***User Authentication by Service Utilisation Profiling***

**Alexandre Aupy & Nathan Clarke**
Network Research Group, University of Plymouth
nclarke@plymouth.ac.uk

## Abstract

In the age of the information society, computer security has become of vital importance in order to maintain the integrity of a system. Many organisations are enforcing new IT policies calling for better security in an attempt to safeguard against misuse of resources and data. These IT polices are beginning to utilise a range of technologies to provide front-line authentication of the user, including fingerprint scanning and iris recognition. These technologies however only really assist in providing improved point-of-entry security. This paper presents a novel authentication technique that is able to provide transparent and continuous authentication of the user based upon the way in which a user interacts with their computer – which applications a person uses and when, and which websites a user browses for instance. Through the application of neural networks, results are very encouraging with an overall Equal Error Rate (EER) of 7%.

**Keywords:** *Biometrics, Security, Neural Networks, Pattern Classification, Human-Computer Interaction*

## 1. Introduction

The Personal Computer (PC) has become an everyday object both at work and increasingly at home. Access to computing resources has typically always been provided by the username and password. Although in principal, passwords can be secure, assuming they are of sufficient length and random in nature, long random passwords comprising of alphanumerical sequences in upper- and lower-case are unlikely to be remembered by users (Denning, 1999). As such, new authentication techniques are continually being developed.

In the last few years, biometric systems such as fingerprint scanning and iris recognition have evolved from science-fiction to trustable security solutions (Nanavati et al., 2002). Along with passwords and tokens they provide a suite of authentication techniques that system administrators can use to further protect their system from abuse. However to date, the majority of computer-based authentication mechanisms provide point-of-entry authentication, which in itself is intrusive, often inconvenient and provides no means of ensuring the identity of the person beyond initial authentication.

This paper describes the construction of a novel biometric which provides a non-intrusive and continuous method to authenticate the user. Through monitoring user's everyday actions, for instance, which applications they open, websites they browse and words they type, it is hypothesised that users can be differentiated and subsequently authenticated. The technique, known as Service Utilisation Profiling utilises these interactions and a pattern classification engine to evaluate the authenticity of the user. This paper presents a study into the feasibility of Service Utilisation Profiling based upon a desktop computing environment. The paper begins by presenting some background information in the field of biometrics, and continues to describe and present the findings of a practical study.

## 2. Background

The use of biometrics has existed for hundred of years in one form or another, whether it is a physical description of a person or perhaps more recently a photograph. Consider for a moment what it is that actually allows you to recognise a friend in the street or allows you to recognise a family member over the phone. Typically this would be their face and voice respectively, both of which are biometric characteristics. However, the definition of biometrics within the IT community is somewhat broader than just requiring a unique human characteristic and describes the process as an automated method of determining or verifying the identity of a person (Woodward et al., 2003). Biometrics systems are typically divided in to two categories dependant upon their underlying characteristic. Physiological biometrics are those based upon classifying the person according to some physical attribute, such as their face, finger or their hand. Behavioural biometrics rely upon a unique behaviour of the person such as their voice or the way in which they write their signature. Service Utilisation Profiling inherently falls within the latter of these categories.

Biometrics all work on the basis of comparing the biometric sample against a known template, which is securely acquisitioned from the user when he or she initially enrolled on the system. However this template matching/pattern classification process gives rise to a characteristic performance plot between the two main error rates governing biometrics. The False Acceptance Rate (FAR), or rate at which an impostor is accepted by the system, and the False Rejection Rate (FRR), or rate at which the authorised user is rejected from the system. The error rates share a mutually exclusive relationship as one error rate decreases, the other tends to increase, giving rise to a situation where neither of the error rates are typically at zero percent (Cope, 1990). Figure 1 illustrates an example of this relationship.

