

# An a-contrario Approach for Face Matching

Luis D. Di Martino, Javier Preciozzi, Federico Lecumberry and Alicia Fernández

*Instituto de Ingeniería Eléctrica, Universidad de la República, Montevideo, Uruguay*

**Keywords:** Face Recognition, Face Matching, a-contrario, STASM, LBP, Extended LBP, Chi-Square.

**Abstract:** In this work we focus on the matching stage of a face recognition system. These systems are used to identify an unknown person or to validate a claimed identity. In the face recognition field it is very common to innovate in the extracted features of a face and use a simple threshold on the distance between samples in order to perform the validation of a claimed identity. In this work we present a novel strategy based in the a-contrario framework in order to improve the matching stage. This approach results in a validation threshold that is automatically adapted to the data and allows to predict the performance of the system in advance. We perform several experiments in order to validate this novel strategy using different databases and show its advantages over using a simple threshold over the distances.

## 1 INTRODUCTION

Face Recognition Systems have acquired great importance in the last two decades. They are needed in a world where automatic identification/tagging systems as those used in security applications and human-computer interaction systems are major topics of research and huge development in the industry. These work by analyzing a face in an input image, they can identify the person whose face is in the image (Identification Mode) or validate the identity claimed by the subject (Verification / Matching Mode). In this work we focus on the verification mode, in this the system works with two sets of data: the gallery ( $G$ ) and the query ( $Q$ ) datasets. The first contains  $N_G$  faces of the people registered in the system, the second is used for testing, it contains  $N_Q$  faces of people who claim one of the registered identities.

A vast amount of works have been done in this area and huge improvement can be seen from the first face recognition systems to the current ones. Nowadays, face recognition systems can achieve a recognition rate over 97% by using local features (Zhang et al., 2007; Ahonen et al., 2006; Zhang et al., 2005; Zhang et al., 2010; Huang et al., 2011) when identifying people and produce a very low rate of false positives when verifying declared identities. This performance is achieved on a database of well controlled conditions such as the FERET database (Phillips et al., 1998). However, some problems as

aging and large variations in lighting and pose still have not been fully resolved and highly degrade the performance of a face recognition system. It is still very difficult for an automatic face recognition system to recognize a person from a picture taken in heavily uncontrolled conditions, what is called “face in the wild”, this has brought great attention in the face recognition field community (Kan et al., 2013), (Zhu and Ramanan, 2012), (Masi et al., 2013). In (Marsico et al., 2013) this issue has been studied with emphasis, the authors introduces two indices in order to evaluate beforehand the quality of the input image in terms of the sample pose (index SP) and sample illumination (index SI). This allows to reject the image or ask for a new capture when the input image does not meet the necessary quality. It is also very difficult to perform the identification or verification process when the aging factor appears between images, e.g., trying to recognize a person from only one picture registered 10 years ago, this matter has been treated extensively in the literature, (Park et al., 2010), (Lanitis et al., 2002), (Ling and Soatto, 2010). This problem has been usually addressed by constructing models that simulate the changes in a person’s face due to aging. These are used jointly with systems that estimate the age of a person from a picture containing its face, the time difference between two face images is used as an input for the aging simulation system. The problem with using such approaches is that lot of face images of each individual are needed in order to build a accurate

model, because the aging process is highly dependent on the person whose face is being treated.

Face recognition systems, as many others biometric systems, are usually divided in three stages: preprocessing, feature extraction and a matching stage<sup>1</sup>. Generally, great efforts are dedicated to the innovation on preprocessing and feature extraction techniques while using classic techniques in the matching stage. In this paper we focus on the matching stage, performing the study, implementation and analysis of a novel strategy based on the *a-contrario* framework (Desolneux et al., 2003b). This approach is independent of the feature extraction process as soon as a dissimilarity distance between samples is defined. *A-contrario* models have been successfully used in applications where the searched features are very unlikely to occur by chance under some background model. Therefore, the meaningful matches can be found, in a *contrario* fashion, by training models that make possible to reject these matches that are due to chance with high precision.

