Artificial Immune Systems for text-dependent speaker recognition

K. M. Faraoun[1]
A. Boukelif [2]
[1] Evolutionary Engineering and Distributed Information
Département d'informatique, Djillali Liabès University.
Systems Laboratory, EEDIS -SBA– Algeria
kamel_mh@yahoo.fr
. [2]Communication Networks ,Architectures and Multimedia laboratory
Département d'électronique, Djillali Liabès University.
University of S.B.A. Algeria
aboukelif@yahoo.fr

Abstract. This paper shows the potential accomplishments of artificial immune systems (in particular, the negative selection algorithm) application to the problem of speaker recognition. Both the use of binary representation of original signal and that of its Fast Fourier Transform in a real-number representation are analysed. A number of experiments are performed on different datasets to examine the performance evolution with respect to the different system parameters. It is found that substantial enhancements of the system capabilities are possible by means of the exploitation of the Fast Fourier Transform.

Keywords: Speaker recognition, Artificial immune system, Negative selection

(Received December 30, 2005 / Accepted July 06, 2006)

1 Introduction

Speaker recognition is a biometric-based technology (technology that verifies or identifies individuals by analyzing a facet of their physiology and/or behavior) that refers to automatic voice detection technologies, including speaker identification and speaker verification. Speaker identification is the process of finding the identity of an unknown speaker by comparing the voice of that unknown speaker with voices in a database of speakers. It entails a one-to-many comparison. Speaker verification is the process of determining whether a person is who she/he claims to be. It entails a one-to-one comparison between a newly input voiceprint (by the claimant) and the voiceprint for the claimed identity that is stored in the system.

Speaker recognition systems have a large set of applications in everyday life: Time and Attendance Systems, Access Control Systems, Telephone-Banking, Biometric Login to telephone aided shopping systems, Information and Reservation Services, Security control for confidential information and Forensic purposes.

All speaker recognition systems contain two main modules: feature extraction and feature matching. Feature extraction is the process that extracts a small amount of data from the voice signal that can later be used to represent each speaker [1]. Feature matching involves the actual procedure to identify the unknown speaker by comparing extracted features from his/her voice input with the ones from a set of known speakers. Speaker recognition is a difficult task and it is still an active research area. In practice, this task has been challenged by the highly variant of input speech signals: a speech signal includes the presence and type of speech pathologies, the physical and emotional state of the speaker and a can be also impregnated with the acoustical noise and environment where the recording is done. Often, humans are able to extract the identity information when the speech comes from a speaker they are acquainted with. However, in an automatic verification, a routine must be developed to accommodate this kind of parasitic alterations.

As a result, different kinds of speaker recognition systems tools and methods were built based on different methods like:

- Neural networks learning, [3];
- The Bayesian Maximum A Posteriori (MAP) Adaptation Method [4];
- Statistical analysis and vector quantization [5];
- Gaussian mixture models (GMM) [6];
- Hidden Markov models (HMM)[7].

In this work, an attempt is made to show the use of the negative selection algorithm (an artificial immunology based algorithm) to build an efficient speaker recognition system. Section 2 gives an overview on the theory developed around the speaker recognition technology along with a short synopsis on artificial immunology and the negative selection algorithm. Details on the method developed in the present work and descriptions of the experiments conducted are given in section 3. Section 4 presents a discussion of the obtained results. The paper is finally concluded with a summary of the most important points.

2 The speaker recognition

A speech signal is a very complex function of the speaker and his environment that can be captured easily with a standard microphone. Each voice signal is represented in a waveform. After its acquisition by a microphone, a sound is converted to electrical current. Continuous oscillations of air pressure become continuous oscillations of voltage in an electrical circuit. This fast-changing voltage is then converted into a series of numbers by a digitizer. A digitizer acts like a very fast digital voltmeter. It makes thousands of measurements per second. Each measurement results in a number that can be stored digitally (that is, only a finite number of significant digits of this number are recorded). This number is called a sample and the whole conversion of sound to a series of numbers is called sampling. The result of the sampling operation is a numbers vector, witch represent the voice signal waveform. The numbers range depends on the sampling bit-rate (16-bit or 8-bit in our work). Any speaker recognition system, use on the obtained vector to extract different features and characteristics of the voice signal, and to do any kind of analysis [2].

