Strategy-based behavioural biometrics: a novel approach to automated identification

Roman V. Yampolskiy*

Department of Computer Engineering and Computer Science,
University of Louisville, JB Speed Hall, #106,
Louisville, KY 40292, USA
E-mail: roman.yampolskiy@louisville.edu
*Corresponding author

Venu Govindaraju

Department of Computer Science and Engineering,
CUBS, University at Buffalo, 520 Lee Entrance,
Suite 202, Buffalo, NY 14228, USA
E-mail: govind@buffalo.edu

Abstract: Behavioural intrusion detection is a frequently used for insuring network security. We expend behaviour based intrusion detection approach to a new domain of game networks. Specifically, our research shows that a behavioural biometric signature can be generated based on the strategy used by an individual to play a game. We wrote software capable of automatically extracting behavioural profiles for each player in a game of Poker. Once a behavioural signature is generated for a player, it is continuously compared against player's current actions. Any significant deviations in behaviour are reported to the game server administrator as potential security breaches.

Keywords: behavioural biometrics; dissimilarity functions; strategy biometric.

Reference to this paper should be made as follows: Yampolskiy, R.V. and Govindaraju, V. (2009) 'Strategy-based behavioural biometrics: a novel approach to automated identification', *Int. J. Computer Applications in Technology*, Vol. 35, No. 1, pp.29–41.

Biographical notes: Roman V. Yampolskiy holds a PhD Degree from the Department of Computer Science and Engineering at the University at Buffalo. After graduating, he served as an Affiliate Academic at the University of London, College of London until finally accepting an Assistant Professor position at the University of Louisville in 2008. His main areas of interest are computer security, artificial intelligence, behavioural biometrics and intrusion detection. He is an author of over 40 publications including multiple books.

Venu Govindaraju is Professor of Computer Science and Engineering at the University at Buffalo (SUNY Buffalo). He received his BTech (Honours) from the Indian Institute of Technology (IIT), Kharagpur, India in 1986, and his PhD from UB in 1992. He has co-authored more than 230 scientific papers. He has been the PI/Co-PI of projects funded by government and industry for over 50 million dollars in the last 15 years. He is the founding director of the Centre for Unified Biometrics and Sensors (CUBS) and associate director of the Centre for Document Analysis and Recognition (CEDAR).

1 Introduction

Behavioural biometrics are a valuable tool in many security tasks that require identification or verification of an individual. Behavioural biometrics are often employed because they can be easily collected non-obtrusively and are particularly useful in situations that do not provide an opportunity for collection of stronger (more reliable) biometric data. We investigated strategy used while playing a game as a type of a behavioural biometric. The game of poker is used as an example of a game with a clearly identifiable player strategy. The profile produced for each player is used as the person's behavioural-biometric profile.

Our approach can be used by online casinos to detect a hacker who is using a stolen account on a game server or a player trying to cheat by using an AI bot in order to win more money. Both are currently big problems in the world of online gaming and a successful solution is beneficial not just from theoretical but also from a practical point of view. Advantages of the developed solution are listed below:

- no special hardware required
- no noticeable enrolment period
- provides continuous player verification
- identifies user not the system or geographic location.

2 System description

Based on the idea of using strategy followed while playing a game as a type of a behavioural biometric we propose a complete system for player verification. First a player profile is generated either by data mining an existing database of poker hands or by observing a live game of poker. Next a similarity measure is obtained between the feature vector generated based on the recently collected player data and the data for the same player obtained in previous sessions. A score is generated indicating how similar the current style of play is to the historically shown style of play for a particular player. If a score is above a certain threshold, it might indicate that a different user from the one who has originally registered is using the account and so the administrator of the casino needs to be alerted to that fact. If the score is below some threshold, the system continues collecting and analysing the player data.

As can be seen from Figure 1 using a previously generated database of poker hands does not provide an option of continuous monitoring and so is an inferior alternative, which might be valuable in terms of initial experimentation, but which must be replaced by live data collection for the completed product.

Figure 1 A diagram of the developed system

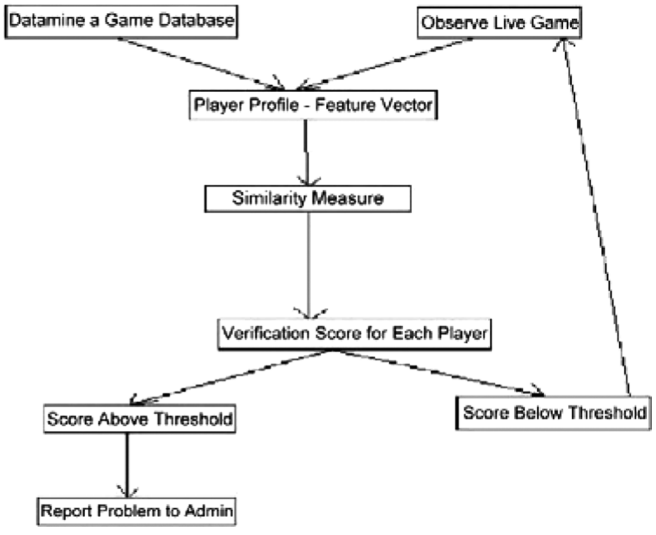

2.1 Data

Data for poker related experiments could be obtained in several ways: by observing real games played by human opponents in casinos, home games and at online gambling sites or by utilising existing poker-hand databases.

2.1.1 IRC poker dataset

Long before online casinos became prevalent on the Internet, there exists the Internet Relay Chat (IRC) poker server. Michael Maurer developed a program he called the Observer that resided on the IRC poker channels and monitored and logged the details of every game it observed. This resulted in the collection of the more than 10 million complete poker hands (from 1995 to 2001) that comprise the IRC Poker Database (Maurer, 2005).

2.1.2 Data from observing human play

Additional data for poker related experiments could be obtained by observing many real games played by human opponents in casinos, home games and at online gambling sites.

Data from observing online human play. This is probably the best source of data since the games are multiple in number, stakes and quality of players. The data is same as the data potentially generated in the desired field of application of the final algorithm. The data can also be easily collected automatically by creating a simple observer bot or even easier collected by the online casino itself if it desires to do so, perhaps for security reasons. In fact most casinos already do collect some game data if not the complete information about every hand played.

Data from observing off-line human play. Another alternative is to attempt to collect data from real brick and mortar casinos, home games and tournaments. This approach is interesting because many additional factors may be collected which are not available in an online setting and are generally referred to as tells. But this information is neither objective nor useful for our purposes since it will not be possible to obtain similar data while employing our algorithm online. Finally it is a daunting manual task to collect data from real life human play, which results in expensive and potentially full of errors set of statistics.

3 Generation of synthetic profiles

Many biometric technologies are still in their infancy and do not yet have large reliable data sets which are needed for further development of such systems. One solution to the problem of insufficient availability of training and testing biometric data is the creation of the so-called simulated or synthetic biometric data using sophisticated computer algorithms (Yanushkevich et al., 2004).