Algorithms based on *a-contrario* framework were first used in the detection of alignments (Desolneux et al., 2003b), contrasted edges and grouping (Desolneux et al., 2003a). Later, its use has been extended to more complex tasks, e.g., the detection of line segments (Gioi and Jakubowicz, 2010), matching of shapes (Musé et al., 2006), matching of SIFT-like descriptors (Rabin et al., 2008) and biometric identification systems based on iris templates (Mottalli et al., 2010). In this work we extend the use of the *a-contrario* framework in order to improve the performance of a face recognition system working in the verification mode where the system is used to validate an identity claimed by a person.

The rest of the paper is organized as follows: in Section 2 we present the proposed face matching technique based in the *a-contrario* strategy, in Section 3 we provide details of the used framework, databases and conducted experiments. In Section 4 the obtained results are presented, and finally, we conclude our work in Section 5.

## 2 *a-contrario* Framework

The *a-contrario* framework is based on the *Helmholtz Principle*, which in its most general form states that whenever some large deviation from randomness occurs, a structure is perceived. Thus, we can find sig-

<sup>1</sup>This stage is commonly referred, in the pattern recognition area, as the classification stage. In particular, in biometrics applications, it is called the matching stage as it involves the match of a query subject with a known identity.

nificant events as those who are far from the random or background model. This principle can be formalized as follows:

**Definition 1** ( $\epsilon$ -meaningful event). *We say that an event of the type “a given configuration of objects has a property” is  $\epsilon$ -meaningful if the expectation of the number of occurrences of this event is less than  $\epsilon$  under the background model.*

**Definition 2** (Number of False Alarms - NFA). *Given an event of the type “a given configuration of objects has a property”, the number of false alarms (NFA) is the expectation of the number of occurrences of this event under the background model.*

Definition 1 can be rewritten in terms of the NFA defined before:

**Definition 3** ( $\epsilon$ -meaningful event). *An event  $E$  of the type “a given configuration of objects has a property” is  $\epsilon$ -meaningful if the NFA is less than  $\epsilon$ :*

$$NFA(E) < \epsilon \quad (1)$$

Let  $H_0$  be the background model: “a given configuration of objects has a property and is produced by some known model”. We define a random variable  $\mathcal{E}$ , and we analyze the observation  $E$  of this random variable considering the number of false alarms. The correct definition of this NFA is a central problem in all *a-contrario* methods. However, quite often this definition can be reduced to an expression of the following form, which gives an upper bound of the actual NFA as defined before:

**Definition 4** (Number of false alarms - NFA). *The number of false alarms (NFA) of an event  $E$  is defined as:*

$$NFA'(E) = \mathcal{N} \cdot P(\mathcal{E} \geq E | H_0) \quad (2)$$

where  $\mathcal{N}$  is the number of all possible configurations of the event  $E$ .

Moreover, we can often show that the expectation of the number of occurrences of an event  $E$  satisfying  $NFA'(E) < \epsilon$  is actually less than  $\epsilon$ . For this reason, defining an event as  $\epsilon$ -meaningful, whenever  $NFA'(E) < \epsilon$ , is still consistent with Definition 1 and ensures that the method is robust in the sense that no more than  $\epsilon$  “false detections” will be obtained due to noise.

### 2.1 *a contrario* Matching Approach

As explained before, *a-contrario* methods are based on the identification of events which are very improbable under some background model. In the case of face recognition, given two faces, we can define the

background model  $H_0$ : “The faces correspond to different persons” (we can also define  $H_1$ : “Both faces correspond to the same person”). Note that even if we do not have an explicit formulation for the  $H_0$  model, we can build it empirically from the data, since the number of samples is overwhelming. This is one of the reasons that explains the popularity of *a-contrario* methods with respect to classical hypothesis testing: the model we test against is not the one that describes the rare events but the *a-contrario* one that is in general much more easier to obtain. Given the comparison of  $N_Q$  query faces against  $N_G$  faces in the gallery dataset, it is clear that there are much more representatives that comply the  $H_0$  hypothesis than the ones that comply with  $H_1$ . This explains why if we build an empirical distribution from the data is better to work with the *a-contrario* model (because of the law of large numbers, the empirical distribution of  $H_0$  will be closer to the real distribution than  $H_1$ ).