Speaker recognition systems are classified as text-dependent (fixed-text) and text-independent (free-text). The text-dependent systems require a user to re-pronounce some specified utterances, usually containing the same text as the training data, like passwords, card numbers, PIN codes, etc. There is no such constraint in text-independent systems. Speaker models capture characteristics of somebody's speech that show up irrespective of what one is saying. In the text-dependent system, the knowledge of knowing words or word sequence can be exploited to improve the performance.

The used representation is also a very important element for the recognition system. Some techniques use the initial form of the voice signal (obtained directly by the sampling phase). Others employ a transformed form of the signal, mainly via Fast Fourier transform (FFT). The FFT transformation permits to work in the frequency domain and hence to use the frequency spectrum of the voice signal instead of the wave form. This form gives more information about the signal structure and can be more characterized to each speaker.

Some other applications use the MFCC (Mel-Frequency Cepstrum Coefficients) [8]. MFCC's are based on the known variation of the human ear's critical bandwidths with frequency. The main MFCC role is to construct a voiceprint for each speaker, based on its voice signal characteristics. This voiceprint is then used by the identification process. Only voiceprints are stored in the database instead of storing the whole voice signal. When a new voice is recorded, its voiceprint is extracted and then compared to all voiceprints in the database until identification or reject. MFCC voiceprint combined with neural network techniques has been used to build a robust and flexible speaker recognition system [9].

In the present work, a new form of voiceprint is created for each speaker, given by his voice signal, using the negative selection algorithm. This algorithm is one of the main components of artificial immune systems.

3 The artificial immune systems

Artificial immune systems (AIS) are adaptive systems inspired by theoretical immunology and observed immune functions, principles and models. They form the basis of solutions for various real world problems, in particular, intrusion detection.

The natural immune system is a network of cells, tissues, and organs that work together to defend the body against attacks by "foreign" invaders. So any artificial immune system must give a model for each element and each inspired mechanism. Mechanisms are implemented as algorithms, when elements (cells, tissues...) are represented by binary strings or real vectors (depending on the problem definition). There are various mechanisms in the artificial immune system such as clonal selection, affinity maturation, somatic hyper-mutation, receptor editing and negative selection.

4 The negative selection algorithm

Through the use of the negative selection process, there have been a number of works attempting at building artificial immune systems for virus detection [10], computer security [11], hardware fault tolerant systems [12] and Time series analysis [13]. The original work by Forrest, Perelson et al. in 1994 [14], in which the negative selection algorithm was proposed, has been a starting point for almost all the research in the AIS related to the computer security.

The negative selection algorithm is inspired by the maturation of T-cells in the thymus gland [15]. The algorithm consists of two stages: censoring and monitoring. The censoring phase caters for the generation of change-detectors. Subsequently, the system being protected is monitored for changes using the detectors generated in the censoring stage. The basic principle of a negative-selection algorithm is as follows:

- Define **self** as a multi-set N_S of strings of length l over a finite alphabet, a collection that we wish to protect or monitor. For example, N_S may be a segmented file, or a normal pattern of activity of some system or process.
- Generate a set N_R of detectors, each of which fails to match any string in N_S. A partial matching rule is used to compare the strings.
- Monitor N_S for changes by continually matching the detectors against N_S. If any detector ever matches, a change (or deviation) must have occurred.

Matching between detectors and self-strings is done via a matching rule witch indicate for each two strings of the same length l, and with the same alphabet, if they match or not. Obvious approximate matching rules include Hamming distance and Euclidian distance, but the more adopted actually and the more plausible rule in the immunology concept, is the so called r-contiguous bits [15]: Two strings match if they have r contiguous bits in common. The parameter r is the threshold of the matching rule that determines the specificity of the detector. It is an indication of the size of the subset of strings that a single detector can match. If r = l, then the matching is completely specific, that is, the detector will match only a single string (itself); but if r = 0, the matching is completely general, that is, the detector will match every single string of length l.

