Attacking Iris Recognition: An Efficient Hill-Climbing Technique

Christian Rathgeb and Andreas Uhl
Department of Computer Sciences, University of Salzburg, 5020 Salzburg, Austria
{*crathgeb, uhl*}*@cosy.sbg.ac.at*

Abstract

In this paper we propose a modified hill-climbing attack to iris biometric systems. Applying our technique we are able to effectively gain access to iris biometric systems at very low effort. Furthermore, we demonstrate that reconstructing approximations of original iris images is highly non-trivial.

1 Introduction

In the past years several approaches to image reconstruction from biometric templates have been proposed [2]. Ratha *et al.* have summarized existing points of attack in biometric systems [6]. Depending on these points of attack, biometric modalities, and structures of templates, several attempts to image reconstruction have been proposed out of which hill-climbing (HC) has proven to be one of the most effective.

The first to come up with HC attacks applied to fingerprints was Soutar [7]. The key idea behind HC is to consecutively modify an input image which is presented to the biometric recognition system in order to access a distinct account. The attacker observes the matching score returned by the system at the time of each authentication and retains changes in the input image which increase the matching score. The process of changing the input image is repeated until no significant improvement in the matching score is observed. In order to perform a HC attack to any biometric system an attacker must have access to the internal match score (MS) calculated by the system. HC attacks assume attackers are able to observe any communication channels in a biometric system. That is, HC may be performed by modifying an input image presented to the biometric sensor, as well as, for example, by manipulating an extracted feature vector. Adler [1] successfully applied a HC attack to a face recognition system. Starting from an initial face image fractions of several eigenfaces of other users are added. Approximations of the target image which contained most distinct face features were presented in experimental results. Until now, several approaches have been proposed implementing hill-climbing attacks for different biometric modalities including behavioral biometrics as well. Even for biometric cryptosystems and systems which employ quantized MSs HC attacks have been proposed. Due to a lack of space we are not able to discuss all of these (for further details see [2]). We apply a modified HC attacks to iris recognition and consider the aspect of gaining access to iris recognition systems as well as iris texture reconstruction by applying our HC strategy. We show that image reconstruction is highly non-trivial applying HC to iris biometrics. To our knowledge there is no other work investigating HC attacks in iris recognition.

2 Hill Climbing in Iris Recognition

We propose a modified HC attack which we apply to iris recognition. Since the applied feature extraction is rather sensitive to small changes in the extracted iris texture we decided to apply pixel-wise modifications to preprocessed iris textures and observe improvements in the MS. HC attacks require that attackers are able to tap communication channels in a biometric system and manipulate transfered data, thus, we modify preprocessed iris textures, which represents a realistic scenario. Since less redundant data remains after preprocessing, manipulation of preprocessed iris textures is expected to be most effective. For each HC attack we either use an initial texture where each pixel value is 128 (Fig. 1 (b)) or an initial eigeniris texture averaged out of five randomly chosen textures (Fig. 1 (c)).

2.1 Target System

For this purpose we employ our own implementation of the algorithm described by Masek [5] in the feature extraction process. In the algorithm of Masek, which is a simplified implementation of Daugman's algorithm, the upper 512×50 pixel of the preprocessed iris tex-

1051-4651/10 $26.00 © 2010 IEEE
DOI 10.1109/ICPR.2010.303

tures (e.g. Fig. 1 (a)) are examined and mean values of blocks of 1×5 pixel are processed. Ten 1-D intensity signals are analyzed where complex values in the transform-domain of the Log-Gabor transformation are encoded with 2 bits per 1×5 pixel block extracting an iris-code of $512 \times 10 \times 2 = 10240$ bits. Since feature extraction is rather similar to that proposed by Daugman, which is implemented in most commercial iris recognition systems, we found applying HC to this algorithm reveals representative results. In order to provide a zero FMR the algorithm rejects iris images which produce a Hamming distance (HD) higher than ~ 42 % according to the claimed identity. That is, a HC attack aims at generating an input image which produces a MS $> 58\%$ (MS=1-HD) for a distinct target image.

2.2 Proposed Hill Climbing

Starting from one of the initial textures, pixels are modified step by step. We start modifying the texture in the top left. If the buttom right is reached we start over again. The whole process is summarized as follows:

1. The value of a pixel is increased by a predefined constant c and authentication is performed.
2. The resulting MS is observed and if it has increased the modification is retained.
3. In case the MS decreases or remains the same the pixel value is decreased by c and the resulting MS is observed.
4. By analogy, if a decrementation of the pixel value leads to a higher MS the modification is retained.

