Face Recognition in the Virtual World: Recognizing Avatar Faces

Roman V. Yampolskiy[1], Brendan Klare[2], and Anil K. Jain[2]
[1]Department of Computer Engineering and Computer Science, University of Louisville, KY, USA 40292; [2]Department of Computer Science, Michigan State University, MI, USA 48824;

ABSTRACT

Criminal activity in virtual worlds is becoming a major problem for law enforcement agencies. Forensic investigators are becoming interested in being able to accurately and automatically track people in virtual communities. In this paper a set of algorithms capable of verification and recognition of avatar faces with high degree of accuracy are described. Results of experiments aimed at within-virtual-world avatar authentication and inter-reality-based scenarios of tracking a person between real and virtual worlds are reported. In the FERET-to-Avatar face dataset, where an avatar face was generated from every photo in the FERET database, a COTS FR algorithm achieved a near perfect 99.58% accuracy on 725 subjects. On a dataset of avatars from Second Life, the proposed avatar-to-avatar matching algorithm (which uses a fusion of local structural and appearance descriptors) achieved average true accept rates of (i) 96.33% using manual eye detection, and (ii) 86.5% in a fully automated mode at a false accept rate of 1.0%. A combination of the proposed face matcher and a state-of-the art commercial matcher (FaceVACS) resulted in further improvement on the inter-reality-based scenario.

Keywords: Virtual world, Second Life, avatar, face recognition, local image features

1. INTRODUCTION

"The bomb hit the ABC's headquarters, destroying everything except one digital transmission tower. The force of the blast left Aunty's site a cratered mess. Just weeks before, a group of terrorists flew a helicopter into the Nissan building, creating an inferno that left two dead. Then a group of armed militants forced their way into an American Apparel clothing store and shot several customers before planting a bomb outside a Reebok store. This terror campaign has left a trail of dead and injured, and caused hundreds of thousands of dollars' in damages. The terrorists belong to a militant group bent on overthrowing the government"[1]. Unfortunately the terrorists involved will not be prosecuted for their crimes as they have committed them away from the jurisdiction of the law enforcement agencies in the lawless virtual world known as the Second Life[1].

Quick investigation of the Second Life world shows that it is populated by numerous terrorist organizations, including international jihadist terrorist groups associated with al-Qaeda and local groups of radicals such as Second Life Liberation Army[1]. Similar to the 9/11 terrorists who practised flying planes on simulators in preparation for their deadly attack on civilian buildings and the Pentagon, security experts believe some of the people behind the Second Life terror campaign are real-world terrorists who are rehearsing for future strikes against non-virtual targets. Al-Qaeda and Jemaah Islamiah are well known for sending their recruits to training camps in other countries, but with the presence of US forces in those countries, new training options are being looked at by the leaders of terror networks, including the idea of training in the virtual worlds. Terrorists can train in virtual worlds such as Second Life in a simulated environment, using weapons that are identical to their real-world counterparts[1].

Virtual worlds are also extremely attractive for the run-of-the-mill criminals interested in conducting identity theft, fraud, tax evasion, illegal gambling and other traditional crimes. United States intelligence officials are very concerned over the unprecedented growth of online virtual communities populated by avatars which "offer the opportunity for religious/political extremists to recruit, rehearse, transfer money, and ultimately engage in information warfare or worse with impunity"[2]. The U.S. Congress held hearings about the potential use of virtual worlds for money laundering operations by Al-Qaida and other terrorist groups[3]. The United States government's growing concern are likely to make online virtual worlds the next frontier in the battle over the limits on the government's right to improve security via data collection and analysis and the surveillance of commercial computer systems[2]. As more people join virtual communities and create avatars, it will become more difficult to identify potential criminals and terrorist. "As in the real world, one of the central difficulties is establishing the identity of individuals" [2].

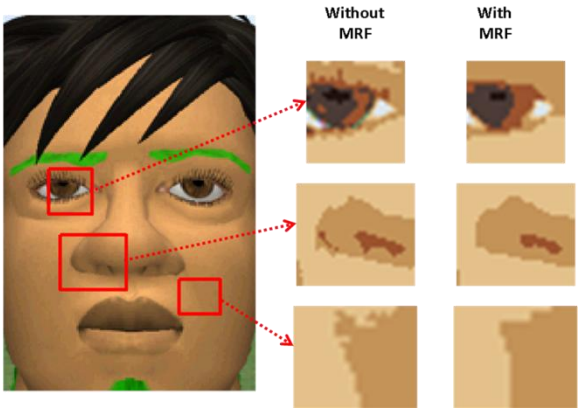

Figure 1. Difference between color segmentation with and without the proposed Markov random field model. Use of MRF results in a smooth segmentation.