Lets now formalize all these concepts: independently from the features used to compare two faces, we always have a distance associated to this comparison. Let  $D(q_i, g_j)$  be the distance between two faces  $q_i$  and  $g_j$  from the query and gallery datasets respectively. If we compute the distance of  $q_i$  against all the faces in the gallery, we can build the empirical distribution of the distance  $D(q_i, g_j)$  given the background model  $H_0$ . Let be  $p_{q_i|H_0}(x)$  this probability density. Given a value of distance  $D(q_i, g_j) = \delta$  between two faces, the probability that the compared faces satisfies the  $H_0$  hypothesis is computed using Equation 3.

$$P(D(q_i, g_j) \leq \delta | H_0) = \int_{-\infty}^{\delta} p_{q_i|H_0}(x) dx \quad (3)$$

In our *a-contrario* framework, the meaningful event (the one that has very low probability under the background model  $H_0$ ) is that the two compared faces correspond to the same person. Thus, a matching between two faces  $q_i$  and  $g_j$  is considered to be correct if the “Number of False Alarms” is smaller than a given threshold, which in fact is a threshold on the expectation of the improbable event. The “Number of False Alarms” for the face matching problem can be defined as follows:

**Definition 5** (Number of false alarms - NFA). *The number of false alarms (NFA) of the event E: “faces  $q_i$  and  $g_j$  correspond to the same person” is defined as:*

$$NFA(q_i, g_j) = N_Q(N_G - 1)P(D(q_i, g_j) \leq \delta | H_0) \quad (4)$$

where  $N_Q$  and  $N_G$  are the size of the query and gallery datasets respectively, therefore  $N_Q(N_G - 1)$  corresponds to all possible configurations of the event described before.

Having defined a global threshold  $\epsilon$ , the match between  $q_i$  and  $g_j$  is said to be  $\epsilon$ -meaningful if  $NFA(q_i, g_j) \leq \epsilon$ , in this case the face images are validated as corresponding to the same person. From these definitions is easy to prove that the expected NFA when testing all the possible combinations between the people in the query and gallery datasets is smaller than  $\epsilon$ .

The definition of a threshold  $\epsilon$  in the *a-contrario* framework, instead of thresholding the distance, has several advantages. First, it represents an intuitive indicator of the expected number of false alarms and therefore allows to control the performance of the system in advance. Second, this threshold is automatically adapted to the database as, for each person in the query dataset, it is applied over the trained probability function and therefore does not need to be adjusted each time the database is changed.


### 3 EXPERIMENTS

In this section we first present the techniques used in the preprocessing and feature extraction stages. Then, we provide details of the used databases and conducted experiments.


#### 3.1 Developed Framework

##### 3.1.1 Preprocessing Module

The preprocessing module automatically validates the face in an input image, discards all unnecessary information, and resizes the resultant image to a predefined size with fixed eye positions. This module is very important in any automatic face recognition system since this is the only stage where an input image can be rejected, e.g., if there isn't a face in the image. It is also very important because the quality in the registration of the face affects directly the performance of the system when local features are used. For this purpose we use both the manually marked eyes positions and a version of the technique Active Shape Models called STASM (Milborrow and Nicolls, 2008). This software takes as an entry a passport style image, validates a face in it, and then returns the position of seventy two pre-determined facial landmarks, from which two corresponds to the eyes positions. The location of these seventy two points might seem as an exaggeration since only two of these are finally used for the registration of the face. But this strategy allows to obtain better results finding eyes positions, even if they are closed, gaining in robustness. The STASM technique has been widely used in the face



(a) Application of the classic LBP operator.



(b) The circular (8,2), (8,1), and (16,2) neighborhoods in the ELBP approach.



(c) Result of applying the ELBP operator.