The whole process is repeated for pixel values until the modified texture is accepted by the system or no significant improvements are observed. For the target image in Fig. 1 (a), examples for applying the proposed HC until no significant improvement in the MSs are obtained for both initial images are illustrated in Fig. 2 (a) and Fig. 2 (b), respectively (notice that the algorithm of Masek only processes the upper 512×50 pixels). Fig. 2 (c) shows the average texture shown in Fig. 1 (b) which was modified until successfull authentication for the target texture of Fig. 1 (a) was obtained (for visibility modifications were applied to the center rows of the texture). For each attack c was set to 10.

2.3 Block Detection

Modified textures only reveal high MS with respect to the described feature extraction of Masek. If modified textures are presented to other biometric systems

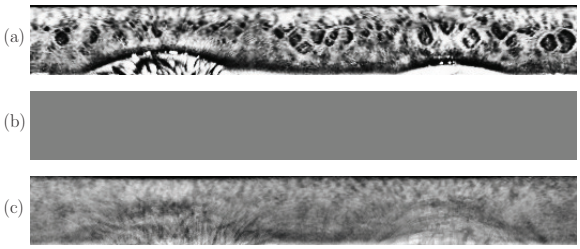

Figure 1. (a) Preprocessed iris texture (b) Average texture (c) Eigeniris texture.

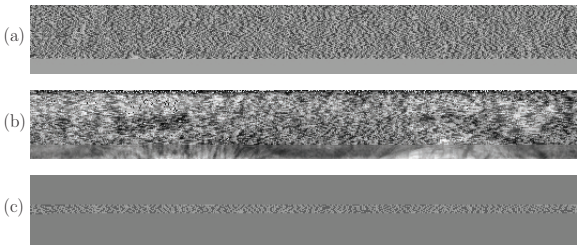

Figure 2. (a) Average texture (MS: 94.95%) (b) Eigeniris texture (MS: 92.88%) (c) Average texture (MS: 58.00%).

high match scores will only be obtained if the feature extraction of the according system would be rather similar to that of Masek. For example, if we match the original texture of Fig. 1 (a) and the generated texture of Fig. 2 (a) the algorithm of Masek reveals a MS of 94.95%. If we match both textures with our own implementation of the algorithm of Ma (described in [4]) a MS of 80.54% is obtained. Since, the algorithm of Ma is very similar to that of Masek (1-D wavelet transform is applied to ten 1-D intensity signals of length 512 at two subbands) obtained MSs turn out to be high. In contrast, if we match both textures with our implementation of the algorithm of Ko [3] which is based on a completely different feature extraction (changes in cumulative sums of grayscale pixel-blocks are observed) a MS of 45.73% is obtained. Based on 50 generated iris textures we have found that these observations hold in general as can be seen in Tab. 2. We conclude that the structure of generated iris textures highly depends on the algorithm from which MSs are observed.Besides that, the proposed algorithm does not produce useful reconstructions of the target iris texture (see Fig. 2). To accelerate modifications and to improve approximations of target iris textures we aim at genericly identifying the block-dimension on which the feature extraction operates. Since most iris recognition algorithms tend to average certain pixel blocks the identification of block dimensions would lead to faster modifications. If the block dimension is known several pixels can be modified during one iteration while accuracy is kept. Hence, a faster attack is possible (see Sect. 3). Furthermore,

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Row	9	6	5	4	0	8	3	4	1	2	6	2	2	0	0	6	5	3	4	0
Column	4	6	5	2	2	3	2	2	1	1	0	1	3	7	2	5	0	3	2	4

Table 1. Number of modifications in rows and columns for a 20 $\times$ 20 pixel block.

Algorithm	Mean Match Score	Standard Deviation
Masek	95.0012	0.0013
Ma	80.3572	4.4269
Ko	47.7788	2.4781

Table 2. Means and standard deviations of MSs for 50 generated textures applying our HC to the algorithm of Masek.

modifications based on pixel blocks are expected to reveal better approximations of target images. If pixel-wise modifications are performed, the resulting pixel blocks do not necessarily have to look alike those of target images while block-wise modifications might yield better approximations (see Sect. 2.4).

