User Authentication via Keystroke Dynamics: An Artificial Immune System Based Approach

Kenneth Revett

British University in Egypt
Faculty of Informatics and Computer Science
El Sherouk City
11837 Cairo, Egypt

ken.revett@bue.edu.eg

Abstract. Keystroke dynamics is a behavioral biometric that is based on *how* a user enters their login details. In this study, a set of eight attributes were extracted during the course of entering login details. This collection of attributes was used to form a reference signature (a biometric identification record) for subsequent authentication requests. The algorithm for the authentication process entails the deployment of an artificial immune based approach. The approach uses self-reactivity to discriminate self from non-self from the enrollment data. During the classification task, the system relies on deploying a pool of non-self reactive antibodies to perform a very general classification task. The results of this study indicate that the error rate is less than 5% in many cases.

Keywords: artificial immune systems, auto-immune sensitivity, behavioral biometrics, biologically inspired computation, di-graphs, keystroke dynamics

Introduction

Keystroke dynamics is a particular instance of a behavioral biometric that captures the typing style of a user. The dynamics of a user's interaction with a keyboard input device yields quantitative information with respect to dwell time (how long a key is pressed) and time-of-flight (the time taken to enter successive keys). By collecting the dynamic aspects acquired even during the login process, one can develop a model that captures potentially unique characteristics that can be used for the identification of an individual. To facilitate the development of the model of how the user enters their details, an enrollment phase is required, when the user is asked to enter his/her login id/password until a steady value is obtained (usually limited to 10-15 trials – but this is implementation dependent). Once this data has been collected, a reference 'signature' is generated for this user. The reference signature is then used to authenticate the user account on subsequent login attempts. The user with that particular login id/password combination has their keystroke dynamics extracted and then compared with the stored reference signature. If they are within a prescribed tolerance limit – the user is authenticated. If not - then the system can decide whether to lock up the workstation - or take some other suitable action.

In this study, a preliminary investigation of the application of the artificial immune system (AIS) is applied as a means of performing the authentication process. The AIS derives its inspiration from biological immune systems, and seems to serve as a natural approach to user authentication generally. In intruder detection systems, the AIS approach has been deployed with some success to identify. It is a natural approach that is designed to distinguish self from non-self. In the case of intrusion detection, non-self is a foreign entity which is attempting to gain access to the computer resources. The immune system acts as a surveillance mechanism, constantly vigilant looking for entities that are not typical of the host environment. Likewise, with an application to keystroke dynamics, the AIS approach monitors the keystroke dynamics of the user in order to determine whether they belong to the host (equally suitable for both static and continuous authentication modes). To date, there is a dearth of publications in this domain. This paper presents an overview of our system, which implements many aspects of the AIS approach, along with some preliminary results. We start by presenting a brief perspective on keystroke dynamics and artificial immune systems.

1.1 Introduction to keystroke dynamics

Gaines was the first to report the results of a properly controlled study in the field of keystroke dynamics [1]. His study examined the typing patterns of seven professional typists – with the goal of determining if there were unique typing styles that could be used to distinguish between the typists. Although the results were not on par with current techniques, the deployment of digraphs – the time taken to enter two successive characters was a breakthrough. Joyce & Gupta presented in 1990 [2] an algorithm based on digraphs – but with a larger cohort and the results were significantly improved with respect to the Gaines study [1]. In 1997 Monrose and Rubin use the Euclidean Distance and probabilistic calculations based on the assumption that the latency times for one-digraph exhibits a Normal Distribution