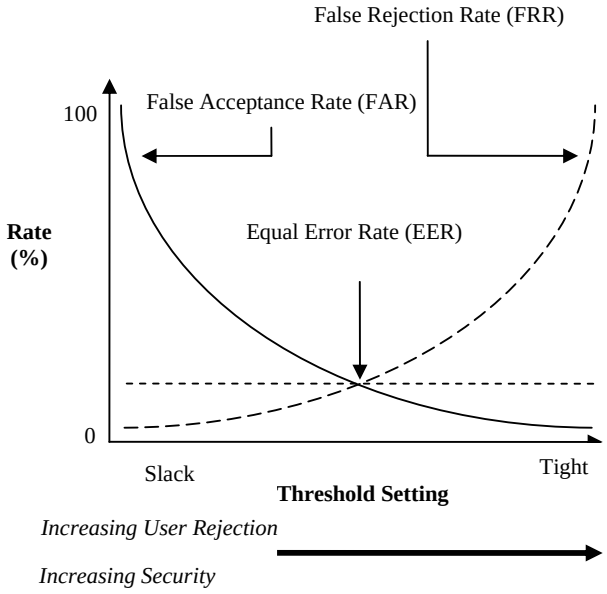


**Figure 1 Mutually exclusive relationship between the False Acceptance & False Rejection Rates**

This leads to a trade-off situation between high security and low user acceptance (due to the fact that the authorised user is being rejected a large proportion of the time) and low security and high user acceptance, to which a decision has to be made about what threshold setting to set that meets both the security requirements of the device and acceptance levels of users.

## 3. Methodology

The objective of this investigation was to evaluate the feasibility of authenticating a user based upon their natural interaction with their desktop computer. Are the way in which we use our computer and what we use it for discernable characteristics for authentication? To this end, the methodology of the trial was split into two phases; the Data Collection & Extraction and Pattern Classification.

### 3.1 Data Collection & Extraction

To enable data collection an application was developed that was capable of capturing keyboard, mouse interactions and system events so a profile of user activity could be built. In order to obtain real input data, this application ran in the background on 21 participant computers for a period of sixty days.

Specifically the *"Logger"* application captured a wide range of interactions and system information. However, due to large volumes of data (in excess of 500MB), it was decided upon for this initial feasibility study to restrict the number of interactions. This would also help in reducing the classification input and subsequent complexity of the pattern classification engine. Table 1 illustrates what features were extracted and utilised in the pattern classification and evaluation stage.

| Code | This action is raised each time… |
|---|---|
| KEY | A full word has been typed in. The word is recorded, along with the title of the window where it has been entered. |
| OPN | A window is opened. The name and class of the window are recorded. |
| CLO | A window is closed. The name and class of the window are recorded. |

**Table 1: Principal user actions captured by Logger**

Low-level Windows API programming was used to perform the capture of the various system events. These actions are then saved due to the development of a proprietary simple database management system, which was required to allow little memory/CPU resources consumption, and on-the-fly compression. The compression technique employed is known as a string indexing technique. It associates a cardinal index with each string (and stores this association in the dictionary, and thus saves only the 4-byte index instead of the full 255-character string.

The extraction of this data was performed by a second application, *"ReadDB"*, which permitted the network designer to extract all the relevant information, illustrated in Figure 2. The output from this application enabled the data to be directly inserted into the pattern classification engine.