5 Negative selection algorithm and the speaker recognition

The main goal to be achieved in almost all applications of the negative selection algorithm is to detect abnormal deviation from normal behavior. The algorithm generates detectors from a segmented version of the original data (the self set Ns) whose representation differs from one application to another. In general, a binary representation is used to codify the self-elements. The elements have the same fixed length (witch can be a parameter of the algorithm).

Accordingly, the negative selection algorithm is used in the present work to construct a set of detectors for a given speaker voice signal. The generated detectors are then used as a voiceprint to monitor the acquired new voice signals (identification phase). If the signal is produced by the same speaker, the form and the data distribution must be very similar to the original one (used to generate detectors), so a very low anomalies rate will be obtained (null in the ideal case).

According to the obtained anomalies rate, the automatic speaker recognition system decides of the new voice speaker identity. To achieve that, a database of voiceprints of different speakers is used. A voiceprint is given by the set of detectors obtained when applying the negative selection algorithm to the corresponding voice signal (the learning phase). If the lowest obtained anomalies rate is higher than a fixed threshold ∂ (witch a parameter of the system that determine the highest accepted value of the detected anomalies rate) , the system decides that this voice signal does not belong to any speakers of the database, and so the speaker is not identified. Else, the speaker is identified as the one with the lowest anomalies rate (lower than ∂). Fig.1 resumes how the recognition system operates.

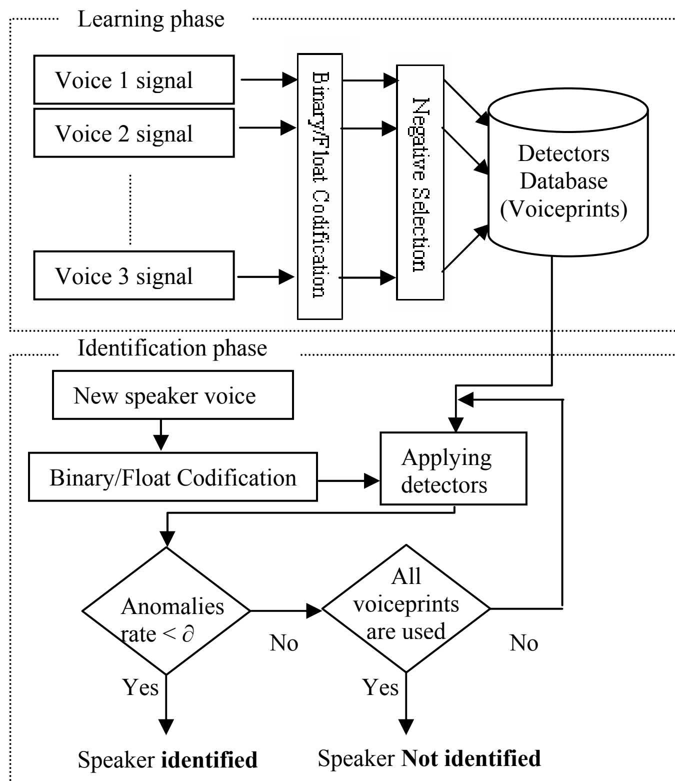

Figure 1: Speaker recognition system based on the negative selection algorithm

6 Data representation

Two approaches were used to built the automatic speaker recognition system. The first one uses the voice single waveform as input when the second uses its Fast Fourier Transform (FFT).

In the first approach, the result of the sampling operation of the voice signal is a vector that contains either integer or byte values in the case of 16-bits or of 8-bits sampling rate respectively. In both cases, a binary representation can be used. In our implementation, the sampled vector is decomposed into bit-strings of length l, where l can be varied from 8 to 64 bits (a parameter of the system). The negative selection algorithm is then applied on the resulting strings set to generate detector with the same length l. The number of generated detectors is also a variable parameter of the system; it can be varied to study system performance.

