Keystroke Dynamics: Characteristics and Opportunities

Heather Crawford
Department of Computing Science
Sir Alwyn Williams Building
University of Glasgow
Glasgow, United Kingdom G12 8RZ
Email: crawforh@dcs.gla.ac.uk

***Abstract*—Significant research into the feasibility of keystroke dynamics as a potential biometric authentication method has taken place since the advent of computers. The studies have progressed from examining typing patterns on desktop keyboards using statistical pattern classifiers to mobile keyboards using neural networks as a pattern classifier. The studies do not have a unifying method of comparing results, which limits comparison between the methods presented. Without the ability to compare studies within research areas, future study is limited in its ability to provide important modifications to the work in question. This paper reviews a representative subset of the current research in keystroke dynamics, and provides recommendations on the potential direction of future work in this area. This will provide a set of guidelines that can be followed by researchers intending to do further work in the area of keystroke dynamics.**

I. Introduction

Authentication serves two primary purposes. The first is to correctly identify those users who are authorized to access a resource such as a web page or database and to deny access to those who are not correctly identified. The second is to reassure the user that the resource is, in fact, protected from access by unauthorized users. The latter purpose is subtle, but its loss can result in a systematic loss of security when users attempt to ease the burden that typical authentication methods place on them by reusing passwords, writing them down, or using simple, easy to break passwords. It is the insufficient level of security provided by passwords [1] that has fueled the search for alternative forms of authentication. As we become increasingly dependent on computers for accessing potentially private or sensitive information, the need for a more secure authentication method becomes clear. In recent years, mobile devices such as PDAs and smartphones have grown in popularity, which has led to a rich set of tools that allow us to access this same private or sensitive information on our mobile devices as well as from our desktop and laptop computers [2]. Personal Identification Numbers (PINs) suffer from the same flaws as do passwords [3]. A need for improved authentication methods exists on a wide range of devices, from servers that store personal information, desktop and laptop computers, as well as mobile devices that are used to access information on demand.

Authentication methods fall into three categories: something you *know*, something you *have*, and something you *are* [4, p.28–32]. Passwords and PINs fall into the first category, since they are something that is memorized and provided at the correct time in order to prove your identity and subsequently to access the resource in question. Something you have can take many forms, all of which are collectively called *tokens*. Examples of tokens include ATM cards, smartcards, and the keys to your car. Tokens are often used in conjunction with something you know; for example, ATM cards are used with a corresponding PIN. The third category, something you are, is known as biometrics and can be further divided into two types. *Physiological* biometrics include those characteristics of the human body that can be used to uniquely identify someone. Examples include facial recognition, fingerprints, and retinal scanning [5]. *Behavioral* biometrics use a person's actions or habits to uniquely distinguish them from others. Examples of behavioral biometrics include paper-and-pen signatures [6, p.10], use patterns for devices [7], [8], and keystroke dynamics [9], [10]. This area of research grew from the limitations that traditional physiological biometric methods have: high cost for users to enroll, the requirement for sophisticated (and often expensive) equipment, lack of support for remote access. As such, it is a widely researched area with great potential as mobile devices allow us to access an increasing number of resources quickly and easily. However, even the best research may be overlooked if there is no way to determine whether it has made a significant improvement over previous work.

This paper consolidates the research that has been done to date in the area of keystroke dynamics, which is identifying users based on their typing patterns. The purpose is to examine the methods used, to explore their successes and failures, and to provide a list of recommendations for future work in the area. The research presented here has varied greatly on many levels; significant differences exist in the keystroke attributes used to identify a particular user, in the various pattern matching algorithms used, and the types of error measurement methods employed. Without a unified way of comparing the results of studies in this area, it is likely that this potentially powerful authentication method

978-1-4244-7550-6/10/$26.00 ©2010 IEEE

will simply be overlooked because the differences between the studies makes comparing the results a near-impossible task. Future researchers will be better able to guide their work accordingly if a standard set of research practices is suggested and followed. Furthermore, if such guidelines were applicable to other areas of behavioral biometrics, it would allow comparisons between them that would have benefits throughout other research areas. This paper is intended to provide such recommendations as a first step towards the growth of behavioral biometrics research.