While ideally we want all our biometric systems to be tested on real data to insure the highest standard of quality and system security is obtained it is not always possible for a variety of reasons. Testing a biometric system requires many thousands of samples in order to establish system's false reject and False Accept Rates (FARs). Obtaining biometric data in sufficient quantity is a time consuming and expansive process. Volunteers quickly get bored with a repetitive task of biometric data collection

and paying people for their cooperation is often beyond the budget of many research centres. Also dealing with real biometric data brings up issues with privacy of individuals providing their data and with security of biometric databases (Sumi and Matsuyama, 2005).

Synthetic data addresses many of the concerns presented above. Once a system for producing simulated data is developed it is fast and cheap to produce large quantities of high quality biometric data which adheres to statistical distributions desired by the investigator without any privacy or security concerns to worry about. Such data can be used for testing newly developed biometric systems, benchmarking well developed security systems, testing scalability of authentication systems or for certification of commercially available packages. Production of synthetic biometrics allows researchers to better understand individuality of biometric patterns and allows parametric sensitivities of algorithms to be investigated in greater detail (Ma et al., 2005; Makthal and Ross, 2005; Orlans et al., 2004).

A number of approaches exist for the generation of the artificial biometric data. Most of them are concerned with the creation of a simulated image depicting a particular physical biometric such as a fingerprint, face or iris. The existing approaches can be grouped into the following categories:

- distortion of an existing image to generate numerous similar images
- combination of multiple images to produce a novel image with partial properties of all the seed images
- generation of images based on physical models for the biometric in question.

Because features of a strategic profile have meaning, unlike minutiae points of fingerprints or colours of iris, it is possible to apply a fourth methodology for creation of synthetic game-based behavioural biometrics, namely parameterised design. It may not make any sense to design a synthetic fingerprint with all the minutiae points located in just the left half of the fingerprint but it makes sense to have a strategic profile for a player who is only aggressive in the first two rounds of the game. A fifth and final option we are not really considering here is generation of a synthetic-behavioural-biometric based on purely random approach. Such simulated data would correspond to unrealistic strategies not encountered at real world poker tables and as such would be completely useless for our purposes of generating realistic artificial biometric data possessing all properties of the authentic samples.

The first approach to the creation of the synthetic behavioural biometric we have implemented is based on taking an existing player profile and modifying it to make numerous additional profiles similar to the given one (Yampolskiy and Govindaraju, 2008). Poker is a game of high variance and so even by following the same exact strategy it is possible for a player to play a slightly different number of cards at every stage based on the actual cards being dealt to him. We have estimated poker variance for a reasonably large number of played hands to be around 3%. By taking all the feature points in a given profile and replacing them with new randomly generated values in the range of +/−3% of the given ones we are able to obtain multiple artificial profiles for the same player which are representative of the authentic profiles which could be produced by the same player due to the degree of natural variance in the game of poker. Any resulting values outside of the range from 0% to 100% are changed to the closest values falling in the range. Obviously it is also possible to adjust the variance rate to accommodate different playing styles and types of poker games. An example of one profile generated in such a way based on a temporal-seed-profile from player Bob is shown in Table 1.

Table 1 A sample profile generated from a seed profile

Player name: Synthetic Bob			*Hands dealt: n/a*	
	Pre-flop (%)	*Flop* (%)	*Turn* (%)	*River* (%)
No. of hands played	n/a	n/a	n/a	n/a
Folded	68.4	26.7	23.8	20.1
Checked	5.7	52.3	51.8	53.8
Called	21.3	32.1	28.2	33.4
Raised	4.9	3.5	4.4	5.7
Check-raised	2.2	2.9	2.0	2.7
Re-raised	1	0.2	0.0	0.4
All-in	1.5	3.4	4.9	39.5

This methodology is best for generation of multiple profiles for the same individual which can be used for testing of verification or even identification abilities of strategy based behavioural biometric systems.

Second approach to the generation of artificial behavioural biometric data we implemented is based on combining feature points from two or more different seed profiles. This can be done in two different ways either simply picking one of the profiles as providing a particular data point for the profile being created or taking average of the values in the seed profiles to serve as the new value. This methodology is similar to the crossover operations used in genetic algorithms for production of the next generation of solutions from the currently available distribution of genetic strings (Yampolskiy et al., 2004; Goldberg, 1989).

In case of strategic profiles this approach leads to the production of profiles representing somewhat averaged strategies. For example, combining an overly aggressive and a passive profile results in a solid profile typical of many good players at low-level tables. This methodology is best for generation of multiple profiles needed to make sure our database is sufficiently large to make verification of particular individuals of interest a non-trivial task. It also works well for generation of novel strategic profiles not yet encountered during the collection of authentic data and so

insures diversity of strategies encountered by the biometric processing system.

An approach corresponding to the generation of synthetic biometrics based on physical models with respect to poker strategies is achieved by creation of realistically behaving artificial poker players. It is up to the poker experts to develop multiple strategically interesting poker bots. Typically a number of basic strategies are used for the initial design and by adding some behavioural variation at particular stages of the game new strategies are introduced. Most popular basic profiles are Solid, Rock, Maniac, Fish, and Typical (Online, 2006). Once such players are created they are allowed to play against each other or against human opponents while the system generates corresponding strategy based behavioural profiles for them which also include the contextual information about the flop, player's position and stages of the hand. Profiles produced in such a way show a very high degree of over time consistency as computerised players are not subject to psychological swings so typical of human players commonly referred to as going on tilt (Schoonmaker, 2005).

Our implementation of poker bots was done using the statistical package called Online Hold'em Inspector version 2.26d4 (Online, 2006). By specifying such conditions as tendency of our bots to bluff, slow play, check raise and their aggressiveness level as well as their pre-flop card selection we were able to generate numerous valid poker strategies. Validity of our poker bots was tested at low-stakes real-money online poker tables against human opponents where our bots consistently scored around three big bets per hour in profits (Yampolskiy and Govindaraju, 2008). Figures 2 and 3 demonstrate bot's characteristics, which we were able to manipulate.

Figure 2 Flop playing strategy menu

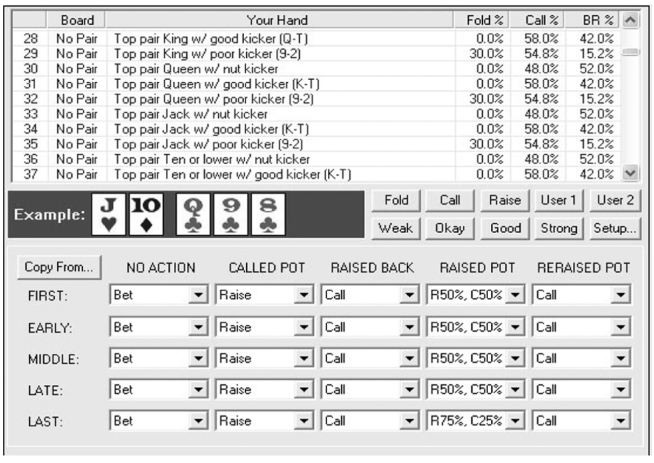

By manipulating hundreds of variables associated with our bots playing strategy and combining them in numerous ways we were able to generate a multitude of realistically behaving poker players and as a result collected behavioural biometric data on all such strategies. Also by statistically analysing our bot's strategy we were able to predict some characteristics of the bot's behavioural profile as shown in Figure 4.