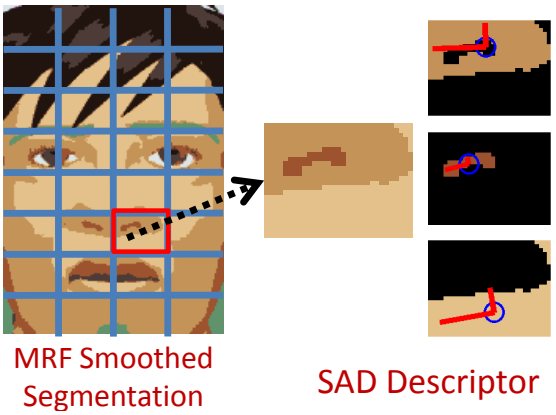

Figure 2. Illustration of the SAD descriptor. The spatial variation for each color label is computer at each image patch. Image above is the same image as Fig.1, but after reducing the color space to 8 levels and smoothing the labeling with the MRF.

Relationships between social, economic, and psychological status of game players and their respective avatars in the virtual environment[4] show that avatars for the most part resemble their "owners" rather than being completely virtual creations[5] and exhibit a high degree of permanence. This obviously does not hold true for all terrorists training for a real-world attack, but maybe true for some less sophisticated fanatics. With the convergence of the physical and the virtual worlds, the distinction between the two begins to fade and the need arises for security systems capable of working in the contexts of inter-reality and augmented reality. However, available face recognition systems have neither been designed nor evaluated on non-human agents that exhibit large visual and behavioral variations. In this paper we describe utilization of state of the art face recognition systems[6] and development of novel face recognition algorithms for face-based avatar authentication.

In the context of investigating criminal and terrorist activity in the virtual world, we see six (4 non-symmetric) scenarios requiring an automated face matching algorithm.

1. Matching a Human Face to an Avatar face

This capability is useful to connect a person's real identity to their virtual persona. It is increasingly common for a person to upload his photograph to serve as a prototype for a 3D avatar as well as to create drawings closely resembling him to serve as his online persona.

2. Matching face of one Avatar to another Avatar

This capability is useful for continuously tracking a virtual persona through cyberspace at different times and in different places.

3. Matching an Avatar's face from one virtual world to the same Avatar represented in a different virtual world(s)

A recent development in the world of virtual communities is the desire to interconnect different virtual worlds. One such world, called HiPiHi, is being created in China which makes it doubly hard to track down the identity of avatars [2].

4. Matching an Avatar sketch to an Avatar face.

Just like matching a forensic sketch based on the description provided by a victim of a crime or a witness, it is important to be able to match a sketch of a virtual criminal to an actual avatar responsible for the crime.

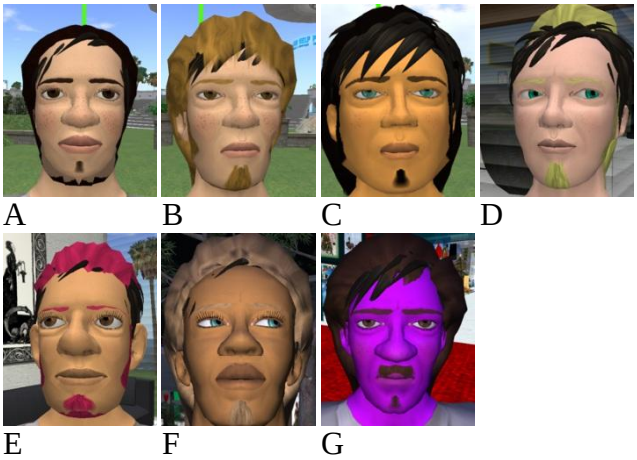

Figure 3. Examples of the different subjects in the Second Life avatar dataset. Each image corresponds to one of the different pose sets (A-G). In our matching experiments, the frontal image from group A was used as the gallery image. The remaining sets were all used as probe images.