The next section will discuss keystroke dynamics in general to provide a basis for comparisons given in Section III. Section IV will outline the recommendations for further research in this area, and Section V provides a summary of this paper.

II. Background Concepts

The research into keystroke dynamics as an authentication method relies on developing a technique that is robust, inexpensive, and has the potential to be transparent to the user. It was researched as early as 1975 [11] and has been the subject of several patents [12]–[14]. Its feasibility has been examined on desktop computers [15]–[17], web applications [18], and mobile devices such as smartphones and PDAs [10], [19]. Keystroke dynamics are not expected to be unique to each individual since there are likely to be similarities between individuals' typing style, particularly on mobile devices, but it is known to be sufficiently different between users to be useful as a method of verifying a user's identity, as is evidenced by the low error rates seen in the studies summarized in Section III. In essence, keystroke dynamics has the potential to be used as an authenticator but is not powerful enough to be used as an identifier.

The following sections describe the background information needed to meaningfully discuss the results of the studies highlighted in this paper. Section II-A defines the common error rates used in biometrics research; Section II-B describes the two different styles of text entry used. Section II-C and Section II-D outline the measurements used and the ways of using the measurements to match a typing sample to that of a particular user.

A. Error Rates

The main concept behind using keystroke dynamics as an authenticator is to detect the unique patterns that exist when a user types on a keyboard. These patterns can be recognized in many different ways, including statistical classifiers and neural networks. The results of the classification (and thus the performance of the classifier) can be measured using two error rates: the False Accept Rate (FAR) and False Reject Rate (FRR). FAR expresses the likelihood that an unauthorized user (i.e., an impostor) will be granted access to the protected resource, and FRR represents the likelihood that an authorized user will be denied access to the protected resource. As such, the FRR represents an annoyance to the user since being falsely rejected simply means that the user will have to make another attempt to authenticate. FAR represents an intruder who has been granted access to the protected resource, which is a much more significant problem, and thus FAR must be significantly lower than FRR. The acceptable values for FAR and FRR are variable and depend on the sensitivity of the protected resource. The European Standard for Access Control (EN-50133) states that FRR should be less than 1% and FAR should be less than 0.001% for commercially viable authentication systems [20]. There is a distinct trade-off between FAR and FRR — a low FAR in combination with a high FRR represents a high-security environment that is not likely to be accepted by users since they are likely to experience a significant number of re-authentication requests. The opposite environment has a high FAR and a low FRR and represents a low-security environment since it has the potential for accepting many unauthorized users, but rarely requires an authorized user to re-authenticate. In theory, it is possible to achieve a 0% FAR and FRR simultaneously. In practice, however but in practice it is not likely since their relationship is linked: as one rate increases, the other rate tends to decrease [21].

The relationship between FAR and FRR has been described as mutually exclusive since it is impossible to both reject and accept an authentication attempt [21], [22]. While such a statement is clearly true, care must be taken when using such a description for these two related error rates. FAR and FRR also share an inverse relationship - it should not be assumed from the use of the term "mutually exclusive" that no relationship exists between the error rates. The proof of such a relationship lies in the definition of Equal Error Rate (EER).