Figure 3 Pre-flop hand selection strategy menu

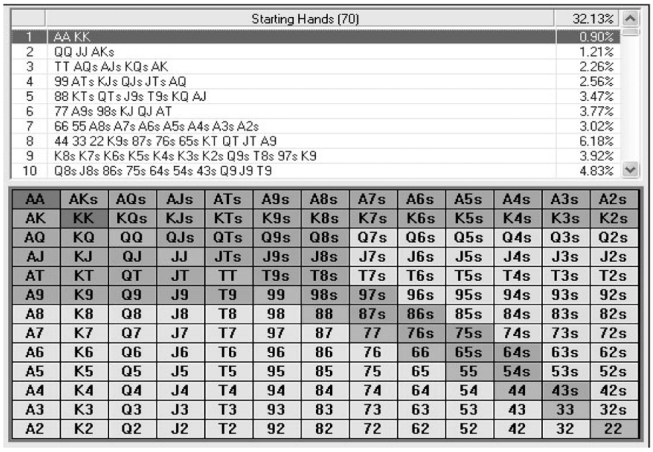

Figure 4 Estimates of a statistical profile for a bot's strategy

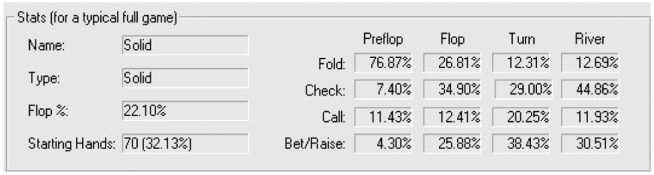

Finally, we get to an approach we call parameterised design. Because we are not generating a raw biometric image but rather a set of feature measurements we are able to declare with statistical parameters in which ranges we wish all feature points to reside. This is a somewhat inverted approach from the one utilising artificially intelligent poker playing programs. Instead of designing a poker strategy which can be observed to produce a statistical profile describing player's behaviour we are directly generating the statistical feature vector which is parameterised with the intention of representing a valid game strategy. This is a less tedious approach as instead of prescribing particular actions to each one of the thousands of possible situations in a game of poker we only have to specify some general trends such as aggression and card selection at different phases of the game.

Generally a style of a poker player is represented as a point on a 2-dimensional styles grid. The y dimension represents the tight/loose score and the x dimension stands for the passive/aggressive behaviour of the player. Players are measured on each dimension from 1 to 9 (Schoonmaker, 2005). This gives us up to 81 different playing styles which is sufficient for production of baseline profiles for testing of verification systems. The resulting baseline profiles can later be used to generate additional profiles using methods presented above.

Our algorithm takes an (x, y) pair and produces a behavioural profile which confirms to a statistically predicted action frequency distribution for this particular playing style. For example a loose and passive player commonly refereed to as a 'Calling Station' is represented by a point $(9, 1)$ on a playing styles grid and would correspond to a behavioural profile which looks at over 89% of flops and bets or raises less than 11% of the time

basically only if he holds the absolutely best cards at the moment. By expanding those ideas to all four stages of the game (pre-flop, flop, turn, river) we are able to produce behavioural profiles corresponding to different styles of play. We can see that each grid value controls about 11% of style space and so different styles are very easy to distinguish using statistical analyser as the variance in the game of poker is around 3%.

This approach provides a way to control the properties of the behavioural biometric profile via specified parameters. This is particularly useful if we wish to run a controlled experiment, for example seeing how our system performs in a large field of tight/aggressive players such as found at high limit games for which actual testing of the system may be beyond the means for many researchers.

3.1 Statistical measure of player's style

If we are going to study the game of poker and more particularly the style of our opponents scientifically, we will need to quantify and statistically analyse our opponents' behaviour. In order to do so we propose and define a number of variables associated with actions of our opponents. The parameters chosen are selected because they can be easily tracked by relatively straightforward methodologies and more importantly they are believed to accurately describe the long-term model of behaviour of our opponents (Software, 2005; Poker-Edge.com, 2006; Brandt, 2005).

The following list of variables represents individual values within each feature vector. Combining individual values in the feature vector in varying ways may generate additional descriptors. Some important statistics such as the total number of hands played are kept for the internal bookkeeping but are not a part of the feature vector.

Fold percentage of times this particular player has decided to give up his claims to the pot.

Check percentage of times this particular player has decided not to invest any additional money into the pot.

Call percentage of times this particular player has paid an amount equivalent to the raise by some other player ahead in position in order to keep playing this hand.

Check-raise percentage of times a player has checked allowing another player to put some money into the pot, just to come over the top and raise the pot after the action gets back to him.

Raise percentage of times this particular player has chosen to raise the stakes.

Re-raise percentage of times this particular player has chosen to re-raise somebody-else's raise. This would include a re-re-raise and re-re-re-raise so on.

All-in percentage of times this particular player has chosen to invest all his money in the current hand.

A combination of such statistical variables taken together produces a feature vector which is used by a pattern recognition algorithm to determine if a current profile is consistent with that previously seen one from this particular player or if a possible intruder has taken the control of the account.

Descriptive accuracy of a behavioural profile can be greatly increased if additional information is included. We have utilised a profile structure which separates player's actions into the four stages of the hand, making temporal information available, and as a result, description of player's strategy more meaningful. Table 2 is an example of such temporal profile.

Table 2 Temporal strategy profile

Player name: Bob				*Hands dealt: 224*
	Pre-flop	*Flop*	*Turn*	*River*
No. of hands played	224	68	46	33
Folded	67%	28%	24%	18%
Checked	7%	54%	52%	52%
Called	21%	32%	28%	33%
Raised	4%	1%	4%	6%
Check-raised	0%	4%	0%	0%
Re-raised	0%	1%	0%	0%
All-in	1%	3%	4%	39%

Source: Yampolskiy and Govindaraju (2006a)

Profiles can be further enhanced with the inclusion of spatial information, essentially making a separate profile for each of the ten positions a player can have around the table. Such profiles clearly demonstrate dependence of player's strategy on position and are shown in Table 3.