Figure 1: LBP and ELBP strategies.

recognition field, its implementation provides accurate results when marking points of interest in a face (a comparison of this method and other implementations of the ASM technique can be found in (Marsico et al., 2013)).

### 3.1.2 Feature Extraction

The feature extraction module receives a normalized image and computes a set of features (a vector) which ideally would uniquely represent the person in the input image.

We used the method presented on (Ahonen et al., 2006) for its simplicity, this technique uses Extended Local Binary Patterns (*ELBP*) to extract the so-called “micro-patterns”. This sets of features has been successfully used in texture and structures description, especially in face recognition (an extensive survey can be found in (Huang et al., 2011)). In the original *LBP* technique, first an operator is applied to every pixel in the image. This operator compares the gray values of each pixel against its neighbors and assign a binary label as shown in Figure 1(a). In the *ELBP* strategy this procedure is extended in order to consider more than only the immediate neighbors by taking the gray values in points uniformly distributed on a circle as shown in Figure 1(b). The pixel values are bilinearly interpolated whenever the sampling point is not in the center of a pixel.

The result of applying the *ELBP* operator is a new image of features as shown in Figure 1(c). This image is divided in regions from which feature histograms

are extracted. These histograms are finally concatenated, this way both local and global information is preserved. The micro-patterns hold the local information and the order in the concatenation of the local histograms assure that when comparing two faces the same regions of the faces will be taken into consideration for the comparison. The feature vector has a size  $s$  that can be computed as

$$s = N_b N_p \tag{5}$$

Where  $N_b$  represents the quantity of bins in the local histograms and  $N_p$  the total number of patches used in the division of the image.

### 3.1.3 Matching

In order to measure the dissimilarity between feature vectors we used the Chi-square distance between histograms corresponding to each one of the patches, this distance has been widely used when working with histograms as feature vectors obtained with operators as *ELBP*. Given a face  $q_i$  in the input image and a face  $g_j$  in the gallery dataset, the distance between them can be computed as the sum of the distances measured between patches as shown in Equation 6.

$$\begin{aligned} D(q_i, g_j) &= \sum_{n=1}^{N_p} \chi_n^2(q_i, g_j) \\ &= \frac{1}{2} \sum_{n=1}^{N_p} \sum_{m=1}^{N_b} \frac{(q_{i,n,m} - g_{j,n,m})^2}{(q_{i,n,m} + g_{j,n,m})} \end{aligned} \tag{6}$$

After the distance between samples is computed, we perform the matching by using two different techniques. The first matching strategy is based in the nearest neighborhood (NN) technique, in which an identity is validated just by applying a threshold in the distance between two faces. This strategy has been widely used in the face recognition area, and therefore we used as a benchmark. The second strategy, which is the novelty presented in this work is based in the use of *a-contrario* model that works as described in Section 2.1.

## 3.2 Databases

In this work we use two different databases to evaluate the performance of the developed face recognition system: *FERET* and *DNIC*.

### 3.2.1 Feret

The *FERET* database was created as part of the program Face Recognition Technology carried on in the years 1993 – 1997 (Phillips et al., 1998). This

database has become very popular in the area of face recognition and is commonly used as a benchmark. It contains a gallery  $f_a$  containing 994 people and three standards test sets:  $f_b$ ,  $dup_1$  and  $dup_2$ . The set  $f_b$  includes images taken on the same day of the ones in the gallery with only differences in expressions, the eyes position coordinates of 849 people of this dataset are provided. The set  $dup_1$  contains images of 736 persons and includes the aging effect when compared to the samples in the gallery. The set  $dup_2$  is a subset of  $dup_1$  that contains 228 persons with images taken, at least, 540 days later than the images in the gallery.