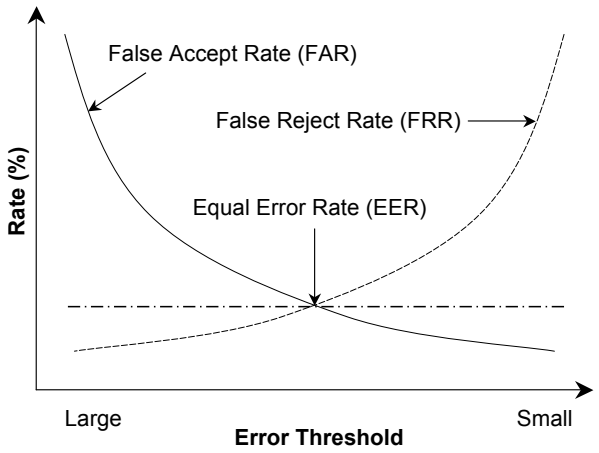

Fig. 1. Relationship between False Accept Rate, False Reject Rate, and Equal Error Rate. Adapted from [23].

While FAR and FRR are valuable metrics for assessing the quality of the biometric system, they do not provide a suitable

measure when comparing two or more systems. A commonly used performance measure for such comparisons is EER, which is defined as the point at which the graphs representing FAR and FRR cross (see Figure 1). Essentially, the study or method with the lowest EER is the most accurate in terms of minimizing both the chance that an unauthorized user gains system access and also the chance that an authorized user will be refused access. While biometrics research has used EER for many years [5] it has only been in the last few years that keystroke dynamics researchers have begun to use it as a comparison method for study results. This makes it difficult to compare the results of these studies meaningfully.

B. Static Versus Dynamic Text Entry

The current keystroke dynamics studies choose either static or dynamic text entry as a basis for comparing a typing pattern to a pattern captured during enrollment. Static text entry requires the user to type a pre-defined text string such as their username and password and compares the keystroke patterns to those gathered at enrollment using the same pre-defined string. Dynamic text entry allows the user to type any text they wish, with authentication taking place by comparing pattern similarities of commonly typed letters.

The studies included in this paper show that authentication based on static text entry are significantly easier to implement and provide much more acceptable error rates. Dynamic text entry is a much better approximation of real-world situations, however, and allows the authentication based on keystroke dynamics to take place in a transparent manner, which increases the likelihood of general user acceptance. As a third alternative, Gunetti and Picardi [24] challenge the standard definition of "dynamic" text entry, stating that most studies perform pseudo-dynamic text entry at best. In their opinion, true dynamic text entry (what they call *free text*) means that the text entered should not be constrained in any way, including allowing the user to choose what text they wish to enter as well as allowing errors, pauses, and other breaks in the flow of text entry. While it is clear that this method more closely mimics typical text entry, and thus the results of the study are more likely to mimic the results given in a real-world environment, the standard definition of dynamic text entry is more common in the literature, and will be used throughout this paper.

C. Metrics

The standard metrics that are used in keystroke dynamics research are *inter-key latency* and *key hold-time*. The former is a measure of the amount of time between when a key is released and the subsequent key is pressed. The latter is a measure of the amount of time between when a key is pressed and when the same key is released. The relationship between these metrics can be seen in Figure 2. Both metrics are common in studies that examine keystroke dynamics on desktop and laptop keyboards, but recent studies have shown that key hold-time is not a discriminating feature on mobile keyboards with thumb-based input [25]. Other metrics have been used in the studies examined in this paper, but, as shown in Section III-A, most of the alternatives are either unusable or are just different wordings of the two standard metrics discussed above.

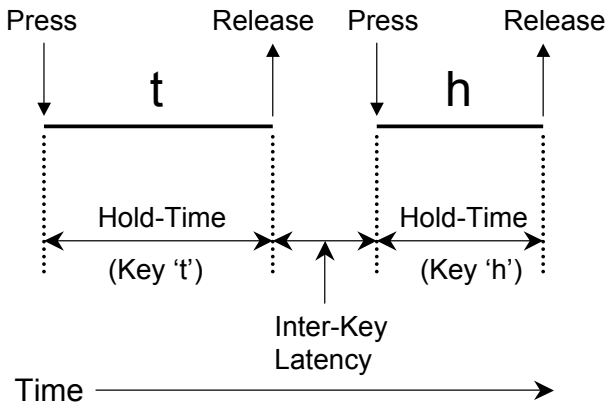

Fig. 2. Relationship between key hold-time and inter-key latency for the digraph "th".