Prior work on visual authentication of avatars is limited. Yanushkevich et al. suggested a need for development of "Robotic Biometrics"[7]. An approach for parameterized generation of avatar face datasets was reported in[8]. "Avatar facial biometric authentication" was published in French and it described results of an experiment in which an avatar face recognition system was constructed for recognizing avatars from MyWebFace.com. The system used wavelet transform for feature extraction and SVM for classification purposes[9]. Yampolskiy et al. provided application of biometric principles to avatar recognition and outlined future directions and potential applications[4]. Boukhris et al.[10] have presented an approach for applying face recognition to avatars as part of security framework for virtual worlds. Mohamed & Yampolskiy have applied wavelet transform with Local Binary Pattern (LBP) to recognize avatar faces in[11]. Yampolskiy et al. evaluate state of the art academic and commercial algorithms developed for human face recognition on the avatar face dataset and obtain recognition accuracy ranging from 53.57% to 79.9% on different systems[12].

2. AVATAR FACE RECOGNITION

Avatar and human faces have similarities and differences. Both have a consistent structure and location of facial components (i.e. the relationship among eyes, nose, etc.). These similarities suggest an avatar face recognition framework designed in the same manner as human FR systems. Avatar faces span a wider range of colors than human faces, and the colors provide strong discrimination between identities (see Fig. 3). In some virtual worlds avatars take non-human forms and so the proposed approach of avatar face recognition will not be applicable. Since Second Life is a dominant force in the space of virtual communities the proposed approach is nonetheless highly applicable.

We propose a framework for avatar face recognition (Scenario 2) that follows the same procedures as standard face recognition systems, consisting of three stages:

1. Face detection and image normalization
2. Face representation
3. Matching

2.1 Face Detection and Normalization

While traveling through a virtual world, real-time face detection will detect the presence of an avatar subject with a frontal to near-frontal face in the field of view. Similar to traditional face recognition, once an avatar face is detected it must be preprocessed by performing both geometric and color normalization in order to reduce variations caused by external parameters such as camera location and illumination.

While a number of approaches to face detection exist, the method of Viola and Jones[13] is often preferred due to the accuracy, speed, and availability through open source software[14]. Indeed, we found this method to work with similar effectiveness on avatar faces. We explored the use of (i) a Haar cascade trained on avatar faces, and (ii) the default frontal face Haar cascade packaged with OpenCV. Despite the default cascade being trained on human faces, we found similar accuracy using this cascade as we did with the avatar trained cascade. As a result, we decided to use the default cascade.

Geometric and Appearance Normalization

Geometric normalization is applied to reduce the effects of scale, rotation, and translations that occur within the detected face window. In order to accomplish this task, we rely on the location of the eyes[15].

The location of the eyes is estimated in the form of prior information from a training set. The position of the eyes is determined using the average relative position of the eyes within the face window from each training image. Once the center of the eyes is estimated, the face image is geometrically normalized by first performing planar rotation to set the angle between the eyes to 0 degrees. Next, the size of the image is scaled so that the distance between the two eyes is 75 pixels. Finally, the image is cropped to 140 x 210 pixels.

To compensate for changes in illumination, histogram equalization is performed to normalize the appearance. The avatar faces contain three color channels, but we merged these into one dimension for normalization because, certain color channels may have a low dynamic range for a particular face due to the abnormal face colors observed in avatars (green, red, etc.), and not due to a global illumination change.

2.2 Face Representation

In order to match two avatar faces, we represent the face in a metric space by first computing a set of local feature descriptors across the face region. Two separate feature descriptors are used to describe (i) the structure of the face, and (ii) the appearance properties of the face.

For computing the local descriptors, the normalized face image is divided into an ordered set of N overlapping square patches $P_i, i = 1 \ldots N$, each of size s_p x s_p, $s_p = 32$. For each patch P_i two feature vectors are extracted: one describing the appearance $(A_i \in \mathbb{R}^{d_a})$, and the other describing the structure $(S_i \in \mathbb{R}^{d_s})$. Computing features across a set of overlapping patches allows for salient descriptions at specific locations of the face that is robust to variations in geometric normalization.

Structural Representation

The structural representation of the face is intended to encode the face shape and morphology. We evaluated a number of local descriptors and found the local binary pattern (LBP) [16,17] representation to be the most effective. A special case of LBP, called the uniform LBP[17], is used in this work. Uniform LBP assigns any non-uniform binary number to the same value, where uniformity is defined based on the number, u, of transitions between the binary values 0 and 1. For $p = 8$ and $u = 2$ (where p is the number of sampling locations), the uniform LBP has 58 uniform binary numbers, and the 59th value is reserved for the remaining 256 – 58 = 198 non-uniform binary numbers. Thus, each pixel will take on a value ranging from 1 to 59. In our experiments, we used the parameter values $p = 8$, $u = 2$, and $r = 1$ (r is the radius of the descriptor).