In the second approach, using FFT, the resulting vector contains real numbers. These values are used with a real value adapted negative selection algorithm [17]. Generated detectors have also real values, and the matching rule used is the Euclidian distance. Implemented algorithm is presented in figure 2.

```
Negative_Selection (In :S,m,r_s; out :D);
S: set of self samples
m: number of detectors
r_s : detection radius
1: D ← ∅
2: Repeat
3:    x ← random number from [0,1]^n
4:    repeat for every s_i in S={s_1,i=1,2,….}
5:             d ← Euclidian distance between s_i and x
6:            if d< r_s then go to 2
7:    D ←D ∪ {x}
8: Until | D | = m
9: Return D
```

Figure 2: Real valued negative selection algorithm

7 Experiment

As explained above, our system is text-dependent. Accordingly, the used dataset must contain different copies voice signals of the same speaker, using the same text. This can be achieved by collecting voice signals at different periods, or in different situations (in noisy environment, the speaker has a cold, etc…).

For our experiments, we have use the YOHO CD-ROM Voice Verification Corpus available at the Linguistic Data Consortium (LDC) [18]. The data is transformed and represented in the format detailed above. The used dataset contain a total of 138 speakers. We have chosen only 10 speakers to perform our tests (from 105 to 115).

Different combinations of parameters were tested to examine the system performance with respect to each one. All tests were performed on an Intel-Pentium 4 CPU 2.66Ghz with 256 Mo Ram size.

8 Results and discussion

8.1 Binary representation

The length of each element of the self set after segmentation (parameter l) was first set to 16 bits. The detectors number was set to 100, the matching threshold r to 8 and the detection threshold ∂ to 0.4. The r-contiguous matching rule is used to match detectors with signal elements (self set).

Once a voiceprint of a speaker is constructed, the identification system can be tested in two ways: (i) using the voice signal of other speakers (user discrimination), (ii) using the voice of the same speaker recorded in a different context (after some days, or with some noise). The system is expected to get high anomalies rate in the first case and low one on in the second.

Accordingly, we generated 100 detectors for the speaker number 112 (its voiceprint) using the segmented version of the waveform. The length of the elements was set to 16. The obtained detectors were applied on the signals of speakers 110, 111 and 114. In each case, the anomaly rate (ranging in the interval [0,1]) was calculated following formula:

$$\textit{Anomaly rate} = \frac{\sum \text{number of detected anomalies}}{\sum \text{number of self elements}} \qquad (1)$$

The obtained values were of 0.63, 0.71 and 0.77 for the speakers 110, 111 and 114 respectively. The same test applied to different versions of the voice signal of the same speaker (same text recorded with a time interval of one week) the anomalies rate was of 0.27. Table.1 resumes different results obtained for the speaker 112. According to the used detection threshold, the signal of the speaker 112, acquired one month after the original recording, was not recognized. This can be explained by the fact that signal characteristics have endured significant changes after this period so that it's self definition is no longer preserved. The new signal is then considered as a new self (A different speaker).

To solve this problem, the generated detectors must be at the same time specific to the signal (in order to insure its discrimination), and sufficiently general to efficiently cover the legitimate variation space. To achieve that, one has to operate a number of changes on the different system parameters.

Table 1: Obtained results using the speaker "112" for speaker data for experiments

Speaker	Anomalies rate	Result
110	0.63	not identified
111	0.71	not identified
113	0.49	not identified
114	0.77	not identified
115	0.70	not identified
116	0.85	not identified
107	0.89	not identified
108	0.79	not identified
109	0.77	not identified
112 after 1 week	0.27	**Identified**
112 after 2 weeks	0.31	**Identified**
112 after 1 month	0.42	not identified

To study the system performance, different combinations of the parameters have been tested. The testes were performed using 10 speakers from 105 to 115. Both negative samples (voice signals of other speakers) and positive samples (voice signals of the same speaker recorded in a different context) have been used. The system is required to identify the positive samples and reject the negative ones.