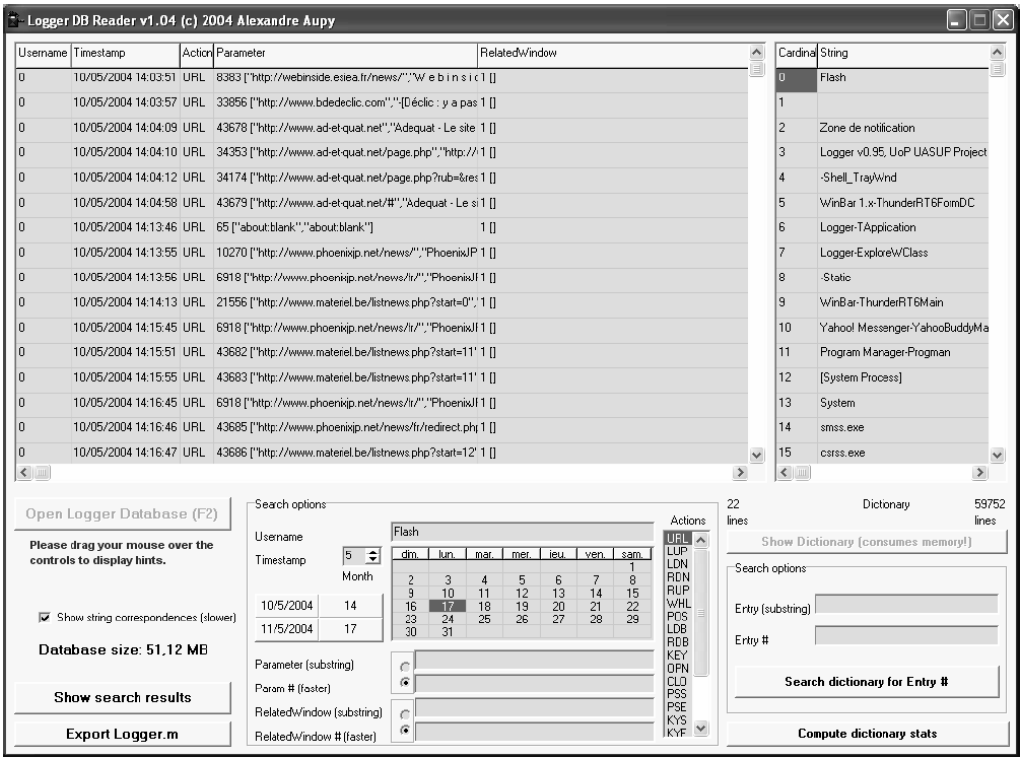


**Figure 2 Screenshot of *ReadDB* – Data Extraction Software**

## 3.2 Pattern Classification

The pattern classification engine utilised in this study is a neural network configuration known as the Feed Forward Multi-Layered Perceptron (FF-MLP). FF MLP networks have particularly good pattern associative properties and provide the ability to solve complex non-linear problems (Bishop, 1995). For more information regarding neural networks refer to Haykin (1999) and Hagan et al. (1996).

An input sample to the neural network is made up of the interactions described in section 3.1, in addition to a timestamp. As it is simplest to present the neural network with numerical data, the timestamps were divided into quarters of an hour, giving rise to a range of input values 0-95 during a day. Initial studies found that four interactions on a single input vector was not sufficient for successful classification. As such, it was decided to utilise 300 repetitions of the aforementioned interactions, giving rise to an input sample size of 1,200. On average these 300 actions corresponded to approximately 10 minutes of user interactivity. Therefore the input to the neural network was 1,200. The size and configuration of the remaining network parameters was somewhat more of a trial and error process. An optimum level found for this study was a two layer FF-MLP with 50 neurons in the first layer and 1 in the output layer. The hyperbolic tangent sigmoid function was utilised as the transfer function due to its non-linear properties, and a back-propagation with momentum and adaptive learning was utilised as the training algorithm.

The collection period of 60 days was split into 4 two-week chunks, with each user taking the turn of the authorised user with the remaining acting as impostors. Due to the quantity of input data the networks could only be designed with 2 users – an authorised user and an impostor, creating a total of 441 networks. Although this is not ideal, for the purposes of a feasibility study it will permit an insight into whether Service Utilisation Profiling is plausible.

The calculation of the error rates and subsequent performance is based upon a smoothing technique, where one classification is based upon three network outputs. Two passes and a fail equal a pass and two fails and a pass corresponds to a fail.

## 4. Results

The results from this study are very encouraging with typical EERs of below 10%, and an overall average EER of 7.1%, as illustrated in Table 2. Figure 4 illustrates the characteristic relationship between the FAR and FRR. Some care must be taken with these figures however as they arguably over-estimate the probable performance you would expect in practice for two reasons:

1. The networks are generated and tested against one impostor user each time and not against all impostors (as would be typical). In practice a single network would have to stop all impostors.
2. The EERs are calculated based upon the complete 60 days of input data. The first fortnight of data however is also utilised in the training of the network. You would expect therefore for the network to perform well against this data.

| User | EER (%) | User | EER (%) |
|---|---|---|---|
| 1 | 3.8 | 12 | 6.9 |
| 2 | 9.4 | 13 | 5.0 |
| 3 | 7.6 | 14 | 5.1 |
| 4 | 3.8 | 15 | 10.7 |
| 5 | 8.5 | 16 | 3.0 |
| 6 | 7.6 | 17 | 6.2 |
| 7 | 6.0 | 18 | 6.0 |
| 8 | 6.1 | 19 | 7.4 |
| 9 | 5.4 | 20 | 6.8 |
| 10 | 10.5 | 21 | 6.2 |
| 11 | 6.7 | **Average** | **7.1** |