### 3.2.2 DNIC

*DNIC* (Dirección Nacional de Identificación Civil) is the Uruguayan government organization responsible for the emission of ID cards and passports. It has a civil identification system, with a database containing more than three millions identities<sup>2</sup>. This database has the special feature of being used in a real application. Its images are acquired under not completely controlled and different environments, therefore, some special characteristics are to be expected as changes in illumination and pose variations. Some others factors are controlled like the lack of face accessories (hats, lenses, scarfs, etc) and neutral face expressions. Summarising, one could say that this is a database of semi-controlled conditions. As the IDs are renewed periodically<sup>3</sup>, this database allows us to study the impact of the aging factor in the performance of the system.

### 3.3 Experiments

In this work we perform several test of the proposed system using the databases presented above. In all the performed experiments the system is evaluated both using the NN approach and the a-contrario proposed technique. In each case, the *Verification Rate (VR)* is plotted against the *False Acceptance Rate (FAR)*. This two indicators are commonly used when evaluating the performance of a biometric system used for identities verification. In order to present how these indicators are computed the following functions are introduced.

$$\mathbb{1}_D^\tau(q_i, g_j) = \begin{cases} 1 & \text{if } D(q_i, g_j) \leq \tau \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

<sup>2</sup>Because of the local laws that protect privacy, sample images of this database can not be shown in this article.

<sup>3</sup>Normally, the IDs are renewed every 10 years, but in special cases (for example, in the case of stolen or lost documents) the IDs are renewed with a smaller time difference.

$$\mathbb{1}_{id}(q_i, g_j) = \begin{cases} 1 & \text{if } id(q_i) = id(g_j) \\ 0 & \text{if } id(q_i) \neq id(g_j) \end{cases} \quad (8)$$

The indicator function in Equation 7 equals 1 when the distance between the faces  $q_i$  and  $g_j$  is below a predefined threshold  $\tau^4$  and therefore are declared by the system as being faces of the same person. The function defined in Equation 8 equals 1 when the compared faces correspond to the same person (their associated identity is the same). The *VR* and *FAR* indicators are computed according to the equations 9 and 10 respectively.

$$VR(\tau) = \frac{\sum_{i=1}^{N_Q} \sum_{j=1}^{N_G} \mathbb{1}_{id}(q_i, g_j) \times \mathbb{1}_D^\tau(q_i, g_j)}{N_Q} \quad (9)$$

$$FAR(\tau) = \frac{\sum_{i=1}^{N_Q} \sum_{j=1}^{N_G} (1 - \mathbb{1}_{id}(q_i, g_j)) \times \mathbb{1}_D^\tau(q_i, g_j)}{N_Q \times (N_G - 1)} \quad (10)$$

It could be said that the *VR* indicator represents the rate of persons who, claiming their true identity, are correctly validated by the system. On the other hand, the value *FAR* indicates the rate of people who, claiming a false identity, can fool the system that validates them as genuine users.

**Experiment 1.** In this experiment the *FERET* set  $f_b$  is used to evaluate the system both in the manual and automatic configurations, in the manual configuration the provided eyes coordinates are used while in the automatic mode the eyes position coordinates are found using the presented *STASM* technique. This first test has two different goals. First, since the images of this set have little variation with respect to the images in the gallery, it allows us to study the impact of the positioning of the eyes in the performance of the system. Second, it validates the performance of the a-contrario proposed technique in a well known database that is usually used as a benchmark.

**Experiment 2.** In Experiments 2 and 3 we perform the evaluation with several subset from the *DNIC* database. As stated before, this database allows to evaluate the system working with a database of semi-controlled conditions used by the *DNIC* in a production environment. In experiment two, we test the developed system using a database of 5000 people, this permits us to evaluate the system working with a mid-size database.

<sup>4</sup>The threshold  $\tau$  is the one applied directly over the distance between samples or in the probability computed using the proposed *a-contrario* technique.