[3]. In 2000, they also present an algorithm for identification, based on the similarity models of Bayes, and in 2001 they present an algorithm that uses polynomials and vector spaces to generate complex passwords from a simple one, using the keystroke pattern [4]. Various fuzzy logic algorithms have been applied – mapping the variability in typing patterns to a fuzzy concept. For instance, Hussein et al [5] used a combination of fuzzy clustering algorithms - obtaining an error rate (EER) of approximately 5-10%-depending on the number of samples they acquired per login id/password combination. Another study [6] employed a fuzzy rule set in order to classify login id/password combination with somewhat better success than Hussein - although they report only their preliminary results. Techniques based on neural networks have been explored - focusing on ART-2 and multi-layer perceptrons trained with the back-propagation algorithm. For instance, Obadiat provides data that suggests that the error rate can be reduced to approximately 2.4-4.2%, depending on the exact pre-processing performed using a non-standard neural network [7]. Other researchers have also applied neural networks (using standard back-propagation) to keystroke dynamics, generating error rates on the order of 2-4%[8]. Other machine learning approaches, based on support vector machines (SVM) have been used to address the classification problem presented by keystroke dynamics. Sung et al. have applied SVM to this domain, reporting an error rate of approximately 8-10% [9]. Bergadano et al. have employed an edit distance approach to user authentication. The edit distance is a measure of the entropy between two characters (in this case trigraphs) contained within two or more strings [10]. Revett et al. have used the rough sets algorithm to extract rules that form models for predicting the validity of a login ID/password attempt [11]. Lastly, the use of various bioinformatics based approaches such as motifs and multiple sequence alignments have yielded success with respect to user authentication and identification [12], [13]. In the next section, we describe the fundamentals of the biologically inspired computational approach deployed in this study: the artificial immune system (AIS) approach.

1.2 Introduction to Artificial Immune Systems

A variety of computational models have their roots in biological processes: these include artificial neural networks, genetic algorithms, particle swarm, and more recently artificial immune systems. The attraction of these biological processes is probably derived from their apparent information processing capabilities. These systems have the innate ability to perform classification based activities, are distributed, and are adaptable. In addition, these biological computation capacities appear to operate automatically and autonomously – a very desirable yet debatable topic. In the present case, the artificial immune system contains these same properties: they are distributed, adaptable systems with memory that provide the organism with the basic ability to distinguish self from non-self [13]. The operational goal of the immune system is to eradicate any non-self matter that enters the organism's biological domain. The end result of this process is the destruction of that which is deemed to be non-self through a series of chemical reactions.

In the present context, we would like to produce an AIS that is able to perform the essential functions of biological immune systems: distinguishing self (authentic users) from non-self (imposters). In the keystroke dynamics domain, self is not a fixed point, but rather a set of entries that the user has successfully been authenticated with. We generally do not repeat the same typing pattern precisely. As a matter of fact, such perfect fidelity may alert a 'replay attack' module that may reject the authentication attempt outright! So variation is expected – the issue is how much can we incorporate into our authentication system in order to maintain false acceptance and false rejection rates within desirable levels?

An artificial immune system is simply an implementation of a biological immune system *in silico* basically. It must implement the salient features of the biological immune system, at some level. In this work, the concept of distinguishing self from non-self is implemented in a fashion that certainly has biological realism, but is not complete in all levels of detail. The AIS presented here implements the antigen-antibody concepts which form the cornerstone of immunology. The antigen is the foreign object which may be recognised by the immune system – if it is, it will be destroyed if possible. The foreign object is initially encountered by the human host through interactions with host generated molecules termed antibodies. Human immune systems contain literally billions of antibodies each capable of interacting with a bewildering array of antigens. In some instances, these antigens may be part of the host, yielding an auto-immune response. This is considered a mistake so to speak, but yields serious repercussions as people with arthritis and related disease well know. More typically, the circulating antibodies identify a substance as foreign, which is truly foreign, and mounts an attack which attempts to destroy the host. In this sense, the immune system operates in a distributed and parallel fashion. This feature must clearly be incorporated into an AIS model, which is true in the current case. Lastly, the system must be adaptable if it is to respond to antigens it has not been previously exposed to – that is like ANNs, it must be able to generalize. This ability should manifest over variable time windows – locally by differentiating into variants that can attack antigens that are also adaptable, and also over the long term, so that a repeat attack will be acted upon more rigorously [14]. This long term aspect of the immune system simply indicates that it has actually *learned* something from the previous interaction. The distinction is like guessing the answer to a question by shouting out random answers, compared to solving the problem analytically. The next section describes the basis of the experiment and describes how an implementation of the AIS was deployed.