**Table 2: Service Utilisation Results**

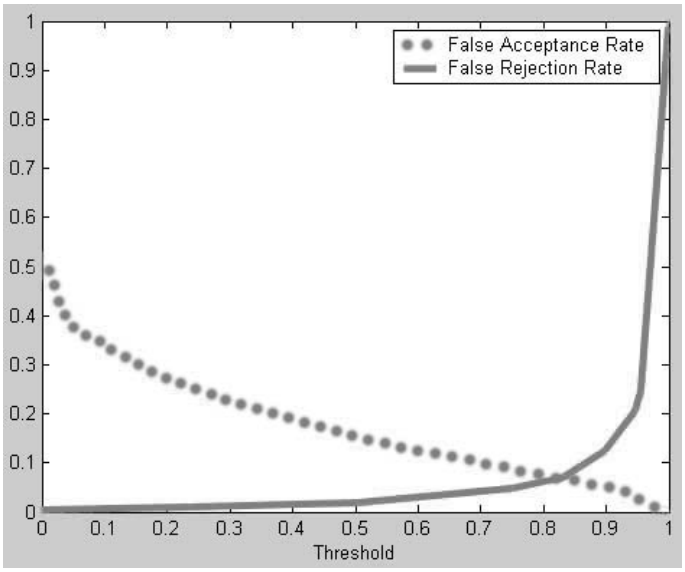


**Figure 4: Average FAR & FRR Plot**

Nonetheless the results are still very encouraging given the relatively small amount of information/interactions that the network is utilising. Further analysis of these problems with the performance show they are not substantial problems. Although training a network with a single impostor is not normal procedure, unfortunately processing constraints restricted using more impostor data. However, tests performed where a network was trained using one authorised user and an impostor, but validated using another impostors' data have shown encouraging results. *User 6* was trained using *User 7* data as impostor, but validated using *User 10 and 15* achieved an EER of 0%.

Figure 5 illustrates the output from the neural network given legitimate and impostor data, split into fortnightly sections. Although the second performance problem certainly has the effect of improving the FAR and FRR, it can be seen from Figure 5 that network performance is still strong given the second fortnight of data, which has not been previously shown to the network. This is one of the primary indicators to the success of Service Utilisation Profiling. What can also be seen from this figure is how the network output becomes noisier as time progresses. This would be expected as people would begin over time to work on different projects. Subsequent re-training of the network would update the network responses to the changes in user behaviour. However, even in the fourth quarter it can be seen the majority of impostor data still resides around zero and with a carefully selected threshold level authentication can still be successful.

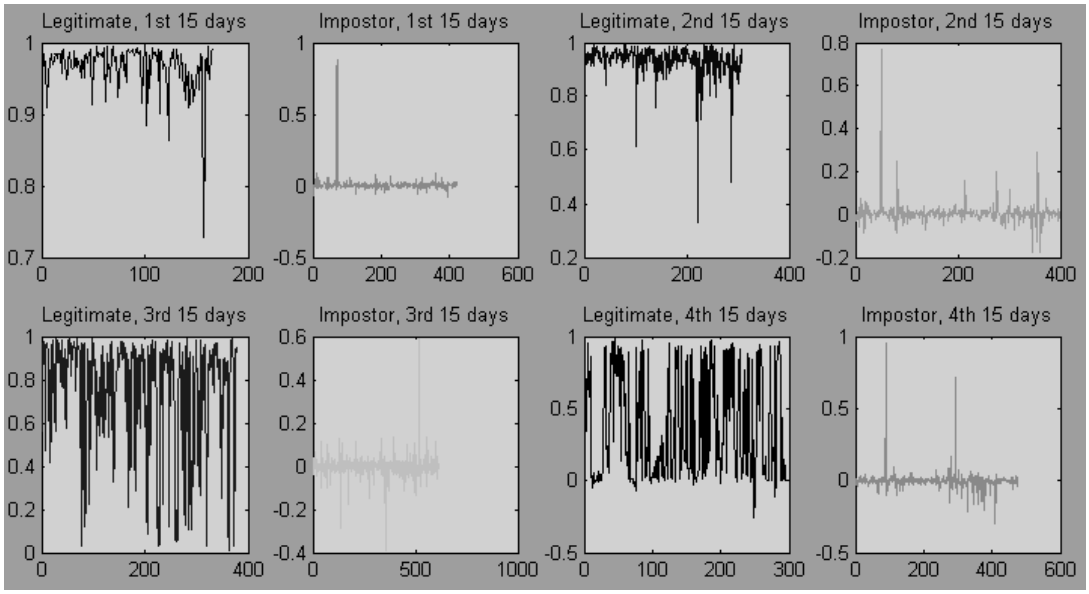


**Figure 5: Example of Network Outputs**

## 5. Discussion

The results presented from this study are not only encouraging but also similar to that expected with other behavioural biometrics (Smith, 2002). It would be suggested due to the variable nature of the input data that this technique not be implemented as the only solution to user authentication, neither should the result be taken as conclusive evidence to lock the user out. However, this technique, in addition to initial point-of-entry authentication would provide an additional continuous confidence measure into the validity of the user. Placed within an appropriate framework, where possible abuse results in the system administrator being made, or a reduction in applications and services available to the user, would assist in ensuring impostors are not able to access sensitive information and services, however, at the same time ensuring a convenient level of security for the authorised user.