LBP values are first computed at each pixel in the (normalized) face image. For each patch i, a histogram of the LBP values $S_i' \in \mathbb{Z}^{d_s}$ is computed (where $d_s = 59$). This feature vector is then normalized to $S_i = \frac{S_i'}{\sum_i^{d_s} S_i'}$. The final structural representation of the avatar face is a set of N feature vectors $S_i, i = 1 \ldots N$.

Appearance Representation

The appearance representation utilizes the discriminative information provided by the color of the avatar faces to help determine their identity. The range and distribution of colors that human faces possess is much smaller and more compactly distributed than that of avatar faces. For this reason, a descriptor called the Spatial Appearance Descriptor (SAD) is developed.

The first step in the appearance descriptor is to reduce the size of the color space. Using the screen capture method of acquiring avatar faces, each pixel can take 256^3 unique values, which is too verbose for an efficient representation. Using a training set of normalized avatar faces, we reduce the color space to K discrete levels $(K \ll 256^3)$ by performing K-means clustering. Each RGB color pixel in the face is mapped to the nearest color in the set of K colors. When selecting K, there were two tradeoffs. Larger values of K made the descriptor more explanatory, but could greatly increase the time for MRF convergence. On the other hand, smaller values of K were fast to compute/converge, but the accuracy and descriptiveness of the descriptor would be compromised. The selected value of K=8 was empirically determined to maximize these two competing factors.

MRF Image Smoothing

In order to ensure that color labels are compact and isolated, it is necessary that the colors be spatially smoothed. This is achieved by modeling the color distribution in the face using a Markov random field (MRF)[18]. In this case, the MRF labels are the K discrete colors $L_k, k = 1 \ldots K$,discovered in the clustering step.

For a (normalized) face image with n pixels, the compatibility of pixel $p_i, i = 1 \ldots n$, with label $l_i \in L_k$ is denoted as

$$\Phi(p_i, l_i) = \exp\{-d(p_i, l_i)^2/\sigma_\Phi\} \tag{1}$$

where $d(p_i, l_i)$ is the Euclidean distance between the pixel p_i and the label l_i in the RGB space.

A spatial smoothness constraint is imposed on the color labeling of the face using the smoothing function $\Psi(l_i, l_j), j \in R_i$, where R_i is the set of pixels in the clique of pixel p_i (in our case R_i is equivalent to the 4-connected pixels of pixel p_i) and $l_i \in L_k$ is the label assigned to pixel i.

The smoothing function Ψ is designed such that a penalty is imposed when (i) two neighboring pixels with similar observations take on different labels, and (ii) two neighboring pixels with distant observations take on the same label. To express this within Ψ, we have

$$\Psi(l_i, l_j) = \exp\{z(i, j)/\sigma_\Psi\} \tag{2}$$

$$z(i, j) = |I_l(l_i, l_j) - I_p(p_i, p_j)| \tag{3}$$

where I_l and I_p are both indicator functions. $I_l(l_i, l_j)$ indicates whether the labels at pixels i and j are the same color label: it is 1 when both inputs are the same label and 0 otherwise. $I_p(p_i, p_j)$ indicates being members of an edge or not, and follows the rule $I_p(p_i, p_j) = 0$ if $||p_i - p_j|| < \tau_E$, else$I_p(p_i, p_j) = 1$. Our experiments used the value $\tau_E = .3$.

Considering Φ and Ψ to represent pseudo-probabilities $\tilde{p}$ (for compactness, we did include a normalization constant - the result is the same but they are not true probabilities), the MRF is solved by selecting the label for each pixel that maximizes the global likelihood of the expression

$$\tilde{p}(p_1 \ldots p_n, l_1 \ldots l_n) = \prod_{i=1}^{n} \left(\Phi(p_i, l_i) \prod_{j \in \tilde{R}_j} \Psi(l_i, l_j) \right) \tag{4}$$

For a given image, p_i is fixed. Thus, we need to select the l_i that maximize Eq. 4. The optimal solution to this problem is NP-hard to solve, but many methods exists to find an approximate solution. We employed the iterative belief propagation algorithm[19] to accomplish this task. Belief propagation utilizes a message passing framework, where information can propagate to neighboring sites through messages. Letting $M(i, j)$ denote a message from pixel i to pixel j (such that $i \in R_j$), we have

$$M(i, j) = \max_{l_i} \left\{ \Phi(p_i, l_i) \Psi(l_i, l_j) \prod_{k \in \tilde{R}_i} M'(k, i) \right\} \tag{5}$$

where M' denotes messages from the previous iteration, and $\tilde{R}_i$ is the four-connected pixels of i, with pixel j excluded. Using belief propagation, messages pass until either convergence or a predefined number of iterations is completed. Because our application necessitated speed over minor segmentation improvements, we fixed the number of belief propagation iterations to 2.