Table 3 Spatial strategy profile

Action	*Small*	*Big*	*3rd*	*4th*	*5th*	*6th*	*7th*	*8th*	*9th*	*Dealer*
Folded	77	73	71	69	67	64	61	59	57	51
Checked	55	53	50	49	48	44	41	39	37	34
Called	14	16	19	22	26	29	33	37	43	53
Raised	2	3	4	6	8	11	13	15	17	20
Check-raised	31	28	23	19	17	15	12	9	6	4
Re-raised	0	1	2	4	6	10	14	18	25	30
All-in	37	39	41	43	47	51	55	59	62	65

Finally with the addition of contextual information about the cards revealed at the flop divided into seven flop types described in the poker literature (as shown in Table 4) (Badizadegan, 1999) we have a 3D information space, which for every stage of the game, every position and every flop provides frequency counts of player's actions as illustrated in Figures 5 and 6.

Dimensionality of such a profile could be extremely high, compared to the basic profiles (Yampolskiy and Govindaraju, 2006b). Table 5 summarises different possible profile types which can be used with strategy based behavioural biometrics along with the information they include and lists the profile's dimensionality. Ideally any similarity measure function we propose to utilise should be flexible enough to handle any of the presented profile types (Yampolskiy and Govindaraju, 2006a).

Figure 5 Poker table with the flow of information

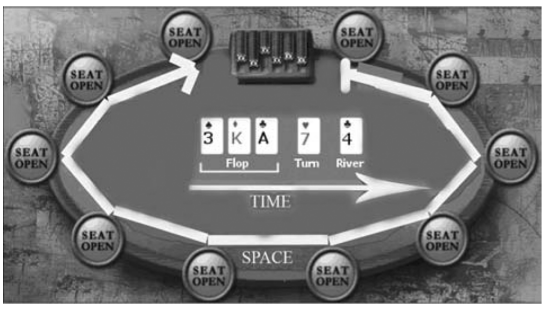

Figure 6 3D profile structure

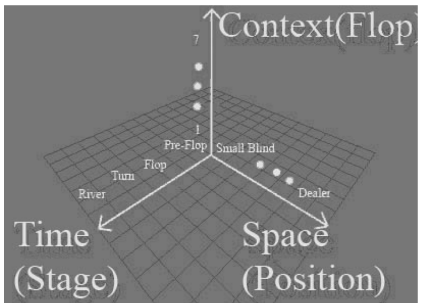

Table 4 Flop types and number of variations for each type

Flop type	*Number of variations*	*Example*
Three cards of the same rank	1	2♣ 2♠ 2♦
Pair plus a 0–3 gapped card	2	J♦ J♣ 10♦
Pair plus a 4+ gapped card	2	K♥ K♠ 7♣
Three cards 0–2 gaps apart	3	9♠ 10♦ J♥
Two cards 0–3 gapped and a third card 4+	3	K♠ 7♣ 8♥
Three cards 3–6 gapped	3	J♥ 10♦ 6♣
Three cards with 4+ gaps between all cards	3	K♠ 8♣ 3♥

Source: Badizadegan (1999)

As the amount of contextual information increases so does the dimensionality of the behavioural profile. This results in what is known as the 'curse of dimensionality'. The matching algorithm needs a large number of feature measurements to account for all the different possibilities of potential situations. The complexity of a high-dimensional space increases exponentially with the number of features. This large collection of features forms a high-dimensional space, in which it is very difficult to find the best decision boundary (Baggenstoss, 2004). One of the similarity measure functions, 2D style measure, examined in this paper is specifically designed to avoid the complications presented by the curse of dimensionality.

Table 5 Profile types by information included and vector dimensionality

Profile type	*Information included*	*Profile dimensionality*
Basic	Frequency counts for actions	7
Temporal	Frequency counts for actions at different stages of the game	$7 \times 4 = 28$
Contextual	Frequency counts for actions with respect to the flop type	$7 \times 7 = 49$
Spatial	Frequency counts for actions at different positions around the table	$7 \times 0 = 70$
Temporal-spatial	Frequency counts for actions with respect to the stage of the game and relative position around the table	$7 \times 10 \times 4 = 280$
Temporal-contextual-spatial	Frequency counts for actions with respect to the stage of the game and relative position around the table and the flop	$7 \times 10 \times 4 + 3 \times 7 \times 7 = 427$

3.2 Similarity measure

When a new biometric data sample is presented to a security system, it is necessary to measure how closely it resembles template data (Yampolskiy and Govindaraju, 2006a). A good similarity measure takes into account statistical characteristics of the data distribution assuming enough data is available to determine such properties (Lee and Park, 2003). Alternatively expert knowledge about the data can be used to optimise a similarity measure function, for example a weighted Euclidean distance function can be developed if it is known that certain features are more valuable then others.

3.2.1 Euclidean distance

One of the most popular similarity distance functions is the Euclidean distance. It is just the square root of the sum of the squared distance between the element of the n-dimensional vectors (x_i, y_i) (Sturn, 2000):

$$d_E = \sqrt{\sum_{i=1}^{n} (x_i - y_i)^2}.$$

Euclidean distance is variant to both adding and multiplying all elements of a vector by a constant factor. It is also variant to the dimensionality of the vectors, for example if missing values reduce the dimension of certain vectors produced output will change. In general the value of Euclidean similarity measure may fall in the range from zero indicating a perfect match to sqrt(n) (where normalised n-dimensional vector is used) indicating maximum dissimilarity of playing styles. Obviously both of

those extreme cases do not occur in real life and represent only theoretical possibilities not related to any viable playing style. In experiments with real life data Euclidean Similarity measure is always in between the two extremes (Yampolskiy and Govindaraju, 2006a, 2006b).

3.2.2 Mahalanobis distance

Mahalanobis distance is defined as (Yampolskiy and Govindaraju, 2006a):

$$d_M = \sqrt{(x-\mu)^T \sum^{-1} (x-\mu)}$$

with mean $\mu = (\mu_1, \mu_2, \mu_3, ..., \mu_n)$ and covariance matrix Σ for a multivariate vector $x = (x_1, x_2, x_3, ..., x_n)$. Mahalanobis distance can also be defined as dissimilarity measure between two random vectors X and Y of the same distribution with the covariance matrix Σ:

$$d_M = \sqrt{(x_i - y_i)^T \sum^{-1} (x_i - y_i)}.$$

If the covariance matrix is the identity matrix then it is the same as Euclidean distance. If the covariance matrix is diagonal, then it is called normalised Euclidean distance:

$$d_{NE} = \sqrt{\sum_{i=1}^{n} \frac{(x_i - y_i)^2}{\sigma_i^2}},$$

where σ_i is the standard deviation of the x_i over the sample set. Mahalanobis distance is not dependent on the scale of measurements (Wikipedia, 2006).

3.2.3 Manhattan distance

The Manhattan distance between two points, in a Euclidean space with fixed Cartesian coordinate system, is the sum of the lengths of the projections of the line segment between the points onto the coordinate axes. In other terms, Manhattan distance is the absolute differences of the elements of the two vectors (x_i, y_i) (Sturn, 2000; Yampolskiy and Govindaraju, 2006a)

$$d_{\text{Man}} = \sum_{i=1}^{n} | x - y |.$$

3.2.4 Weighted Euclidean distance

Performance of the Euclidean similarity measure function can be greatly improved if an expert knowledge about the nature of the data is available. If it is known that some values in the feature vector hold more discriminatory information with respect to others, it is possible to assign proportionally higher weights to such vector components and as a result influence the final outcome of the similarity function (Yampolskiy and Govindaraju, 2006a).