For the purpose of block detection we apply a pixel-wise step-by-step HC to a 20×20 pixel block centered in preprocessed iris textures where the above algorithm is applied to each pixel once. We chose the center of the texture since pixels in this area contribute to resulting iris-codes with the utmost probability. Since we apply pixel-wise HC and pixel blocks are averaged in generic feature extraction algorithms we can make the following assumptions: (1) If the modification of a pixel value increases the MS this applies to all pixels within a $x \times y$ block (averaging is applied during feature extraction). (2) The upper left area of each contained pixel block modified with high propability while the bottom right area (each $x \times y$ block is modified in the top left at first and in the bottom right at last).

Therefore, the first row of a block will be substantially modified while modifications will reduce until the next pixel block is reached. The same holds for columns. Thus, by counting modifications according to rows and columns we obtain the block dimension of the applied feature extraction. Hence, no a priori knowledge about the feature extraction algorithm is necessary. A list of the number of modifications in a 20×20 pixel block in the center of the sample image of Fig. 1 (a) is listed in Tab. 1. Even if this procedure is applied only to a single texture, a cell height of 5 pixels is obtained since modifications decrease until every fifth row. On the other hand, a cell width of 1 pixel is detected because no decrease of modifications is obtained. The knowledge of the applied block dimension during feature extraction highly accelerates HC. Additionally, we will demonstrate that images modified based on pixel blocks reveal better approximations of the target image.

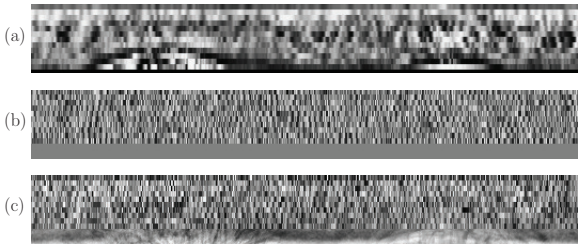

Figure 3. (a) Averaged target image (MS: 100.00%) (b) Average texture (MS: 97.98%) (c) Eigeniris texture (MS: 93.89%).

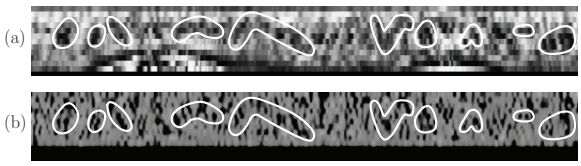

Figure 4. (a) Averaged iris texture of the target image (b) Texture of Fig. 3 (b) prepared with contrast enhancement and a Gaussian filter.

2.4 Image Reconstruction

Applying the above algorithm a block dimension of 1×5 pixels is detected. Hence, the applied iris recognition algorithm processes an approximation of the preprocessed iris texture. A modified input image which leads to a perfect MS does not necessarily look like the target image. For example, Fig. 3 (a) illustrates a texture in which each 1×5 pixel block is averaged leading to a perfect MS. As can be seen in Fig. 3 (a) extracting coarse texture structures of target images is the most we can expect. Based on the determined block dimension we apply block-based HC. Examples for modified images which have been generated for the target texture starting from both initial textures are shown in Fig. 3 (b) and 3 (c). Hence, the grayscale value of each 1×5 block were summed up, modified according to the proposed HC ($c = 10$) and the resulting average value is assigned to each pixel of a processed block. Again, generated textures which produce high MSs do not reveal usefull information about the structure of target images.

Thus, we applied contrast enhancement as well as a Gaussian filter to the calculated textures. We found best structures were extracted using the average texture as initial image. An example for the extraction of coarse texture features is illustrated in Fig. 4 (b) where features are marked with splinegons, just like in the averaged target image. Obviously a coarse texture reconstruction is possible, as pointed out in [1]. However, in iris biometrics the extraction of coarse texture features does not suffices in any case to identify the target image.

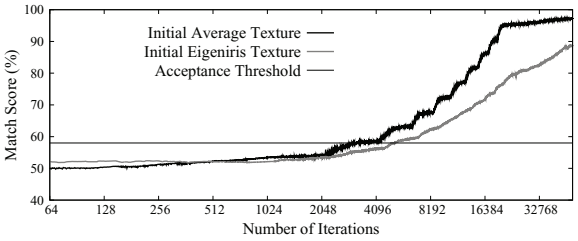

Figure 5. Score prog. (pixel-wise HC).