2. Methods

There were 20 participants in this study, all from a computer science undergraduate students from a Polish University. The users were provided with 8-character login IDs and 8-character passwords, generated randomly by a computer programmer. The characters consisted of all upper and lower case alphabetic characters and the digits. The enrollment process required users to enter their login ID/password 10 times successfully. Each participant enrolled on to a single machine located on campus - for both enrollment and subsequent login attempts. After successfully enrolling (10 trials), the participants were asked to perform 100 self-logins (for FRR) and 100 attacks on other accounts (5 for each account including their own, which was not utilised, for FAR data). The following regime was used for non-enrollment logins: each participant was asked to self-login 100 times over a 7-day period. Therefore, each participant logged into their own account approximately 15 times/day. In addition, students were instructed to login at 3 different periods of the day: morning (09:00-10:00), noon (12:00-13:00) and early evening (17:00-18:00). At each period, students were asked to either perform self-login or non-self login 5 times. This simulates the way users would normally access their computer systems, logging in at various periods during the course of a workday.

The data that was extracted from the login ID and password combination were simply *keypress* di-graphs – that is, the time (in ms) between depressing successive keys. There were a total of 14 di-graphs in the login IDs and passwords that were recorded, and stored in a vector of floats (normalised to [0..1]), along with the actual di-graph characters. This vector of di-graph times serves as the shape space of the authentication attempt (the antigen). The enrollment process provides a sample of self logins, which serve as the means of providing a reasonable set of examples of self. The enrollment samples then serve as the basis for fine tuning the immune system such that it is able to differentiate self (those samples similar to the enrollment entries) from non-self (samples that differ significantly from the enrollment samples). The enrollment data di-graph vector becomes the reference vector for each user of the system. More specifically, a random set of antibodies is produced (1,000 in this study), with a shape space identical to that of the enrollment vectors – containing 14 floats, each element of which is assigned a random number on the interval of [0..1]. The antibodies are then allowed to match up in a lock and key fashion with each of the enrollment vectors and a matching score is obtained. This matching score describes the affinity between the antibody and the antigen. The algorithm for matching (DOC) is the following: the antibody (Ab) and antigen (Ag) are aligned and the sum of the vectors elements, one by one are computed according to the following equation:

$$(1)\quad \text{Degree of Complimentarity} = \text{abs}(((1.0 - Ag) - Ab))$$

The DOC ranges from 0 (perfect match) to 1.0, perfect mismatch (note the use of the absolute value in equation 1). The global matching score (GMS) is simply the sum of the individual di-graph DOC values, which will range from 0 to the number of di-graphs.

$$(2)\quad \text{Global Matching Score} = \Sigma \text{ DOC values}$$

As a first processing step to generate usable antibodies, those that react with self must be eliminated in order to prevent auto-immune reactions. After the 10 enrollment entries are generated, they are exposed to the antibody pool and the GMS values are computed for with respect to each antibody. Those antibodies with GMSs above a threshold, α, are considered to be activated by 'self' antigens (i.e. the authentic user's enrollment data), and are removed from the antibody pool. Only those antibodies that are non-reactive to the enrollment data are kept for subsequent use in the AIS system. Note that this process occurs for each user in the system, as each will produce their own enrollment data. This activity engenders the maturation of the immune system, in which lymphocytes in the thymus become activated by a process of culling out hyper and self-reactive B-cells.