After every iteration of belief propagation, the maximum a posteriori solution (consisting of color labels I'_i for each pixel p_i) is obtained through

$$I'_i = \max_{l_i} \left\{ \Phi(p_i, l_i) \prod_j M(j, i) \right\} \tag{6}$$

Our MRF segmentation used a multi-scale MRF, as described by Wang and Tang[20] and further extended by Klare, Li and Jain[21]. A multi-scale MRF allows messages to pass over larger distances, which generally improves the solution. Our implementation used three scales.

MRF models are sensitive to the selection of the parameters σ_Φ and σ_Ψ, which alter the importance of the image observation (modeled in Φ) and the smoothness constraint (modeled in Ψ). We (heuristically) selected the values $\sigma_\Phi = 0.5$ and $\sigma_\Psi = 1.0$. MRF segmentation is a critical step in the process, as illustrated in Fig. 1.

Table I. The true accept rates at a false accept rate of 0.01 for each probe set in the Second Life data set. Scores from the fully automated framework are listed, with scores using manual eye location listed in parenthesis.

Set	SAD	LBP	SAD + LBP
B	0.84 (0.90)	0.83 (0.94)	0.88 (0.96)
C	0.88 (0.98)	0.86 (0.98)	0.89 (0.98)
D	0.78 (0.89)	0.75 (0.84)	0.80 (0.95)
E	0.85 (0.96)	0.85 (0.98)	0.86 (0.98)
F	0.80 (0.94)	0.80 (0.96)	0.83 (0.98)
G	0.81 (0.98)	0.83 (0.96)	0.85 (0.98)

Appearance Descriptor: Given a face representation using K spatially smoothed color labels, the spatial variation of a color is described using six features, collectively called the SAD (Spatial Appearance Descriptor): (1) direction and (2) magnitude of the major axis of spatial variation, (3) direction and (4) magnitude of the minor axis of spatial variation, and (5) horizontal and (6) vertical location of the center of mass for the color. Given K colors, and six components of information for each color, the dimension of the SAD descriptor A_i for patch P_i is $d_a = 6K$.

The spatial location of each color is computed by taking the mean of the horizontal and vertical location of each pixel in the patch labeled with color L_k. To calculate the spatial variance information, the 2x2 covariance matrix of the locations for each pixel of color L_k is calculated. The range of each component is normalized to fall within [0,1], using the patch size to determine the normalization of the location and magnitude components, and dividing by π to normalize the direction.

The combined runtime for MRF smoothing and the SAD descriptor takes less than 1.5 seconds per image in Matlab using one 3.0GHz core.

2.3 Matching

For a given avatar face, we have two sets of vectors S_i and A_i, $i = 1 \ldots N$, where N is the number of face patches. To determine an avatar's identity, we first concatenate the set of local (patch) descriptors into a single feature vectors of length Nd_s and Nd_a, respectively for S_i and S_a. The concatenated feature vectors are represented as S^j and A^j for the j-th avatar subject.

The distance between two avatar faces corresponding to images j and k is computed using the Chi square similarity measure. Ahonen et al. demonstrated the merits of this measure[16]. With $S^j(i)$ denoting the i-th component of S^j, the Chi square similarity between images j and k using the structural representation is

$$d'_s(j,k) = \sum_{i=1}^{Nd_s} \frac{(S^j(i) - S^k(i))^2}{S^j(i) + S^k(i)} \tag{7}$$

For the appearance representation we have

$$d'_a(j,k) = \sum_{i=1}^{Nd_a} \frac{(A^j(i) - A^k(i))^2}{A^j(i) + A^k(i)} \tag{8}$$

Finally, these distance are normalized using the min-max score normalization scheme

$$d_s(j,k) = \frac{(d'_s(j,k) - s_{\min})}{(s_{\max} - s_{\min})} \tag{9}$$

$$d_a(j,k) = \frac{(d'_a(j,k) - a_{\min})}{(a_{\max} - a_{\min})} \tag{10}$$

$s_{\min}$ and $s_{\max}$ are, respectively, the minimum and maximum distances computed for d'_s. The same relation holds for $a_{\min}$, $a_{\max}$, and d'_a.