In the case of the poker domain, it is believed by the experts in the field, that the style of the poker player is particularly evident in the pre-flop card selection. Before the flop cards are revealed the player has relatively little information to analyse and often acts based on a small set of rules, which dictate how hands should be played based on the hand itself, position of the player and betting action so far observed. Application of such rules is relatively long-term consistent by most players and so has higher discrimination value as compared to action at the later rounds in the game. In such later rounds additional information about communal cards and opponent reading skills become more important than pre-established rules and so are more situation dependent (Yampolskiy and Govindaraju, 2006a).

3.2.5 2D style measure

This is a similarity measure approach used by human poker experts to classify and compare poker players and is included here to investigate feasibility of using such approaches by computerised systems. Generally a style of a poker player is represented as a point on a 2-dimensional styles grid shown in Figure 7. The y dimension represents the tight/loose score and the x dimension stands for the passive/aggressive behaviour of the player. Players are measured on each dimension from 1 to 9 (Schoonmaker, 2005). For example a loose and passive player, commonly known as a 'Calling Station', is represented by a point (9, 1) on a playing styles grid and would correspond to a behavioural profile which looks at over 89% of flops and bets or raises less than 11% of the time basically only if he holds the absolutely best cards at the moment. This gives us only 81 different playing styles, however mathematically we are not restricted to only integer values for expressing the players' style and so in theory the number of styles can be infinite.

The proposed 2D style measure only takes into account the pre-flop selectiveness of the player and the overall aggressiveness expressed in raising, re-raising and going all-in. The two style descriptors chosen (tightness and aggressiveness) are selected because they are least dependent on elements of chance such as the cards revealed by the board and the playing style of the opposing players. The proposed descriptors are computationally easy to obtain.

Tightness = % of cards folded pre-flop

Aggressiveness = average(% raised + % check raised + % re-raised + % all-in).

Aggressiveness value is determined over all stages of the hand, all possible positions and flop types depending on the type of the behavioural profile used to represent the player's strategy and the availability of the contextual information.

Figure 7 The 2-dimenstinal styles

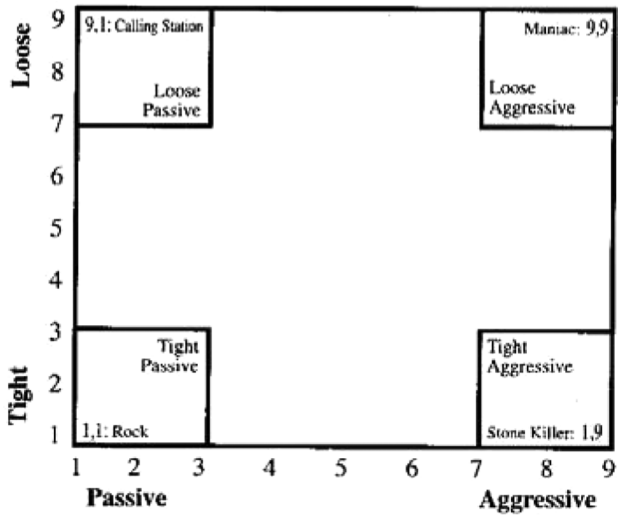

Source: Schoonmaker (2005)

4 Experiments

In this section, we present results of our experiments with user verification and identification as well as our approach to validation of synthetic poker data and results of a spoofing attack on the developed system.

4.1 User verification (Euclidean on temporal profiles)

In a databank of 30 player signatures each one was compared with one profile taken from the same player as the one who generated the original signature and with another profile taken from a randomly chosen player. Giving us an experimental set up in which intruders and legitimate users are equal in number. Using Euclidean similarity measure and a threshold of 75 the original algorithm has positively verified 46.66% (28) users. The FAR was 13.33% (8 users) and False Reject Rate (FRR) was only 8.33% (5 users). This gives us player verification with overall 78.33% accuracy.

We hoped to obtain some improvement in performance of our algorithm as a result of including slow-playing as a part of the feature vector. But possibly because slow-playing is an advanced technique which is not practiced by many players (and even those who do practice it only rarely get the best possible hand to do so), only a small improvement in performance was obtained over previous result (Yampolskiy and Govindaraju, 2006a). This outcome might also be a result of relatively small data set we were working with. It is possible that including other advanced features such as bluffing in addition to slow-playing will produce better results. Using Euclidean measure and a threshold of 75 the algorithm has achieved an accuracy of 80.0%. The FAR was 11.66% (7 users) and FRR was only 8.33% (5 users).

4.2 Intruder identification (Euclidean on temporal profiles)

Once the difference between an expected user's behaviour and the observed behaviour goes over a pre-established threshold a network administrator is notified that an attack may be taking place. While many such occurrences are just false alarms, some do represent accurately detected intrusions. If the network administrator does believe that a real attack took place, he is interested in finding out the identity of the perpetrator. In some instances it may be the case that someone from within the organisation performed an attack, and so the intruder himself has a legitimate account on the same network, probably with fewer privileges (funds) as compared to the compromised account. We investigated the feasibility of determining intruder's identity by comparing the signature for detected deviant behaviour against the database of behavioural signatures from all the users in the system.

For this experiment we used a databank of 30 players. Each player's record contains an original signature from the enrolment period and a second signature from the testing period. Each testing signature was compared against all original signatures in the databank, for a total of 30 comparisons each. The highest matching profile with respect to the similarity measure was recorded as either belonging to the same player (a successful identification) or to a different player (a false identification). From the total of 30 highest matching profiles five were correctly identified and 25 were false matches. This gives us intruder identification with overall 16.66% accuracy. These results are obviously below acceptable industry standards but clearly indicate feasibility of behaviour-based intruder identification. This methodology might give us a certain edge in the fight for network security, but also explains why behavioural biometrics are not typically used for user identification but only for verification.

4.3 Similarity measure function experiments

Experiments were conducted with a 100 authentic user profiles and a 100 impostor profiles used in each. Three different experiments were conducted; in each one a different type of behavioural profile representation was used. Specifically a 28-dimensional temporal profile, a 280-dimensional temporal-spatial profile and a 427-dimensional temporal-spatial-contextual profile were chosen as this allowed us to observe the influence of increasing the amount of environmental information available to the security system on systems performance. We also had an opportunity to observe the effect of the curse of dimensionality with respect to the performance of our similarity measure functions.

For each similarity function a continuously varying threshold curve was generated demonstrating the relationship between FAR and a FRR. Changing threshold

trades the FAR off against the FRR, so the error rates can be adjusted according to the requirements of the security application (Lee and Park, 2003, October). For our experiments the value of the threshold which makes FRR equal to FAR was selected for each similarity measure function and is used as the representative accuracy of the utilised similarity measure function.