Once a processed set of antibodies has been produced, users will enter the authentication (testing) phase, where they will be asked to enter their login ID/password details. The same features will be extracted during authentication, and utilised for the authentication purposes. When the user attempts to authenticate, the digraphs will be collected and form their antigen surface that has the same structure as those collected during enrollment. The authentication sample will be exposed to the pool of primed antibodies and the GMS will be obtained for each antibody. If the score is above a threshold (the same α deployed in the priming stage), then the authentication attempt is classified as *rejected*, otherwise it will be *accepted*. The reason for this decision is that if the antigen (the authentication sample) is identified by the antibodies (via the GMS score being above threshold), then it must be significantly different from the enrollment samples, as the antibodies were selected based on their low GMS scores. If the authentication attempt is not recognised by the antibody pool (that is the GMS is below the threshold), then this sample is considered to be similar to those contained within the enrollment pool. If we stopped at this point, then each authentication sample is classified as being produced either by the actual owner (doesn't activate the immune system) or by an imposter (activates the immune system). Knowing the actual identity of the person entering the authentic login details, we can calculate the FAR and FRR of this stage in the AIS system. Further, by varying the acceptance threshold, α, we can calculate the equal error rate. The results from this experiment are presented in Table 1.

Now, those login attempts that are generally classified as accepted fall into one of two categories: true positive (TP) and false positives (FP). Likewise, those attempts rejected fall either into the true negative (TN) or the false negative (FN) class, which can be summarised conveniently by a confusion

matrix. The true classification rate can be calculated from this data, which is presented in the confusion matrix below. Note this confusion matrix was calculated from a single login account that was checked for FAR and FRR 100 times, selected randomly from the pool of 20 users.

Although deploying a supervised approach is not ideal (we would like to make the system as unsupervised as possible), the purpose of this study is to examine how large the antibody pool must be in order to acquire sufficiently high classification accuracy. That is, how useful is the antibody self-reactivity selection process in the deployment of an AIS?

To address this question, the classification accuracy was assessed with respect to the number of antibodies, DOC and GMS (the free parameters in the model). The number of antibodies was varied from 100 to 1,000,000 in (10-fold) increments (results presented in table 3). In order to estimate this effect, values for the other free parameters are required first. The acceptance threshold (α) was varied from 0 (requiring a perfect match) down to 0.5 in increments of 0.1, and the resulting classification accuracy was computed. These results are presented in table 1. Lastly, the effect of varying the global matching score was varied, from 0 (perfect match across all antigen component – di-graphs) to 7, which requires a 50% match in the worst case. The effect this induced is presented in table 2.

3. Results

The data presented in this section represent the average value across all users (20 in this study), unless otherwise indicated, in which case a sample was randomly selected. No distinctions with regards to samples were made from day to day testing – all samples for each user ID were pooled together with respect to self logins and imposter logins. Please note that the results are produced via calculation of the classification accuracy using equation 3:

$$AC = (TP + TN)/TP + TN + FP + FN \quad (3)$$

Where TN = true negative, TP = true positive, FN = false negative, and FP = false positive. Also note that these free parameters are interdependent, as discussed in the conclusion section in more detail. Briefly, a search across the entire parameter space was required for each calculation was performed.

Table 1. Classification accuracy as a function of the local free parameter, α, which was varied from 0 (perfect match required) to 0.5 (random chance). Note the values are rounded to 1 decimal place.

0.0	**0.1**	**0.2**	**0.3**	**0.4**	**0.5**
87.0%	92.4%	84.6%	77.2%	73.9%	62.1%

Table 2. Classification accuracy as a function of the global parameter GMS, the global matching score. Values were varied from 0 – perfect match across of di-graphs to 50% match in the worst case (in this case with 14 di-graphs, the lower bound was set to 7).

0	**1**	**2**	**3**	**4**	**5**	**6**	**7**
72.4 %	77.9 %	83.2 %	71.3 %	72.1 %	75.2 %	67.3 %	74.8 %

Table 3. The classification accuracy as a function of the number of antibodies deployed in the maturation phase. Note that the number of antibodies are expressed in log10 notation for clarity purposes (2 to 7). Note these results were obtained using the values of α (see table 1) and GMS in terms of classification accuracy.