Let N_g and N_p be the number of negative and positive samples respectively, and AN(i) the anomalies rate obtained for the i^{th} sample computed using equation (1). Three performance measures can be computed for each speaker:

- Negative detection rate:

$$NGR = \frac{\sum_{i=1}^{N_g} AN(i)}{N_g} \quad (2)$$

- Positive detection rate:

$$PSR = \frac{\sum_{i=1}^{N_p} AN(i)}{N_p} \quad (3)$$

- Global detection rate :

$$GR = \frac{NGR}{PSR} \quad (4)$$

According to these equations, a good configuration of the system is achieved when NRG is maximized, and PSR is minimized, and so the global detection rate is in its maximum.Fig.3 shows the variations of NGR and PSR according to the different used values of the self-elements length (Tests operated on the 112 speaker). Fig.4 shows variations of GL according to the same parameter. It can be noted that NGR increases and PSR decreases with increasing self-element length. Hence, the maximum global detection rate is achieved for larger self-element lengths.

In order to have better appreciation of system, the average of each performance measure is taken over all the used speakers (10 in our experiments):

$$AV_{ng} = \frac{1}{10}\sum_{i=1}^{10} NGR_i \quad (5)$$

$$AV_{np} = \frac{1}{10}\sum_{i=1}^{10} PSR_i \quad (6)$$

$$AV_{gl} = \frac{1}{10}\sum_{i=1}^{10} GL_i \quad (7)$$

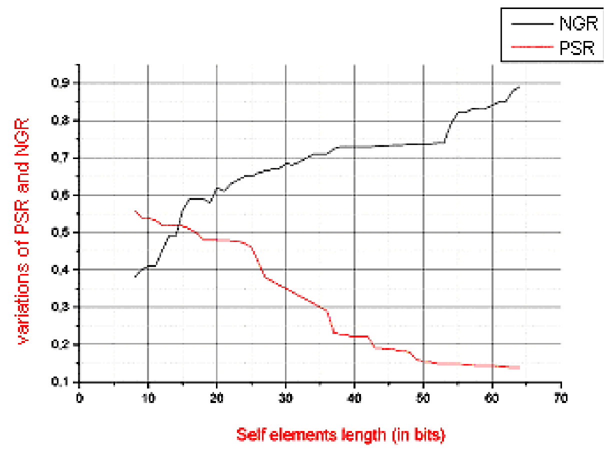

Figure 3: Variations of PSR and NGR according to the self elements length.

Figure 4: Variations of GL according to the self elements length.

Figures 5 and 6 show the variations of AV_{ng}, AV_{np} and AV_{gl} with respect to the elements length and the matching threshold value respectively. These experiments confirm that better performance is reached with increasing length of the self-elements. This can be explained by the fact that longer strings can extract more characteristics from the signal, so the generated detectors are more adapted to the signal information.

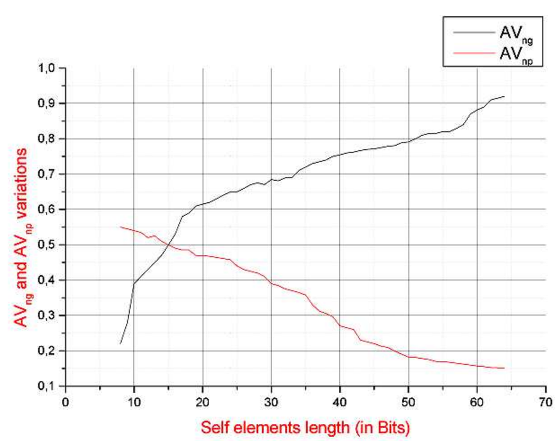

Figure 5: Variations of AV_{np} and AV_{ng} according to the self elements length.

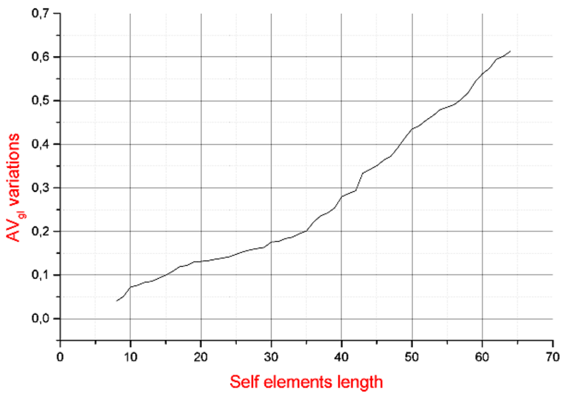

Figure 6: Variations of AV_{GL} according to the self elements length.