We compared three general similarity measure functions (Euclidean, Mahalanobis, Manhattan) with two domain specific functions developed by us (Weighted Euclidean, 2D Style). The weighted Euclidean distance measure we have utilised in our experiments assigns a weight of 3 to all pre-flop features of the vector and weight of 1 to all other features. The weight of 3 has been experimentally established by trial and error of different weights in the range from 1 to 10. The weight is incorporated into the formula by dividing the difference between corresponding values in the two feature vectors by the selected weight.

The 2D style measure approach was designed to counteract the problems with the 'curse of dimensionality' which become particularly taxing with the use of contextual information within the profile. Tightness value is easy to compute as it is simply the average percentage of cards folded pre-flop from all possible positions. Aggressiveness value is slightly more involved but is essentially the average percentage of raised, check raised, re-raised and all in actions from all possible positions at all stages of the game and for each possibly flop type.

As can be seen from Table 6 general similarity measure functions (Euclidean, Mahalanobis and Manhattan) showed a very similar performance, with Mahalanobis distance being slightly inferior to Euclidean and Manhattan distances which showed identical performance of 12% Equal Error Rate (EER). Best performance was shown by a task specific Weighted Euclidean distance which had a 10% EER. 2D Style measure performed poorly in case of temporal profiles, probably because some of the discriminatory power is lost in the averaging process.

Table 6 Verification results using temporal profiles

Similarity measure	*Equal Error Rate* (%)
Euclidean distance	12
Mahalanobis distance	13
Manhattan distance	12
Weighted Euclidean distance	10
2D style measure	14

A great improvement in performance of the strategy based behavioural biometric system was observed with the inclusion of spatial information into the profiles as demonstrated in Table 7. Once again the Weighted Euclidean distance function was the best matching algorithm obtaining 7% EER with general similarity measure functions performing in the range of 9–10% EER. However, performance of the 2D style measure actually became worse to the level of 25% EER.

Table 7 Verification results using temporal-spatial profiles

Similarity measure	*Equal Error Rate* (%)
Euclidean distance	9
Mahalanobis distance	10
Manhattan distance	9
Weighted Euclidean distance	7
2D style measure	25

Improvement in the performance of most similarity measure functions can be explained by a more refined capture of the player's strategy associated with inclusion of information about the spatial location of the player. Decreased performance of the 2D style matcher probably resulted from the influence of zero-value variables on the overall profile average. Zero-value variables are a consequence of not having enough data points in a high-dimensionality profile such as 280-dimensional spatial-temporal profile.

As can be seen from Table 8 with the inclusion of the contextual information the dimensionality of behavioural profile has ballooned to 427D and the influence of the 'curse of dimensionality' became apparent. Performance of all similarity measures has significantly decreased, with that of 2D style measure to almost the point of random guessing. With such a high-dimensionality-behavioural-profile the number of zero-value variables becomes overwhelming as the amount of time needed to collect sufficient data is unreasonable for any real-life security system.

Table 8 Verification using temporal-spatial-contextual profiles

Similarity measure	*Equal Error Rate* (%)
Euclidean distance	33
Mahalanobis distance	36
Manhattan distance	33
Weighted Euclidean distance	29
2D style measure	46

4.4 Validation of synthetic data

We have also performed testing using our biometric verification system which uses a Weighted Euclidean distance measure with an experimentally determined optimal threshold (Yampolskiy and Govindaraju, 2007). For each experiment 100 artificial baseline player profiles have been generated using one of the developed methodologies along with a 100 of testing profiles. In each experiment the number of legitimate users and imposters was equal with no overlap between testing and baseline profiles. Imposter profiles were randomly chosen from profiles unrelated to the baseline one.

Table 9 compares EER obtained on artificial data with that reported from the original experiments on genuine data (Yampolskiy and Govindaraju, 2008). As can be seen, with some data generation methodologies, we have obtained

accuracy levels statistically indistinguishable from those originally produced by the system on genuine data. In particular, approaches based on modifying seed profile, parameterised design and observation of AI players showed the best results. Also the ROC curves of those methods were an almost exact match with the ones from the experiments on genuine data. This leads us to believe that both intra-class and inter-class variation of strategy-based profiles is well simulated with those approaches.

Table 9 EER comparison for genuine and synthetic data

Data type	*Equal Error Rate* (%)
Genuine data	7
Modified seed profile	8
Multi-profile crossover	19
AI players	7
Parameterised design	9

Crossover-based approach did not show good results, which can be explained by the random nature in which multiple profiles are combined during the crossover process.

4.5 Spoofing experiments

Approaches to spoofing behavioural biometrics are similar to those for physical biometrics but with some domain specific variability. Replay attacks are very popular since it is easy to record an individual's voice or copy a signature. Human mimicking or forgery is also a very powerful technique with experts consistently breaching security of signature-based or voice-based authentication systems.

Additionally in the domain of behavioural biometrics it is possible for a parameterised computer generated model to perform the mimicking/forgery of the biometric sample. Such computer produced models of behaviour parameterised with observed target user data steadily improve in their performance.

In order to create an artificial poker player with the strategy of a particular user a number of steps need to be followed. First a long term statistical profile for a large number of players needs to be obtained. We have written special software which observes the game and records every player's action in an individual behavioural profile.

Alternatively this can be easily accomplished as services exist which sell such information for a fee, examples being Poker-Edge.com (2006) and pokerprophecy.com (Pokerprophecy, 2006). These companies have special purpose computers monitoring online casinos around the clock recording every hand of poker played along with actions of individual players and financial outcome. By analysing statistical data provided by poker-edge.com we were able to reverse engineer the strategy employed by different human poker players. Table 10 demonstrates the sample of statistics collected by poker-edge.com along side the analysis of usefulness for those strategy descriptors.

Table 10 Description of key statistics

Statistic	*Description*	*Analysis*
VP$IP	The percent of hands a player voluntarily puts money into the pot (PreFlop). Small blind completions count, Big Blind checks do not count. Roughly PreFlop Call%+Raise%	This stat is the number one indicator for how loose or tight a player is. Higher than 33% loose, and lower than 18% as tight
PreFlop raise	The percent of time a player raises pre flop	PreFlop Aggression/Passiveness. 5% is a median. The higher a player is above 5% the more aggressive, and the lower below 5%, the more passive
PostFlop aggression	The player's combined aggression rating for the Flop, Turn and River. (Bet% + Raise%)/Call%	1.5 is a median for PostFlop Aggression. Players that are much higher than this are very aggressive, and players that are much lower are very passive
Flops seen	The percent of hands a player sees the Flop. (FLseen/HandsPlayed) *100	Another indicator for PreFlop tightness/looseness
Turns seen	The percent of hands a player sees the Turn. (TUseen/HandsPlayed) *100	An indicator for Flop tightness/looseness.
Rivers seen	The percent of hands a player sees the river. (RIseen/HandsPlayed) *100	An indicator for Turn tightness/looseness
Showd-owns seen	The percent of hands a player sees a showdown. (SDseen/HandsPlayed) *100	An indicator for River tightness/looseness

Source: Poker-Edge.com (2006)

A target player's percent of hands for which he voluntarily puts money into the pot is one of best indicators as to what type of player he is. By interpreting this number it can be determined how tight or loose the player is, and what types of cards he is likely to play. To get an idea of what types of hands correspond to what percentage level we can utilise Table 11.