2	**3**	**4**	**5**	**6**	**7**
86.4%	86.2%	85.1%	80.3%	74.9%	67.6%

4. Conclusion

The results from this study are clearly preliminary – yet they provide support to the utility of the AIS approach. Allowing antibody maturation alone produced an AIS system with quite reasonable classification accuracies. This approach to AIS is novel, at least in the biometrics domain. There were only 3 free parameters in this model – but they are interdependent, and hence to optimize their values requires a more sophisticated approach than that deployed in this study, which was an exhaustive search over a reasonably sized parameter space (6x8x7) - 336 calculations. In order to refine the search space, a search mechanism such as a genetic algorithm or some other approach could be utilised quite easily. In terms of the results, the size of the antibody pool indicates that approximately 100- 1,000 antibodies provides the highest classification accuracy note (the run-time for 10^5 antibodies was approximately 2 s) The match threshold at the individual di-graph values (α) indicates that about a 20% mismatch tolerance is optimal. Lastly, the global match score (GMS) yielded similar results as the local measure, which may indicate that this parameter may simply propagate the information content at the local level.

This is an exploratory study, and has provided some useful insight into the deployment of antibody maturation with respect to its effect on classification accuracy. Clearly more work needs to be performed such as: i) exploring the state more thoroughly, ii) using local values for the parameters for α and GMS scores (each di-graph α can be set to a unique value) , iii) a more complete implementation of the AIS will be undertaken and iv) a more comprehensive memory capacity.

References

1. Gaines, R. Lisowski, W., Press, S., & Shapiro, N., Authentication by keystroke timing: Some preliminary results. Rand Report R-256-NSF. (1980) Rand Corp.
2. Joyce, R. and Gupta, G.. Identity authorization based on keystroke latencies. *Communications of the ACM*. Vol. 33(2), (1990) pp 168-176.
3. Monrose, F. and Rubin, A. D., 1997. Authentication via Keystroke Dynamics. *Proceedings of the Fourth ACM Conference on Computer and Communication Security*. Zurich, Switzerland.
4. Monrose, F. and Rubin, A. D., 2000. Keystroke Dynamics as a Biometric for Authentication. *Future Generation Computing Systems (FGCS) Journal: Security on the Web*.
5. Hussien, B., Bleha, S. & McLaren, R.,. An application of fuzzy algorithms in a computer access security system. Pattern Recognition Letters,9:39-43,1989.
6. de Ru, W.G. and Eloff, J.. "Enhanced Password Authentication through Fuzzy Logic". IEEE Expert, 12(6), Nov/Dec, pp.38-45, 1997.
7. Obadiat, M.S. & Sadoun, B. A Simulation Evaluation Study of neural network techniques to Computer User Identification, Information Sciences 102, 239-258, 1997.
8. de Oliveira , M. VS, E. Kinto, Hernandez, E.D.M, & de Carvalho, T.C., User Authentication Based on Human Typing Patterns with Artificial Neural Networks and Support Vector Machines, SBC 2005.
9. Sung, K.S. & Cho S., GA SVM Wrapper Ensemble for Keystroke Dynamics Authentication, International Conference on Biometrics, Hong Kong, pp. 654-660, 2006.
10. Bergadano, F., Gunetti, D. and Picardi, C. (2002) 'User authentication through keystroke dynamics', *ACM Transactions on Information and System Security*, 5(4), pp.367–397.
11. Revett K. Magalhaes, S. & Santos, H., Developing a Keystroke Dynamics Based Agent Using Rough Sets, The 2005 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology Workshop on Rough Sets and Soft Computing in Intelligent Agents and Web Technology, Compiegne, France, 19-22 September, 2005, pp 56-61.
12. Revett, K., A Bioinformatics Based Approach to user Authentication via Keystroke Dynamics, International Journal of Control, Automation, and Systems, 7(1), pp. 7-15, Feb, 2009.(ISSN: 1598-6446).
13. Forrest, S., Hofmeyr,S.A., & Somayaji, A., 1997,.Computer Immunology, Communications of the ACM 40 (10): 88.96 (October).
14. Ebner, M., Breunig, H-G, & Albert, J., On The Use Of Negative Selection In An Artificial immune System, GECCO 2002, pp. 957-964, 2002.