D. Classification Methods

There are several standard classification methods used in keystroke dynamics research. The historical favorite has been pattern classifiers, which are a form of statistical classifier. Commonly used statistical classifiers found in the literature to date include Bayes (including naive Bayes), Mahalanobis distance [10], Hamming distance, Euclidean distance, etc. More recently, studies have begun to use neural networks as a pattern classification method. Common neural network approaches include Feed Forward Multilayered Perceptron Networks (with and without back propagation) [25], Radial Base Function Networks and Generalized Regression Networks. It has been noted by Clarke *et al.* that neural networks are a superior pattern classification method, but that mobile devices (circa 2007) lack the computing power necessary to employ a neural network in situations where the processing is done on the device itself. This means that either statistical methods must be used in mobile device studies, or sample processing and matching must be done on a separate computer that has more resources available. It will be interesting to note whether future studies will attempt to study whether more current mobile devices have sufficient resources to employ neural network-based pattern classifiers, given the ever-increasing memory and resources available on mobile devices.

III. Literature Review

A. Desktop and Laptop Keyboards

Keystroke dynamics studies on desktop and laptop computers have been plentiful since its inception by Spillane in 1975 [11]. The studies differ in duration, number of participants, metrics, and reported error rates, although some comparisons can still be made between the results of the various studies. Table I contains a summary of the main studies performed on desktop and laptop computers. Note that the table is not meant to be exhaustive; instead it contains a representative cross-section of the work to date. Where the Keyboard Type is "Not specified", it is assumed that the study was on a desktop or a laptop computer rather than a mobile keyboard, since studies that use mobile keyboards are generally very explicit in stating what type of keyboard was used. Furthermore, there is no distinction between studies such as Monrose and Rubin (2000) [17] that allowed the user to perform typing on their own keyboard as opposed to studies such as Bleha *et. al.* [16] who required users to type on one particular keyboard in a controlled environment. While this distinction is important if the study results are to be used in a real-world application, controlling the keyboard type does not necessarily affect the study's reported results and thus no distinction is necessary.

The results shown in the table seem to be promising, although it is difficult to choose the "best" results because comparisons between studies are difficult since none of them have reported an Equal Error Rate. In fact, it must be noted that the high error rate percentages seen in the two studies by Monrose and Rubin [17], [29] are due to the fact that they reported what they called Correct Identification Rate. The paper does not relate this rate to the more common FRR and FAR, thus these studies are difficult to compare to those that use the standard measures. Furthermore, there are notable differences in the methods used to calculate FAR and FRR between studies; caution must be used when judging the results reported. For instance, Obaidat and Sadoun [30] report a 0% FAR and a 0% FRR, which is not likely in practice since they are inversely related values. Reporting such results suggests a significant problem with the study, as well as with other studies such as Leggett *et al.* [28].

In these studies, each of the people involved in the study was enrolled as a legitimate user of the system in question, and also performed the role of an impostor. The impostors were given the passwords for the legitimate users and attempted to gain access to the system using them. As such, the impostor's typing patterns were used as training data since each impostor was also a legitimate user. When they played the role of the impostor, the pattern match was to their own legitimate typing pattern, which clearly would not match that of the legitimate user they were trying to imitate. Thus, the system would report a non-matching pattern and disallow access to this impostor. This is notably different from the desired result, which is that the impostor's typing pattern is not in the system already, and the impostor is denied access based on the fact that their (unknown) typing pattern does not match the legitimate user's. This fundamental flaw in the study's design essentially renders useless any results reported. In a real-life authentication system, the typing patterns of the potential impostor would not be known to the authentication system.