The information contained within the structural and appearance representations are fused at the score level. The final similarity between two faces j and k is computed as

$$d(i,j) = d_s(j,k) + d_a(j,k) \tag{11}$$

A determination of a match is made for avatar face images j and k if $d(i,j) > \tau$, where the threshold τ is varied to tradeoff between false match and false reject errors.

3. EXPERIMENTS AND DISCUSSION

3.1 Data

For the purposes of avatar facial data generation, various virtual worlds and avatar creation software were considered based on the needs of this project, including[8]:

- Ability to view avatar from multiple angles
- Selection of contrasting facial features in generating new avatars
- Mutable attributes to avatar facial features
- General ease of use and versatility

Set 1 - AvMaker

The first avatar dataset used in our experiments was based on the images in the FERET[22] dataset. We converted FERET images to images of 3D avatars using the state of the art picture-to-avatar conversion software: Cyber Extruder's AvMaker[23]. AvMaker allows the user to make his own avatar skin for Second Life from a simple passport like photo and then upload the created skin to the user's avatar. In the FERET-to-Avatar dataset an avatar face was generated from every photo in the FERET database. In total, we produced 2,020 avatar images for 725 subjects.

Set 2 – Second-Life

While many virtual worlds such as *Entropia Universe* and avatar creation software such as *Poser*[24] were considered, the Massively Multiplayer Online Role Playing Game (MMORPG) *Second Life* was found to be the best fit to the above criteria. We designed and implemented a scripting technique to automatically collect a dataset of avatar faces. Using the programming language AutoIt as well as a scripting language native to *Second Life*, better known as Linden Scripting Language (LSL), a successful generation of avatars was achieved. The resulting dataset consists of avatars with ten images from different angles at a resolution of 1280 X 1024. Seven facial pictures were taken for each avatar, at differing angles (see Fig 4). The frontal pose shown in Set A was used as the gallery image for each subject. Matching experiments were then performed using the remaining images as probes.

3.2 Results

Set 1 - AvMaker Dataset

The avatar images for each of the 725 subjects created from the FERET dataset were matched against a photograph image for each subject. The photograph image was not the same image from which the 3D avatar was created.

Using Cogntec's FaceVACS face recognition SDK[25], a Rank-1 accuracy of 99.58% was achieved. Testing the prototype avatar system proposed on this dataset was not warranted because: (i) A COTS face matcher (FaceVACS) already achieves a very high accuracy, and (ii) the proposed system is designed to match avatars to avatars using less realistic renderings, such as those in Second Life. These results are reported because it is useful to know that current face recognition technology appears to be sufficiently accurate in the avatar to human matching scenario, when the avatar is rendered using advanced software such as AvMaker.

Set 2 – Second-Life

The first 20 subjects from the Second-Life set were used to as training. Table I lists the true accept rates when operating at false accept rate of 0.01 using the remaining 80 subjects. Table I lists both the accuracy when both manually marked and automated eye locations are used. We separately list the accuracy of the SAD descriptor (second column), the LBP descriptor (third column), and the fusion of the two descriptors as described in Sec. III(c) (fourth column). Based on these results, we observe that (i) automated face and eye detection are currently a bottleneck in performance, and (ii) minor pose variations appear to further degrade the recognition accuracy when using the fully automated framework. While both the structural (LBP) and appearance (SAD) features offer adequate performance when operating individually, combining their discriminative information using match score fusion yields a consistent improvement in accuracy.

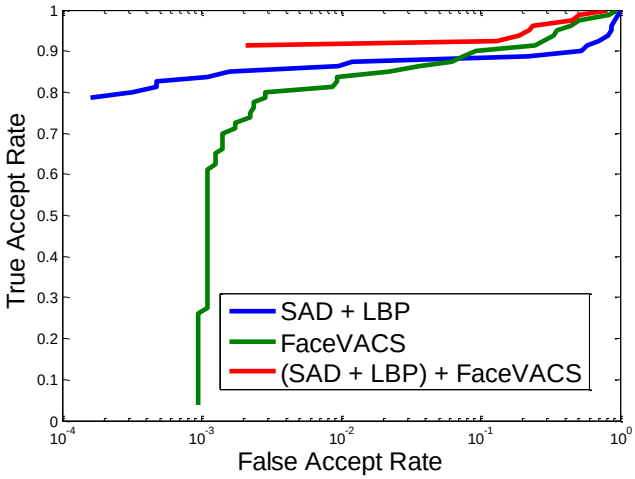

Figure 4. ROC plot of using Second-Life probe set B. Fusion of the proposed framework with FaceVACS further improves the recognition accuracy.