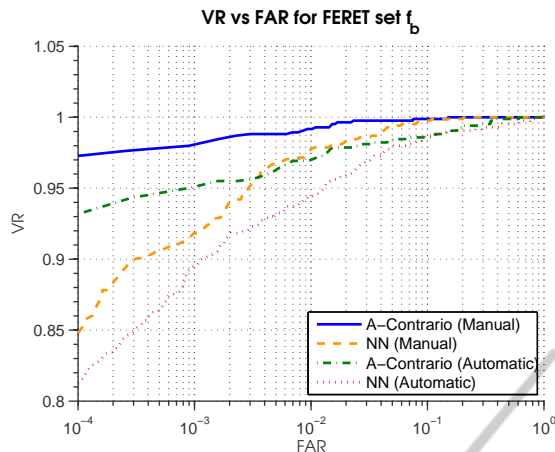


Figure 2: Experiment 1: VR vs FAR for automatic and semi-automatic configurations, FERET  $f_b$  dataset.

**Experiment 3.** In this experiment we analyze the impact of aging in the face recognition process by using datasets composed of 200 people with a time difference between the query images and the samples in the gallery of 2, 3, 4, 5 and 9 years. In all these tests we evaluate the system using both the proposed *a-contrario* framework and the NN approach.

In all the experiments we used normalized images of size  $128 \times 128$  pixels with fixed eyes positions. We used the ELBP technique considering 8 neighbors in a circle of radius 3 pixels, using 256 bins in each local histogram and dividing the image with a  $9 \times 9$  grid.

## 4 RESULTS

The results of Experiment 1 are shown in Figure 2. Clearly the positioning of the eyes affects the overall performance of the system in both the NN and proposed *a-contrario* approaches. The performance is worse when the automatic eyes positions finder technique is used instead of providing the manually marked eyes coordinates. This behaviour is as expected considering that the features are extracted by regions on the face. Therefore, any difference in the registration of a person's face between the images in the gallery and query datasets could result in the comparison of features that are extracted from different face regions negatively impacting the performance.

From this first experiment is clear that the *a-contrario* proposed technique outperforms the classic NN approach in both the manual and automatic configurations of the system. The performance when using both strategies is slightly different for the higher values of the FAR indicator as shown in Figure 2. On the other hand, when we pursue lower values of the

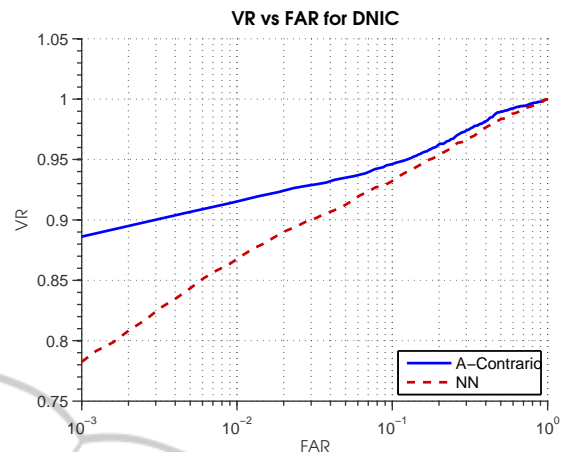


Figure 3: Experiment 2: VR vs FAR for DNIC dataset, subset containing 5000 people.

FAR index, the *a-contrario* technique remains more robust than the NN approach leading to significantly higher values of the verification rate indicator. This is very important, as in these applications usually is convenient to work with a very low FAR value while not affecting the system's genuine users. In the NN strategy the decision threshold needs to be more restrictive in order to achieve these FAR values leading to a degradation in the verification rate. The *a-contrario* strategy is more robust to this change as its threshold is trained for each person taking many samples of distances between different people's faces.

The results of Experiment 2 are shown in Figure 3, it can be seen once more that the proposed *a-contrario* technique outperforms the NN approach. In this case the difference in the performance maintains for all the range of FAR values. This behaviour is to be expected as, in this case, there are many more samples to train the probability density function (pdf) used in the *a-contrario* technique. A better estimation of the pdf allows for a better adjustment of the threshold when comparing two persons increasing the discrimination capability.