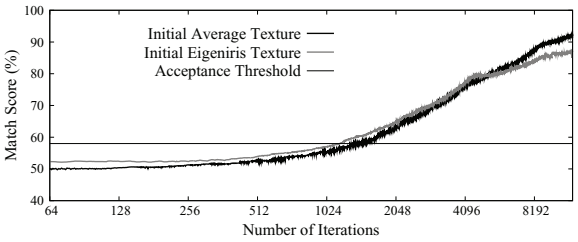

Figure 6. Score prog. (block-wise HC).

3 Computational Performance

In experiments we focus on score progressions of our HC method, that is, the number of iterations which are necessary to achieve a high MS are measured. Experiments are performed using the CASIAv3-Interval iris database, a well-established set of iris images, where iris textures of 512×64 pixels are extracted in the preprocessing step. Starting from both initial textures the proposed HC strategy is applied until no significant improvements in the MS are observed.

For illustration the score progressions for the single target image of Fig. 1 (a) are plotted in Fig. 5 starting from an average texture and an eigeniris texture. The eigeniris texture reveals a slightly higher initial MS. However, as the number of iterations grows the modified average texture generates higher MSs than the modified eigeniris texture. This is because starting from an average texture allows a more directed modification of the resulting iris-codes while the eigeniris texture already generates a distinct iris-code which has to be modified. As shown in Fig. 5, for the target image successful authentication is yielded at a number of 3082 and 5126 iterations starting from an average texture and an eigeniris texture, respectively. For a MS greater than 90% a number of 17969 and 59749 iterations are required for the according initial textures. One important feature of our attack is that modifications are computationally cheap. A single modification only involves an addition or a subtraction to a pixel value, in contrast to more complex HC strategies (e.g. [1]). The average number of iterations necessary to achieve according MSs applying pixel-wise and block-wise HC are summarized in Tab. 3. Results were obtained per-

	Initial Texture / Block Size			
	Eigeniris 1×1	Average 1×1	Eigeniris 1×5	Average 1×5
∅ Iterations until $>58\%$	$\sim$5100	$\sim$3200	$\sim$1500	$\sim$1400
∅ Iterations until $>90\%$	$\sim$58000	$\sim$20300	$\sim$14500	$\sim$9800

Table 3. Average numbers of iterations which yield MSs $>$ 58% and MSs $>$ 90%.

forming over a hundred HC attacks to randomly chosen target textures. If the block dimension is known to an attacker, HC is highly accelerated. For the respective target image now only 1383 and 1168 iterations are necessary in order to generate an image which yields successful authentication, as shown in Fig. 6. Again, starting from an eigeniris texture reveals better initial results while after about 5000 iterations the modified average texture generates higher MS. MSs higher than 90% were generated after a number of 9741 and 14697 iterations. The results shown in Tab. 3 emphasize how the proposed block-detection accelerates HC attacks.

4 Conclusion

We presented an efficient HC attack which we applied to iris biometrics to achieve successfull authentication in an iris recognition system. By extracting additional information about the feature extraction of the iris recognition algorithm proposed HC is highly accelerated. Experiments emphasize the high performance of the attack since only a small number of iterations which comprise only low-cost modifications are necessary to generated images which achieve high MSs.

References

[1] A. Adler. Sample images can be independently restored from face recognition templates. *Proceedings of the Canadian Conference on Electrical and Computer Engineering*, 2:1163–1166, 2003.

[2] A. K. Jain, K. Nandakumar, and A. Nagar. Biometric template security. *EURASIP J. Adv. Signal Process*, 2008:1–17, 2008.

[3] J.-G. Ko, Y.-H. Gil, and J.-H. Yoo. Iris Recognition using Cumulative SUM based Change Analysis. *Intelligent Signal Processing and Communications, 2006. ISPACS'06*, pages 275–278, 2006.

[4] L. Ma, T. Tan, Y. Wang, and D. Zhang. Efficient Iris Recogntion by Characterizing Key Local Variations. *IEEE Transactions on Image Processing*, 13(6):739–750, 2004.

[5] L. Masek. Recognition of Human Iris Patterns for Biometric Identification, Master's thesis, University of Western Australia, 2003.

[6] N. K. Ratha, J. H. Connell, and R. M. Bolle. An analysis of minutiae matching strength. In *Audio- and Video-Based Biometric Person Authentication*, volume 2091, pages 223–228, 2001.

[7] C. Soutar. Biometric system security, white paper, bioscrypt, 1999.