Fig.7 shows the variation of the three parameters AV_{ng}, AV_{np} and AV_{gl} with respect to the matching rule threshold r. The threshold variation was operated keeping fixed elements length of 64 (the maximum), and detectors number of 200. Results show that the system performance measure, AV_{gl}, firstly increases with the increasing r up to a maximum limit. Beyond this limit, it goes decreasing. The optimal value of the matching threshold is found to be 53. This can be explained by a balance between two extreme situations.

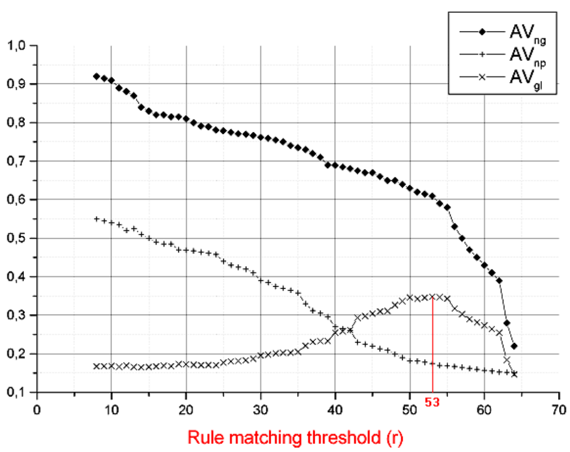

Figure 7: Variations of AV_{ng}, AV_{np} and AV_{gl} according to matching threshold variations

When the value of r is low, the detection is very sensible to small variations, and the detectors set is very specific to the self set (the used voice signal), so small and legitimate variations are interpreted as anomalies. In this case, positive samples could be rejected and incorrectly classified so that the AV_{np} parameter is relatively high. At the same time, the negative samples are easily detected and classified correctly (high AV_{ng}).

For height values of r, the generated detectors need great modification of the initial self set (the original voice signal) to detect anomalies. In this case, the positive samples can be easily classified, and the small variations introduced legitimately are ignored, so the value of AV_{np} is low. At the same time, some negative samples (that are close to the voice signal) are incorrectly recognized so that the value of AV_{ng} is low.

4.2 Real numbers representation

When employing the Fast Fourier Transform (FFT) version of the signal, a real number codification is used. The resulting vector is normalized to range in the interval [0,1]. The resulting vectors components are directly considered as self-elements. For each voice signal sample, the FFT is first computed and then transmitted to the negative selection algorithm (see fig.2). The detectors are then generated as real numbers vectors, with the Euclidian distance as matching rule. The generated set of is used to monitor each positive and negative voice sample from the dataset. The same performance measures are used to test the system performances (NGR, PSR and GR).

When using real values codification, a radius detection parameter determines the coverage of the detection space for each detector. This parameter has the same effect as the matching threshold (from the r-contiguous matching rule used with binary codification). When r increases, the detector is able to match more elements. Hence more error deviation from the original self set can be handled, leading to a greater tolerance for the positive samples, and a difficulty of rejecting negative ones (anomalies could be considered as legitimate changes).

Also, when r decreases, detectors get more specific to the signal characteristics, and the negative samples are easily rejected. At the opposite, many positives samples could be rejected. This definitely leads to and system performances decrease. Our experiments show that the value of 0.33 for the radius r is a good compromise.

The most important parameter in this method is the size N of the sliding window used to decompose the signal. It determines the dimension of the vectors representing the self-elements and the detectors. We tested different values ranging from 4 to 20. The computational complexity of the algorithm grows when using higher values of N.

Fig.8 shows the variations of the parameter AV_{gl} with respect to the size of the used window. The test is achieved with 500 generated detectors. It can be seen that the system reaches it's maximum performances when N is maximum. This can be explained by the fact that large windows can extract more signal transformations characteristics, and then the speaker can be identified more exactly.