We present in Table 1 the relationship between the selected  $\epsilon$  threshold for the application of the *a-contrario* technique and the obtained VR and FAR values when testing the system with the DNIC big dataset. As stated before, the setting of the  $\epsilon$  threshold allows to control the expected number of false alarms, this represents a great advantage as it allows to control the performance of the system in advance. From the presented data, is clear that the FAR indicator that measures the total number of false alarms remains bounded by the chosen epsilon value. This shows empirically that the setting of an arbitrary threshold  $\epsilon$  allows to control the performance of the system be-

Table 1: Experiment 2: FAR in the *a-contrario* approach versus  $\epsilon$  threshold.

| $\epsilon$ | FAR   | VR    |
|------------|-------|-------|
| 0.005      | 0.004 | 0.905 |
| 0.010      | 0.009 | 0.913 |
| 0.050      | 0.047 | 0.934 |
| 0.100      | 0.097 | 0.945 |
| 0.500      | 0.495 | 0.989 |

forehand.

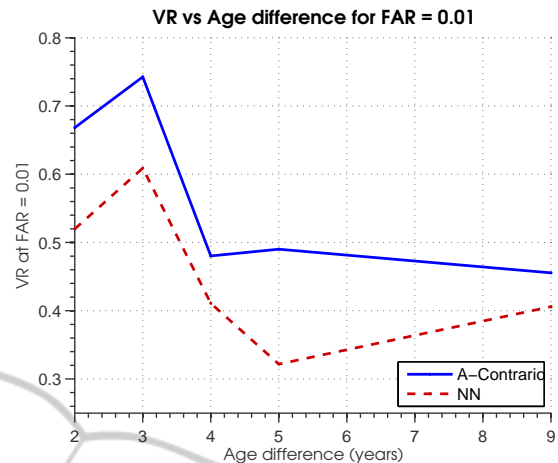
Finally, we analyze how the aging process impacts in the performance of the developed system using both the *a-contrario* technique and NN approach. The obtained results are shown in Figure 4. Clearly, the age gap between the images in the gallery and query dataset highly degrades the performance of the system. There is an improvement in system performance when there is a gap of three years between images with respect to the case where the separation is of two years. This behavior is not as expected as it is normal that the performance worsens as the time difference between images increases. This result can only be explained by the nature of the data, it is important to remark that this database is composed by semi-controlled conditions (lighting, pose, expressions), these conditions could affect the system, in some cases, more than the time difference between images.

While the *a-contrario* technique produces better results than the NN-based strategy, both of them start to fail as the age difference between images increase. This behaviour can be explained by looking at the feature extraction stage. The selected features are very simple to implement and fast to compute when processing images but not robust to the changes produced in a face as a result of the aging process. The *a-contrario* technique does not solve this problem as it relies in the extracted features used when describing a face.

## 5 CONCLUSIONS

The face recognition field is rapidly growing and there is a huge need for new techniques and algorithms that could enable to identify a subject or verify an identity even in very uncontrolled conditions as those presented in pictures taken from surveillance cameras.

In this work we introduce a matching strategy based in the *a-contrario* models that were applied widely in the pattern recognition area but not in the particular problem of face recognition. This technique does not present any constraints for its application and is well suited even in the case there is

Figure 4: Experiment 3: VR vs Age difference for  $FAR = 0.01$  in DNIC database.

only one sample per person in the gallery dataset. It is also independent from the extracted features as soon as a dissimilarity measure between samples is defined, this allows to adapt the matching stage to different feature extraction scenarios. The conducted experiments show that this algorithm outperforms the classic NN approach widely used in the face recognition field as it is adjusted to the input image and the database. It is also shown that it allows to predict with high precision the total number of false alarms by setting a threshold  $\epsilon$ , thus enabling to control in advance the performance of the system.

Several points remains as future work. For instance, we are planning to compare the performance of the proposed technique against others matching algorithms taking in consideration performance, scalability and speed. In particular we are interested in the scalability factor in order to validate how well is this technique suited to the verification process in a citizen database as the one from the DNIC. We are also planning to test the proposed technique with different sets of extracted features in order to know if any advantage can be obtained when using some particular features that can, for example, overcome the aging issue.