The four dynamic text studies report very good results, particularly Gunetti and Picardi [24] and Ahmed *et al.* [32]. Both of these studies show a less than 0.005% FAR (0.0152% in the latter case) and a less than 5% FRR, but the tradeoff is that these methods typically require the user to type several times more characters than the studies that used static text entry comparisons. Gunetti and Picardi required users to type up to 547 characters before they could be authenticated [24], compared to just four strings (first name, last name, username, and password) in Joyce and Gupta's study. The authors reported an average of 24.1 characters total required to authenticate the users [27] - while still reporting similar FAR and FRR rates to Gunetti and Picardi. The good FRR and FAR rates for the static studies supports the idea that using well-known, easy to type, and frequently typed strings give the best results for FAR and FRR - this suggests that using static strings produces better results than using dynamic strings or free text while keeping the number of characters (and thus the level of user frustration) to a minimum. Essentially, there is significant support in the reported results to make a case for continuing the use of static text rather than dynamic, even though dynamic would be preferred in order to increase the transparency (and thus convenience) to the user, which is a stated goal in most of the studies examined here.

Comparisons can also be made between studies that use statistical versus neural network approaches to pattern matching. In general, neural networks tend to produce better results than statistical methods, as can be seen from the data in Table I, although neural networks are highly variable since the number of layers and the number of neurons per layer have a linear relationship with the quality of the results. As the number of layers and neurons increases, so does the complexity of the network and thus the amount of time required to process results. There is a point at which more neurons and layers have a diminishing effect on FAR and FRR, although this has not yet been reported in any of the studies examined here. Also, neural networks have a very high retraining cost - each time a user is added, the network must be retrained, which also increases the amount of processing power required to use these methods. So, despite the somewhat lower quality results seen when using statistical methods, there are still areas where using a statistical pattern matching algorithm may be superior, such as when available processing power is limited. Despite reporting very good results, none of the studies have met both requirements of the

TABLE I
CHRONOLOGICAL REVIEW OF LITERATURE IN DESKTOP AND LAPTOP KEYSTROKE DYNAMICS. NOTE THAT FOR THE SAKE OF BREVITY THE ERROR PERCENTAGES SHOWN ARE THE LOWEST REPORTED WHEN THE STUDY USED MORE THAN ONE CLASSIFICATION METHOD.

Study	Type	Keyboard Style	Metrics			Classifier	Study Size	Error Rates		
			Inter-Key	Hold Time	Other			FAR	FRR	EER
Umphress & Williams (1985) [26]	Static, dynamic	Desktop	✔			Statistical	17	6%	12%	–
Bleha *et. al.* (1990) [16]	Static	Desktop	✔			Statistical	32	0.5%	3.1%	–
Joyce & Gupta (1990) [27]	Static	Desktop	✔			Statistical	33	0.25%	6.67%	–
Leggett *et al.* (1991) [28]	Static, dynamic	Desktop	✔			Statistical	17	5.8%	11.7%	–
Monrose & Rubin (1997) [29]	Dynamic	Desktop	✔			Statistical	31†	84.6%*		
Obaidat & Sadoun (1997) [30]	Static	Desktop	✔	✔		Neural network, statistical	15	0%	0%	–
Cho *et al.* (2000) [18]	Static	Desktop	✔	✔		Neural network	25	1%	0%	–
Haidar *et al.* (2000) [31]	Static	Not specified	✔			Neural network, statistical, fuzzy	–	6%	2%	–
Monrose & Rubin (2000) [17]	Static	Not specified	✔	✔		Statistical	63	92.14%*		
Araujo *et al.* (2005) [9]	Static	Desktop, laptop	✔	✔	✔	Statistical	30	1.89%	1.45%	–
Gunetti& Picardi (2005) [24]	Free (dynamic)	Desktop, laptop	✔	✔		Statistical	205	$<$0.005%	$<$5%	–
Ahmed *et al.* (2008) [32]	Free (dynamic)	Desktop	✔	✔	✔	Neural network	22	0.0152%	4.82%	–

* The authors call this error rate "Correct Identification Rate" and do not distinguish it from FAR, FRR, and ERR.
†This number was reduced from 42 due to timing errors

European Standard for Access Control levels of FAR of less than 0.001% and FRR of less than 1%, with the exception of Obaidat and Sadoun, whose results are somewhat suspect.

