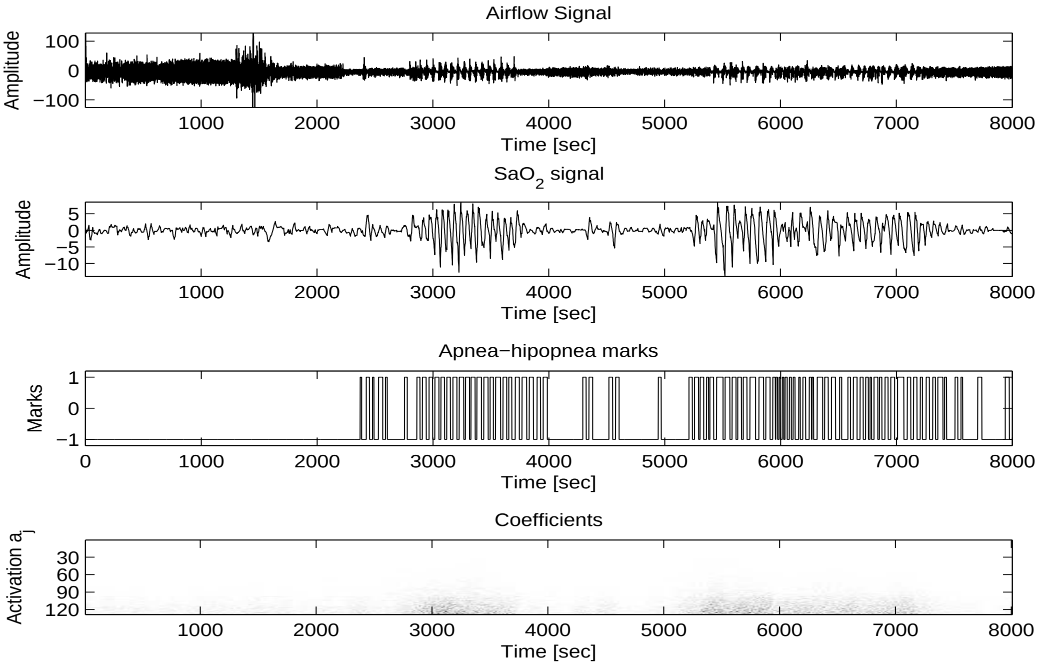

Figure 5: SR Activation. Top: the airflow and the filtered SaO_2 signal. Next the apnea-hipopnea marks where "-1" is a normal and "1" is a apnea-hipopnea event. Down: the activation of the atoms.

On other hand in a more general example, we could see in Fig. 5, 8000 seconds of the SaO_2 signal, the apnea-hipopnea reference marks and the absolute value of the amplitude of the a_j. It can be seen that different atoms are activated for the areas of normal breathing and apnea-hipopnea. This difference in the activation patterns of some atoms for each class motivate us to, in future works, train the NN only with the activation of the most discriminative ones.

At last we trained the multilayered perceptron, with the same dataset used to learn the dictionary. We use the maching pursuit algorithm and the trained dictionary to obtain the coefficient of the atoms of each frame, and we used that coefficients as the input of the NN. The target output are the apnea-hipopnea marks; we evaluate if there is a respiratory event of more than 10 seconds to label that segment as [-1,1], where "-1" is a normal frame and "1" is a apnea-hipopnea frame. For postprocessing we applied a simple threshold at different levels and we check if the respiratory event detected by the NN has more than 10 seconds; after that we count all the checked respiratory events and divide by the total time to measure the ODI.

In Fig. 6 we show a linear regression between the AHI and the ODI that we predict from the NN output using a threshold level of 0.8. It can be seen that the AHI and ODI grows similarly, also we obtain a correlation coefficient of 81.74% (p-value = 4.4335^{-13}). Finally we obtain a sensitivity of 84.27% and a specificity of 85.44% for the detection of patient with a severe OSAH (AHI > 15), calculated as the minimun distance to the point [0,1] at the ROC curve.

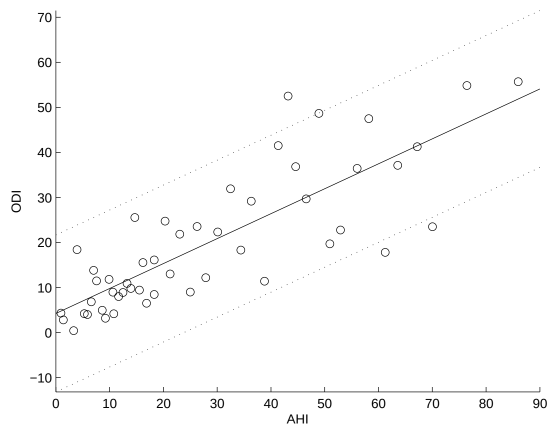

Figure 6: A regression plot between AHI and ODI. In circles the scatter plot of the studies. In solid line, the linear regression between ODI and AHI, and in points two times the standard deviation.