## ACKNOWLEDGEMENTS

The authors would like to thank Rafael Grompone and Pablo Musé for their valuable contributions, the DNIC agency for its collaboration and for granting us the permission to access their valuable data. This work has been supported by an investigation grant provided by the ANII (Uruguay Agency of Investigation and Innovation).

## REFERENCES

- Ahonen, T., Hadid, A., and Pietikäinen, M. (2006). Face description with local binary patterns: application to face recognition. *IEEE transactions on pattern analysis and machine intelligence*, 28(12):2037–41.
- Desolneux, a., Moisan, L., and Morel, J. (2003a). A grouping principle and four applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(4):508–513.
- Desolneux, A., Moisan, L., and Morel, J. (2003b). Maximal meaningful events and applications to image analysis. *The Annals of Statistics*, 31(6):1822–1851.
- Gioi, R. V. and Jakubowicz, J. (2010). LSD: A fast line segment detector with a false detection control. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 32:722–732.
- Huang, D., Shan, C., Ardabilian, M., Yunhong, W., and Liming, C. (2011). Local Binary Patterns and Its Application to Facial Image Analysis: A Survey. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 41(6):765–781.
- Kan, M., Xu, D., Shan, S., Li, W., and Chen, X. (2013). Learning Prototype Hyperplanes for Face Verification in the Wild. *Image Processing, IEEE Transactions on*, 22:3310–3316.
- Lanitis, A., Taylor, C., and Cootes, T. (2002). Toward automatic simulation of aging effects on face images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 24(4):442–455.
- Ling, H. and Soatto, S. (2010). Face verification across age progression using discriminative methods. *Information Forensics and Security, IEEE Transactions on*, 5(1):82–91.
- Marsico, M. D., Nappi, M., Riccio, D., and Wechsler, H. (2013). Robust Face Recognition for Uncontrolled Pose and Illumination Changes. *IEEE Transactions on Systems, Man, and Cybernetics*, 43(1):149–163.
- Masi, I., Lisanti, G., Bagdanov, A., Pala, P., and Bimbo, A. D. (2013). Using 3D Models to Recognize 2D Faces in the Wild. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2013 IEEE Conference on*, pages 775–780.
- Milborrow, S. and Nicolls, F. (2008). Locating facial features with an extended active shape model. *Computer Vision ECCV 2008*, pages 504 – 513.
- Mottalli, M., Tepper, M., and Mejail, M. (2010). *A contrario detection of false matches in iris recognition*, volume 6419. Springer Berlin Heidelberg.
- Musé, P., Sur, F., Cao, F., Gousseau, Y., and Morel, J.-M. (2006). An A Contrario Decision Method for Shape Element Recognition. *International Journal of Computer Vision*, 69(3):295–315.
- Park, U., Tong, Y., and Jain, A. K. (2010). Age-invariant face recognition. *IEEE transactions on pattern analysis and machine intelligence*, 32(5):947–54.
- Phillips, P. J., Wechsler, H., Huang, J., and Rauss, P. J. (1998). The FERET database and evaluation procedure for face-recognition algorithms. *Image and Vision Computing*, 16(5):295–306.
- Rabin, J., Delon, J., and Gousseau, Y. (2008). A contrario matching of SIFT-like descriptors. *2008 19th International Conference on Pattern Recognition*, pages 1–4.
- Zhang, B., Gao, Y., Zhao, S., and Liu, J. (2010). Local derivative pattern versus local binary pattern: face recognition with high-order local pattern descriptor. *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society*, 19(2):533–44.
- Zhang, B., Shan, S., Chen, X., and Gao, W. (2007). Histogram of Gabor phase patterns (HGPP): a novel object representation approach for face recognition. *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society*, 16(1):57–68.
- Zhang, W., Shan, S., Gao, W., Chen, X., and Zhang, H. (2005). Local Gabor binary pattern histogram sequence (LGBPHS): a novel non-statistical model for face representation and recognition. *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, pages 786–791 Vol. 1.
- Zhu, X. and Ramanan, D. (2012). Face detection, pose estimation, and landmark localization in the wild. *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2879–2886.