The smoothing technique previously employed improves the overall performance of the technique. However, practically under the current configuration of the technique a user could only be authenticated (on average) once every thirty minutes. Evaluating the performance without the smoothing technique enabled achieved an EER of 7.5%, and would permit authentication every 10 minutes. This may or may not be suitable depending upon the requirements of the system administrator. This study however has only presented a feasibility study into the technique and suggested methods by which both security and user convenience can be met. A reduction in the number of samples required before authentication from 300 will ensure the period required for a single authentication on average is what the system administrator feels is acceptable.

An additional practical problem associated with requiring a large number of samples before authentication is circumvention. Given a period of 10 minutes, an impostor may be able to misuse the system for 1 or 2 minutes without necessarily affecting the authentication result. This has two subsequent problems. The first is the system is being abused and secondly all the impostor interaction will get logged as the

authorised users actions. Therefore over time, the profile will begin to incorporate impostor interactions thinking they are legal actions by the authorised user, making it easier for the impostor to continue to abuse the system. Again, reducing the number of samples present will assist in reducing this problem.

## 6. Conclusions & Future Work

The results from this feasibility study have been very encouraging, demonstrating users interactions with their desktop computer are indeed discriminatory. This study has only just begun to explore which interactions can be used, using only four pieces of information. There are therefore a number of areas that would require further research and exploration.

1. To thoroughly investigate the complete range of human-computer interactions and system messages to evaluate which would provide the most discriminative information. The feature extraction process which this addresses is always a key factor in the design of a successful pattern classification process.
2. Provide a mechanism where the neural networks can be presented with impostor data from more users. Although initial tests shown illustrate having a single impostor does not overly affect the results, more general data from more users will help the classifier generalise better to new data.
3. Research different methods of pattern classification, different neural network configurations and different methods of evaluation. For instance, an expert system could be used to evaluate particular interactions during different parts of the day, or after initial log on. For example, after initially logging on, many users were found to perform similar actions every day, such as checking email and reading particular on-line news services etc.
4. Address the privacy issue surrounding Service Utilisation Profiling. Obviously as every interaction with the computer is logged, a number of participants were concerned with the privacy of their data. In this particular study the databases were designed to ensure Instant Message conversations and email were not easily reproducible and the databases were held securely. From a practical perspective this issue will have to be addressed before wide-spread adoption would take place.
5. Research into how Service Utilisation Profiling could be incorporated within current Intrusion Detection Systems (IDS) as a means of monitoring user behaviour and increasing the security of the overall system.
6. Research into how Keystroke Analysis (the ability to authenticate a person by the typing rhythm) (Clarke et al., 2003; Dowland et al, 2004) could be used in conjunction with Service Utilisation to provide an intelligent two-tier multi-modal authentication technique.
7. Research into how Service Utilisation Profiling can be adapted for use on the growing number of mobile devices and handsets currently in circulation.

Although much work is required in the field of Service Utilisation, this study has shown it to be a promising method of authenticating a user continuously and transparently, increasing the current security provided by point-of-entry systems such as the password.

## 7. References


Bishop, C. M. (1995). *Neural Networks for Pattern Classification*. Oxford University Press, New York, USA

Clarke, N., Furnell, S., Lines, B., Reynolds, P. (2003). *"Using keystroke analysis as a mechanism for subscriber authentication on mobile handsets"*. Proceedings of the IFIP SEC 2003 Conference, pp. 97-108.

Cope, B. (1990). *"Biometric Systems of Access Control"*. Electrotechnology, April/May: pp.71-74.

Denning, D. (1999). Information Warfare and Security. Addison & Wesley.

Dowland, P., Furnell, S. (2004). *"A Long-Term Trial of Keystroke Profiling using Digraphs, Trigraph and Keystroke Latencies"*. Proceedings of IFIP SEC 2004, pp. 275-290.

Hagan, M., Demuth, H., Beale, M. (1996). Neural Network Design. PWS Publishing Company

Haykin, S. (1999). Neural Networks: A Comprehensive Foundation (2nd Edition). Prentice Hall

Nanavati, S., Thieme, M., Nanavati, R. (2002). Biometrics: Identity Verification in a Networked World. John Wiley & Sons.

Smith, R. (2002). Authentication. From Passwords to Public Keys. Addison-Wesley.

Woodward, J., Orlans, N., Higgins, P. (2003). Biometrics. Identity Assurance in the Information Age. McGraw-Hill/Osborne.