Fig. 4 shows the ROC plot of accuracy of FaceVACS compared to the proposed method. We input the same automated eye locations for FaceVACS as used in the proposed method because the eyes in input avatar images could not be automatically detected by FaceVACS. Fusion of the proposed SAD + LBP framework with FaceVACS resulted in improved recognition rates.

Fig. 5 shows eight avatar pairs. Four correctly matched ones and four mismatched pairs. We believe that erroneous matches are a byproduct of the avatar face generation process which does not provide sufficient facial diversity in the available modeling software. Certain facial structures (noses, lips, etc.) may be reused in the process of avatar generation and consequently lead to false matches. Similar issues arise in face recognition of twins and other close relatives. As artificial face modeling and generation technology improves this type of problem should become less prominent.

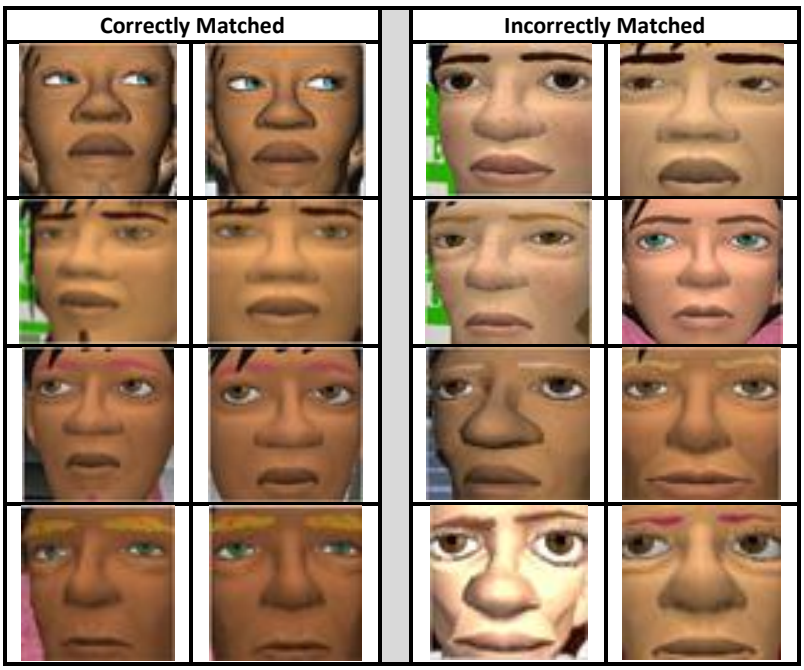

Figure 5. Examples of correctly and incorrectly matched avatar faces.

4. CONCLUSIONS AND FUTURE WORK

This paper addressed the problem of recognizing avatar faces. We have reported results of experiments aimed at within-virtual-world avatar authentication and inter-reality-based scenarios of tracking a person between real

and virtual worlds. Potential directions for future research include the investigation of other visual and behavioral approaches to virtual world security based on appearance of new characteristics and abilities in the avatars of tomorrow. As virtual reality technology progresses, it will require new security solutions for identity management across worlds populated by both human and artificial entities[4].

In our future work we plan to develop automated solutions for a number of scenarios related to avatar authentication, including: matching of an avatar representing the same person from different virtual worlds, matching of a sketch to an image of an avatar, simultaneously profiling multiple independent visual and behavioral characteristic of an avatar to increase the authentication accuracy and conducting experiments on developing a multimodal system capable of authenticating both biological (human being) and non-biological (avatars) entities.

REFERENCES

[1] N. O'Brien, "Spies watch rise of virtual terrorists", Available at: http://www.news.com.au/spies-watch-rise-of-virtual-terrorists/story-e6frfkp9-1111114075761 (July 31, 2007).

[2] R. O'Harrow, "Spies' Battleground Turns Virtual", Available at: http://www.washingtonpost.com/wp-dyn/content/article/2008/02/05/AR2008020503144.html (February 6, 2008).