A player's strategy can be reversed engineered from statistical observations. We will use an example from poker-edge.com's statistical analysis page. Suppose we have a player with VP$IP of 18%, and a PreFlop Raise of 3.5%. To make 18% for VP$IP, this player is likely playing Big Pocket Pairs, Big Cards, Other Broadway Cards (suited), and half of the Other Broadway Cards (unsuited). If we add the percentages together, we get 2.26 + 6.03 + 1.51 + 2.26 = 12.06%. Assuming this player

calls about half of the small blinds, that adds another 5%. We arrive at 17.06% which is very close to his VP$IP of 18%. His PreFlop Raise of 3.5% indicates that he is probably only raising Big Pockets, and AK (Poker-Edge.com, 2006).

Table 11 Pre-flop indicators analysed

Group	*Hands*	*Number of Combinations*	*Percentage of seen*
Big pocket pairs	AA, KK, QQ, JJ, TT	6 + 6 + 6 + 6 + 6 = 30	2.26
Big cards	AK, AQ, AJ, KQ, AT	16 + 16 + 16 + 16 + 16 = 80	6.03
Other broadway cards (suited)	KJs, KTs, QJs, QTs, JTs	4 + 4 + 4 + 4 + 4 = 20	1.51
Other broadway cards (unsuited)	KJ, KT, QJ, QT, JT	12 + 12 + 12 + 12 + 12 = 60	4.52
Mid pocket pairs	99, 88, 77, 66	6 + 6 + 6 + 6 = 24	1.81
A–x suited	A9s, A8s, A7s, A6s, A5s, A4s, A3s, A2s	4 + 4 + 4 + 4 + 4 + 4 + 4 + 4 = 32	2.41
Suited connectors	T9s, 98s, 87s, 76s, 65s	4 + 4 + 4 + 4 + 4 = 20	1.51
Low pocket pairs	55, 44, 33, 22	6 + 6 + 6 + 6 = 24	1.81
A–x (not suited)	A9, A8, A7, A6, A5, A4, A3, A2	12 + 12 + 12 + 12 +12 + 12 + 12 + 12 = 96	7.24
Small blind calls	N/A	N/A	10.0

Source: Poker-Edge.com (2006)

After obtaining statistical measurements of the player's style and performing analysis similar to the one just described we were able to obtain information sufficient to program an artificial poker player with a strategy similar to that used by the target player. For each human player we had two statistical profiles collected from separate game sets. One was used for training artificial player and the other was used for verification experiments. The two data sets were completely different and had no overlap of any kind.

Our implementation of poker bots was once again done using the statistical package known as Online Hold'em Inspector (Online, 2006). We were able to generate a set of 50 artificially intelligent poker players with observable actions mimicking those of human poker players participating in our study.

Because we have adjusted each artificial poker player to act just like its human counterpart the resulting statistical profiles look almost identical to a human eye. We have essentially stolen the behavioural identity of our human poker players and have given it to artificially intelligent programs to mimic. In other words we have obtained a set of parameters for a strategic behaviour and have passed it on to a generative model to synthesise the desired behaviour. Not surprisingly we have obtained very good results for our verification experiments.

To get the desired statistical profiles from artificially intelligent poker players we had them play against each other for a minimum of 10,000 poker hands, which is probably equivalent to 100 of hours of human play. This was done so we could obtain the long term statistically consistent behavioural profiles. With a set of 50 genuine players and 50 spoofed profiles we performed 100 verification comparisons.

We used a decision threshold obtained in finding the best possible FAR on true identity and random (non-spoofed) impostor tests (Yampolskiy, 2006; Yampolskiy and Govindaraju, 2006a, 2007). Similarity of each artificial profile was compared to that of a human profile it was modelled after and against another human profile which was randomly chosen from the set. All 50 artificial profiles were positively verified as profiles of target users they were spoofing, giving us 100% FAR if we keep in mind that we are comparing profiles from a human and a bot. Five profiles were incorrectly positively verified than compared to a randomly chosen human profile. This can be explained by a significant degree of similarity between playing styles of some people. The system used in the experiment gives a FRR of about 8% if only real human profiles are submitted. Our experiments show that spoofing behavioural biometrics is a definite possibility (Yampolskiy, 2008).

Our experiments show that with respect to strategy-based biometrics it is possible to secretly observe the target user during play, generate a statistical profile of his actions and train a behaviour generating model to mimic target's behaviour. This is equivalent to stealing of the individual's behavioural identity and some measures need to be taken to prevent this from happening, particularly as similar approach can be used in domains beyond game networks.

5 Conclusions

A number of conclusions can be drawn from the results of our experiments. First the poker player style measure used by human experts, 2D style measure, is not well suited for use in behavioural biometric systems. It is not capable of coping with insufficient amount of data in high-dimensionality behavioural profiles and is really only suitable for describing the four basic types of poker players encountered in the poker literature (Schoonmaker, 2005). Regardless of the type of profile representation used in the experiments it was the worst performing similarity measure outperformed even by the general similarity measure functions.

Examined general similarity measure functions showed an acceptable profile verification performance with Euclidean and Manhattan distances being indistinguishable from each other in terms of their accuracy. Mahalanobis distance function performed slightly worse possibly as a result of the normalisation procedure which took into

account variance of the data in each profile. Since the degree of variance in each user profile is different it is possibly that normalisation was not evenly distributed and so produced a slight decrease in the performance of this general similarity measure function.

Customised weighted Euclidean measure function specifically designed for the domain of poker-based behavioural profiles showed the best performance on all types of data representation. Heavier consideration for pre-flop player's actions allowed this similarity measure function to pick out the fundamental tendencies of the player's strategy and as a result improve algorithms verification accuracy to as low as the 7% EER for the behavioural profiles enhanced with temporal and spatial information (Yampolskiy and Govindaraju, 2006a).

The use of biometric technologies is growing at an increasing rate. In order to properly test such systems we need a consistent supply of readily available biometric data. Synthetic data generation provides a time and cost effective way of obtaining benchmark and test data not just for biometric systems but also for security and intrusion detection systems in general (Barse et al., 2003; Chinchani et al., 2004; Garg et al., 2006; Kayacik and Zincir-Heywood, 2005; Lundin et al., 2002; Debar et al., 1998; Rossey et al., 2002).

Our spoofing experiments demonstrate that with respect to game-strategy biometrics it is possible to secretly and automatically monitor the target user during play in an online casino, generate an accurate statistical profile of his actions and train an artificially intelligent poker playing program to mimic target player's behaviour. This is equivalent to stealing of the individual's behavioural online identity and is a matter of serious concern for both privacy advocates and security specialists.