4 CONCLUSIONS

We applied the SR technique to segments of the SaO_2 signals for the detection of OSAH. SR was never used for this application. The advantages of the detection of AHI by the SaO_2 signal is that it is easy to adquire, cheaper and do not annoy the patient as a polysomnogram.

We find a dictionary which can represent the desaturations of the SaO_2 signal, and it is able to produce a SR of the SaO_2 signal. We find that we could reconstruct the signal from this dictionary with little error, moreover, if we only use the 32 most important coefficients the information is preserved. We used a multilayered perceptron to non-linearly regress the apnea-hipopnea marks by the desaturations at the SaO_2 signal and we obtain good results. Finally, a linear regression between AHI and ODI show that this procedure could be used to detect OSAH because there is a high correlation between them.

For future works, we will select the atoms of the dictionary which has the most discriminative information to train the multilayered perceptron in order to simplify the model. Also, to improve the results of the NN, we could modify the topology, the activation functions and the training method or use others type of classiffiers. We are also interested in studying the advantages of the overcomplete and subcomplete dictionaries in the application of the apnea-hipopnea detection. Moreover, we are exploring the online apnea-hipopnea detection, and for this reason we need to use or develop faster algorithms to test a study.

5 ACKNOWLEDGEMETS

The authors wish to thank: Agencia Nacional de Promoción Científica y Tecnológica and Universidad Nacional de Litoral (with PAE-PICT 00052), Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET) from Argentina, and M.D. Luis Larrateguy for their support.

REFERENCES

[1] S.A. Abdallah. *Towards music perception by redundancy reduction and unsupervised learning in probabilistic models*. Ph.D. Thesis, Department of Electronic Engineering, King's College London, 2002.

[2] American Academy of Sleep Medicine. *The international classification of sleep disorders, revised. Diagnostic and coding manual*. AASM, 2nd edition, 2001.

[3] Consenso Argentino de Trastornos Respiratorios Vinculados al Sueño. Trastornos respiratorios durante el sueño. *Medicina*, 61(3):351–363, 2001.

[4] F. Guspi and B. Introcaso. Soluciones ralas de sistemas lineales indeterminados. *El Ingeniero en la Red 1(VII) (Revista Electronica FCEIyA, UNR, Argentina)*, 16:1–10, 2000.

[5] J.L. Hossain and C.M. Shapiro. The prevalence, cost implications, and management of sleep disorders: An overview. *Sleep and Breathing*, 6:85–102, 2002.

[6] T.P. Jung, S. Makeig, T.W. Lee, M.J. McKeown, G. Brown, A.J. Bell, and T.J. Sejnowski. Independent component analysis of biomedical signals. Helsinki, 2000.

[7] O.-W. Kwon and T.-W. Lee. Phoneme recognition using ICA-based feature extraction and transformation. *Signal Processing*, 84(6):1005–1019, 2004.

[8] J.-H. Lee, H.-Y Jung, T.-W. Lee, and S.-Y. Lee. Speech feature extraction using independent component analysis. In *Acoustics, Speech, and Signal Processing, 2000. ICASSP '00. Proceedings. 2000 IEEE International Conference on*, volume 3, pages 1631–1634, 2000.

[9] F. Lestussi, L. Di Persia, and D. Milone. Comparison of on-line wavelet analysis and reconstruction: With application to ECG. In *5th International Conference on Bioinformatics and Biomedical Engineering, (iCBBE)*, pages 1–4, may 2011.

[10] M.S. Lewicki and B.A. Olshausen. A probabilistic framework for the adaptation and comparison of image codes. *J. Opt. Soc. Am.*, 16:1587–1601, 1999.

[11] M.S. Lewicki and T.J. Sejnowski. Learing overcomplete representations. *Advances in Neural Information Processing 10 (Proceedings NIPS'97)*, pages 556–562, 1998.

[12] S. Makeig, A.J. Bell, T.-P. Jung, and T.J. Sejnowski. Independent component analysis of electroencephalographic data. *Adv. Neural Inform. Process. Systems*, 8:145–151, 1996.

[13] S. Makeig, T.-P. Jung, A.J. Bell, D. Ghahremani, and T.J. Sejnowski. Blind separation of auditory event-related brain responses into independent components. *Proceedings of the National Academy of Sciences*, 94(20):10979–10984, 1997.

[14] S.G. Mallat and Z. Zhang. Matching pursuits with time-frequency dictionaries. *Signal Processing, IEEE Transactions on*, 41(12):3397–3415, dec 1993.

[15] B.A. Olshausen and D.J. Field. Natural image statistics and efficient coding. *Workshop on Information Theory and the Brain*, 7(2):524–531, 1995.

[16] B.A. Olshausen and D.J. Field. Vision and the coding of natural images. *Am. Sci.*, 88:238–245, 2000.

[17] M.D. Plumbley, S.A. Abdallah, J.P. Bello, M.E. Davies, G. Monti, and M.B. Sandler. Automatic music transcription and audio source separation. *Cybernetics Systems*, 33:603–627, 2002.

[18] Sleep Heart Health Study. *Reading center manual of operations: scoring procedures*. SHHS, 2002.