[3] A. Reuters, "US Congress launches probe into virtual economies", Available at: http://secondlife.reuters.com/stories/2006/10/15/us-congress-launchs-probe-into-virtual-economies (October 15, 2006).

[4] R. Yampolskiy, and M. Gavrilova, "Applying Biometric Principles to Avatar Recognition", Proc. International Conference on Cyberworlds (CW2010). Singapore (October 20-22, 2010).

[5] M. Lyons, A. Plante, S. Jehan *et al.*, "Avatar Creation using Automatic Face Recognition",Proc. ACM Multimedia 98. Bristol, England (Sept. 1998).

[6] S. Z. Li and A. K. Jain (eds.), "Handbook of Face Recognition", Second Edition, Springer, 2011.

[7] S. N. Yanushkevich, S. Stoica, and V. P. Shmerko, "Synthetic Biometrics," Computational Intelligence Magazine, 2(2), 60 - 69 (2007).

[8] J. N. Oursler, M. Price, and R. V. Yampolskiy, "Parameterized Generation of Avatar Face Dataset", Proc. 14th International Conference on Computer Games: AI, Animation, Mobile, Interactive Multimedia, Educational & Serious Games. Louisville, KY (2009).

[9] S. Ajina, R. V. Yampolskiy, and N. E. B. Amara, "Authentification de Visages D'avatar", Proc. Confere 2010. Sousse, Tunisia (July 1-2, 2010).

[10] M. Boukhris, M. Beck, A. A. Mohamed *et al.*, "Artificial Human Face Recognition via Daubechies Wavelet Transform and SVM", Proc. 16th International Conference on Computer Games (CGAMES). Louisville, KY, USA (2011).

[11] A. A. Mohamed, and R. V. Yampolskiy, "An improved LBP algorithm for avatar face recognition", Proc. 23rd International Symposium on Information, Communication and Automation Technologies (ICAT2011). Sarajevo, Bosnia & Herzegovina (2011).

[12] R. V. Yampolskiy, G. Cho, R. Rosenthal *et al.*, "Evaluation of Face Recognition Algorithms on Avatar Face Datasets", Proc. International Conference on Cyberworlds. Banff, Canada (2011).

[13] P. Viola, and M. J. Jones, "Robust real-time face detection," Int. Journal of Computer Vision, 57, 137-154 (2004).

[14] Nevadaigtf.org, Player id, age verification and border control technology forum. (2005).

[15] J. R. Beveridge, D. Bolme, B. A. Draper *et al.*, "The CSU face identification evaluation system," Machine Vision and Applications, 16, 128-138 (2005).

[16] T. Ahonen, A. Hadid, and M. Pietikainen, "Face description with local binary patterns: Application to face recognition," IEEE Trans. Pattern Analysis and Machine Intelligence, 28, 2037-2041 (2006).

[17] T. Ojala, M. Pietikainen, and T. Maenpaa, "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns," IEEE Trans. Pattern Analysis & Machine Intelligence, 24, 971-987 (2002).

[18] S. Z. Li, [Markov Random Field Modeling in Image Analysis] Springer, (2009).

[19] W. T. Freeman, E. C. Pasztor, and O. T. Carmichael, "Learning low level vision," Int. Journal of Computer Vision, 40, 25-47 (2000).

[20] X. Wang, and X. Tang, "Face photo-sketch synthesis and recognition," IEEE Trans. Pattern Analysis & Machine Intelligence, 31, 1955-1967 (2009).

[21] B. Klare, Z. Li, and A. K. Jain, "Matching Forensic Sketches to Mugshot Photos", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 33, No. 3, pp. 639-646, March 2011.

[22] P. J. Phillips, H. Moon, P. J. Rauss *et al.*, "The FERET evaluation methodology for face recognition algorithms," IEEE Transactions on Pattern Analysis and Machine Intelligence, 22(10), 1090-1104 (October 2000).

[23] CyberExtruder, "AvMaker", Available at: http://www.cyberextruder.com/avatars (September 2010).

[24] SmithMicro, "Poser - Easily Create 3D Character Art and Animation", Available at: http://poser.smithmicro.com/poser.html (September 2010).

[25] R. Dhamija, and A. Perrig, "Deja Vu: A User Study. Using Images for Authentication." Proc. 9th USENIX Security Symposium. Denver, Colorado. (2000).