Acknowledgement

This work was supported in part by the National Science Foundation Grant No. DGE 0333417 "Integrative Geographic Information Science Traineeship Program", awarded to University at Buffalo.

References

Badizadegan, M. (1999) *Texas Hold'em Flop Types*, Goldstar Books, Los Angeles, California.

Baggenstoss, P.M. (2004) 'Class-specific classifier: avoiding the curse of dimensionality', *Aerospace and Electronic Systems Magazine*, Vol. 19, No. 1, pp.37–52.

Barse, E.L., Kvarnstrom, H. and Jonsson, E. (2003) 'Synthesizing test data for fraud detection systems', *19th Annual Computer Security Applications Conference (ACSAC 2003)*, Las Vegas, Nevada, p.384.

Brandt, K. (2005) *Player Profiling in Texas Holdem*, Available at: http://www.soe.ucsc.edu/~kbrandt/pubs/prof.pdf, Retrieved 29 May.

Chinchani, R., Muthukrishnan, A., Chandrasekaran, M. and Upadhyaya, S. (2004) 'RACOON: rapidly generating user command data for anomaly detection from customizable templates', *Annual Computer Security Applications Conference*, Tucson, AZ, pp.189–202.

Debar, H., Dacier, M., Wespi, A. and Lampart, S. (1998) *An Experimentation Workbench For Intrusion Detection Systems*, IBM Research Report RZ2998.

Garg, A., Sankaranarayanan, V., Upadhyaya, S. and Kwiat, K. (2006) 'USim: a user behavior simulation framework for training and testing IDSs in GUI based systems', *39th Annual Simulation Symposium (ANSS 06)*, Huntsville, AL, pp.196–203.

Goldberg, D.E. (1989) *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley Publication, Reading, MA.

Kayacik, G.H. and Zincir-Heywood, A.N. (2005) 'Generating representative traffic for intrusion detection system benchmarking', *IEEE CNSR 2005*, Halifax, Canada, pp.112–117.

Lee, K. and Park, H. (2003) 'A new similarity measure based on intraclass statistics for biometric systems', *ETRI Journal*, Vol. 25, No. 5, pp.401–406.

Lundin, E., Kvarnstrom, H. and Jonsson, E. (2002) 'A synthetic fraud data generation methodology', *Proceedings of the 4th International Conference on Information and Communications Security (ICICS 2002)*, Lecture Notes in Computer Science, Singapore, Springer, pp.265–277.

Ma, Y., Schuckers, M. and Cukic, B. (2005) 'Guidelines for appropriate use of simulated data for bio-authentication research', *4th IEEE Workshop on Automatic Identification Advanced Technologies (AUTO ID)*, Buffalo, NY, pp.251–256.

Makthal, S. and Ross, A. (2005) 'Synthesis of iris images using Markov random fields', *Proc. 13th European Signal Processing Conference (EUSIPCO)*, Antalya, Turkey.

Maurer, M. (2005) *IRC Database*, Available at: http://games.cs.ualberta.ca/poker/IRC, Retrieved 19 May.

Online (2006) *Online Holdem Inspector*, Available at: http://www.pokerinspector.com/, Retrieved 2 May.

Orlans, N.M., Buettner, D.J. and Marques, J. (2004) 'A survey of synthetic biometrics: capabilities and benefits', *Proceedings of the International Conference on Artificial Intelligence (IC-AI'04)*, CSREA Press, McLean, VA, pp.499–505.

Poker-Edge.com (2006) *Stats and Analysis*, Available at: http://www.poker-edge.com/stats.php, Retrieved 7 June.

Pokerprophecy (2006) Available at: http://www.pokerprophecy.com, Retrieved 26 September.

Rossey, L.M., Cunningham, R.K., Fried, D.J., Rabek, J.C., Lippmann, R.P., Haines, J.W. and Zissman, M.A. (2002) 'LARIAT: Lincoln adaptable real-time information assurance testbed', *Aerospace Conference Proceedings*, Anaheim, CA, pp.2671–2682.

Schoonmaker, A.N. (2005) *The Psychology of Poker*, 1st ed., Two Plus Two Publishing, Henderson, NV. Software (2005).

Software (2005) *Software – Statistics*, Available at: http://www.ultimatebet.com, Retrieved 4 May.

Stats and Analysis (2006) *Poker-edge.com*, Available at: http://www.poker-edge.com/stats.php, Retrieved 7 June.

Sturn, A. (2000) *Cluster Analysis for Large Scale Gene Expression Studies*, Masters Thesis, The Institute for Genomic Research, Rockville, Maryland, USA.

Sumi, K. and Matsuyama, T. (2005) 'Privacy protection of biometric evaluation database – a preliminary study on synthetic biometric database', *Japan-Korea Joint Workshop on Frontiers of Computer Vision*, Gwangju, Korea, pp.189–194.

Wikipedia (2006) *Mahalanobis Distance*, Available at: http://en.wikipedia.org/wiki/Mahalanobis_distance, Retrieved 22 August.

Yampolskiy, R., Anderson, P., Arney, J., Misic, V. and Clarke, T. (2004) 'Printer model integrating genetic algorithm for improvement of halftone patterns', *Western New York Image Processing Workshop (WNYIPW)*, Rochester, NY.

Yampolskiy, R.V. (2006) 'Behavior based identification of network intruders', *19th Annual CSE Graduate Conference (Grad-Conf2006)*, Buffalo, NY.

Yampolskiy, R.V. (2008) 'Mimicry attack on strategy-based behavioral biometric', *5th International Conference on Information Technology: New Generations (ITNG2008)*, Las Vegas, Nevada, pp.916–922.

Yampolskiy, R.V. and Govindaraju, V. (2006a) 'Use of behavioral biometrics in intrusion detection and online gaming', *Biometric Technology for Human Identification III. SPIE Defense and Security Symposium*, Orlando, Florida.

Yampolskiy, R.V. and Govindaraju, V. (2006b) 'Similarity measure functions for strategy-based biometrics', *International Conference on Signal Processing (ICSP 2006)*, Vienna, Austria, pp.174–179.

Yampolskiy, R.V. and Govindaraju, V. (2007) 'Dissimilarity functions for behavior-based biometrics', *Biometric Technology for Human Identification IV. SPIE Defense and Security Symposium*, Orlando, Florida.

Yampolskiy, R.V. and Govindaraju, V. (2008) 'Generation of artificial biometric data enhanced with spatial-temporal and environmental information', *Biometric Technology for Human Identification V. Symposium*, Orlando, Florida, pp.837–842.

Yanushkevich, S., Stoica, A., Srihari, S., Shmerko, V. and Gavrilova, M. (2004) 'Simulation of biometric information: the new generation of biometric systems', *Proc. BT2004 Int'l Workshop on Biometric Technologies*, Calgary, AB, Canada, pp.87–98.