4.3 Binary vs. real-number representation

Table.2 summarizes the system parameters considered as optimal according to the results presented above. To perform comparison between this implemented methods (binary and real vectors coding), we studied the performance variation with respect to the generated detectors number. This one is the only common parameter that could allow deciding which method is more efficient.

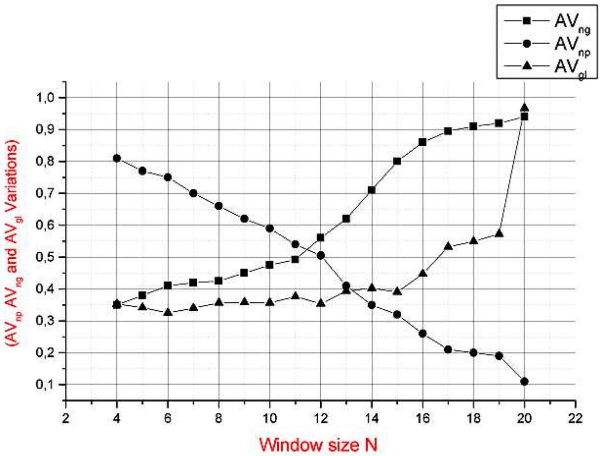

Figure 8: AV_{np} AV_{ng} and AV_{gl} variations according to the windows sliding size N.

Table 2: Optimal parameters values used for the speaker recognition system using the binary and the real values approaches

	Elements length	Matching threshold	threshold ∂
Binary representation approach	64	53	0.4
FFT transformed vectors approach	20	0.33	0.4

Fig.9 shows the variation of the performance of the recognition system with respect to the number of generated detectors in both cases. It is clear that for the same detectors number, using the FFT with real vectors coding allows better performance than simple binary coding. The detection rate is visibly improved and the false alarms rate is minimized. This can be explained by the fact that the use of the FFT carries a better definition of the self set using only the important signal components. The constructed voiceprint is more representative and gives more information about the signal characteristics.

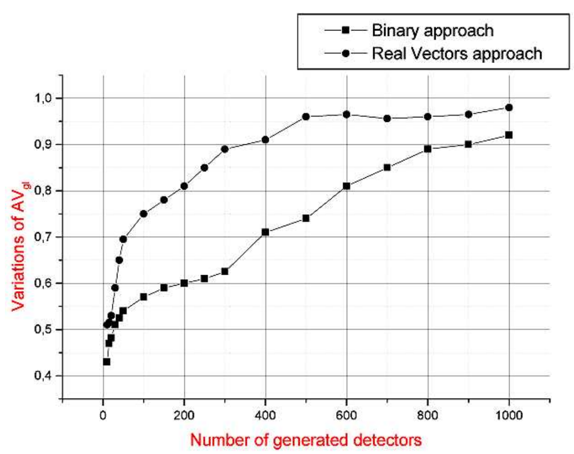

Figure 9. Variations of AV_{gl} for the two approaches according to the number of generated detectors

5 Conclusion

In this work, the negative selection algorithm from artificial immune system technology was applied to build a speaker recognition system. The voice identification problem is considered as anomaly detection one, and the negative selection algorithm is employed to construct voiceprints for speakers. The monitoring phase of the negative selection algorithm corresponds to the detection phase of the speaker recognition system. This approach does not take into account the phonetic properties of the speech, but only the information presented by its numerical form.

Two situations were analysed: (i) the use of binary representation of original signal and (ii) the use of its Fast Fourier Transform (FFT) in a real-number representation.

In both cases, a series of experiment was performed in order to examine the system functioning variation with respect to the different system parameters, in particular the self-elements length and the matching threshold. It is found that increasing the elements length significantly enhance the system performance. On the other hand, there is an optimal value of the matching threshold that maximises the system global detection parameter, hence supplying the optimum performance.

Experiments were carried out in order to assess the relative efficiency of the implemented methods, for instance the binary and the real-number representations. According to the obtained results, it can be asserted that substantial improvements of the performance of speaker recognition system are feasible by means of the use of the FFT along with real number representation.

6. References

[1] Minh N. Do "An Automatic Speaker Recognition System", Audio Visual Communications Laboratory Swiss Federal Institute of Technology, Lausanne, Switzerland, February 2003.

[2] J.R. Deller Jr., J.H.L. Hansen, and J.G. Proakis, Discrete-time Processing of Speech Signals, Macmillan publishing Company, New York, 2000.

[3] J.P Hosom, R.Cole & M.Fanty, "Speech Recognition Using Neural Networks at the Center for Spoken Language Understanding", NSF-Graduate Research Traineeships project, Center for Spoken Language Understanding (*cslu*), Oregon Graduate Institute of Science and Technology , USA, July 1999.

[4] J. L. Gauvain and C.-H. Lee, "Maximum a posteriori estimation for multivariate gaussian mixture observation of markov chains," in *IEEE Transactions on Speech Audio Processing*, April 1994, vol. 2, pp. 291–298.

[5] F. Bimbot and L. Mathan. Second-order statistical measures for text-independent speaker identification. In Proc. ESCA Workshop on Automatic Speaker Recognition Identification and Verification, pages 51-54, 1994.

[6] Douglas A. Reynolds, Thomas F. Quatieri, and Robert B.Dunn, "Speaker verification using adapted gaussian mixture models," Digital Signal Processing, vol. 10, no. 1–3, 2000.

[7] Kremer, H. "Incremental Speaker Adaptation Techniques for HMM Based Speech Recognition Systems". M.Sc dissertation, Ben-Gurion University, Beer-Sheva, Israel, 2002.

[8] Minh N. Do “An Automatic Speaker Recognition System”, Audio Visual Communications Laboratory Swiss Federal Institute of Technology, Lausanne, Switzerland, February 2003.

[9] R.Boite, H.Boulard, T.Dutoit, J.Hancq & H.Leich, « Traitement de la parole » Book, Collection Electricité, Presses Polytechniques et Universitaires Romandes, 1999.

[10] P. K. Harmer and G. B. Lamont. Agent Based Architecture for a Computer Virus Immune System. GECCO 2000, Las Vegas, Nevada, USA, July 8, 2000.

[11] J. Kim and P. Bentley. Towards an Artificial Immune System for Network Intrusion Detection: An investigation of Clonal Selection with a negative Selection Operator. The Congress on Evolutionary Computation. (CEC-2001), Seoul, Korea, May 27-30, 2001.

[12] D. Bradley and A. Tyrrell. Hardware Fault Tolerance: An Immunological Solution. In the proceedings of IEEE International Conference on Systems, Man and Cybernetics (SMC), Nashville, October 8-11, 2002.

[13] D. Dasgupta and S. Forrest. Novelty Detection in Time Series Data using Ideas from Immunology. In ISCA 5th International Conference on Intelligent Systems, Reno, Nevada, June 19-21 1996

[14] S. Forrest, A. S. Perelson, L. Allen and R. Cherukuri. Self-Nonself Discrimination in a Computer. Proceedings of IEEE Symposium on Research in Security and Privacy. 202-212, Oakland, CA. (1994).

[15] P. Helman and S. Forrest. An Efficient Algorithm for Generating Random Antibody Strings. Technical Report CS-94-07, The University of New Mexico, Albuquerque, NM (1994).

[16] Tool Breakage Detection in Milling Operations using a Negative-Selection Algorithm. Dipankar Dasgupta and Stephanie Forrest Dept. of Computer Science University of New Mexico Albuquerque, NM 87131 Technical Report No. CS95-5 August 15, 1995

[17] F. González, D. Dasgupta, and R. Kozma, “Combining negative selection and classification techniques for anomaly detection,” in Proceedings of the 2002 Congress on Evolutionary Computation CEC2002, D. B. Fogel, M. A. El-Sharkawi, X. Yao, G. Greenwood, H. Iba, P. Marrow, and M. Shackleton, Eds. USA: IEEE Press, May 2002, pp. 705–710.

[18] J. Campbell, “Testing with the YOHO CD-ROM voice verification corpus,” Proc. Intl. Conf. on Acoustics, Speech, and Signal Processing, Detroit, May 1995, pp. 341–344.