Another way the studies varied is in the number of people included. Gunetti and Picardi's study included 205 people – by far the largest reported here. A closer look at their study reveals that only 40 people were used as authenticated users, while the remaining 165 were unauthorized and were used to attempt to be incorrectly seen as authorized users. When considered in this light, the size of the study is similar to the other studies in Table I. The smallest study in this list is 15 people. In some cases, such as with Haidar *et al.* [31] the study size was not reported at all, which makes it impossible to compare the results of this study to those of any other, since study size has a significant impact on reported results [22]. In the majority of the work presented here, the study participants not only performed the role of authorized users, but also of unauthorized users when they attempted to spoof the typing patterns of the others. Such a practice can cause problems if used incorrectly, as noted previously in this paper.

Finally, the studies can be compared in terms of the metrics used to uniquely identify users. Most of the studies use the inter-key latency, which is the measure of the time between when a key is released and the subsequent key is pressed. Many of the studies, beginning with Obaidat and Sadoun in 1997 [30], also used key hold time - the measure of the time between when a single key is pressed and released - with good results. In fact, the combination of hold time and inter-key latency provides significantly better results than using just one of the two. In the case of Araujo *et al.* [9] the authors state that they have used a novel combination of metrics to produce their low error rates, when in actuality they are simply rebranding the typical inter-key latency and hold time. In their paper, they state that they use Down-Down Time (a measure of the time interval between successive keystrokes), Up-Down Time (some feature of keyboard latency - assumed to be the amount of time between a key being released and the subsequent one being pressed), and Down-Up Time (the time interval that a key is pressed). Down-Down Time is the combination of Up-Down and Down-Up times, which are the inter-key latency and the hold-time, respectively. Thus, while they claim to innovate by combining new metrics, Araujo *et al.* have renamed the standard metrics used in the other studies. Their low error rates can be attributed to an adaptation mechanism that is applied after every successful authentication. It is used to continually update the user's typing pattern in order to limit the effects of normal variances

in typing patterns for the same purpose.

B. Mobile Devices

While the desktop and laptop keyboard studies have enjoyed a relatively long history of prior research, keystroke dynamics studies on mobile keyboards are relatively new. The first known study of this type was by Clarke *et al.* in 2002 [21], and as mobile devices become more and more powerful, further studies try to take advantage of the idea of a transparent authentication mechanism by further studying the applicability of keystroke dynamics to mobile devices. There are three main types of mobile keyboards: numeric, in which the user must successively press a number key in order to get a different letter; thumb-based, such as the QWERTY keyboard on a Blackberry device, and soft keyboards such as those on an iPhone. The range of keyboards gives a hint as to the types of devices that are considered "mobile"; they include standard mobile phones, PDAs, as well as smartphones. The majority of the studies in Table II were performed on numeric keyboards, with a few of the studies using thumb-based keyboards. Few studies into the feasibility of keystroke dynamics on a soft keyboard have yet been seen.

Similar comparisons to those made in the previous section on desktop and laptop keyboards can be made when examining the results of mobile keyboard studies. The most striking thing to note is that, with the exception of Campisi *et al.*, all of the studies used neural networks as the pattern matching classifier, which is the likely reason for the high 14.46% error rate seen in that study. All of these studies performed data analysis and pattern matching outwith of the mobile device. The only study to note that the large amount of processing required to effectively use a neural network exceeded the capacity of the mobile device was Buchoux and Clarke in 2008 [10]. It would be interesting to determine whether the latest generation of smartphones now have the necessary processing power to make neural networks viable. If so, then processing the pattern matching on the device itself will make keystroke dynamics a viable, if partial, solution to the problem of remotely authenticating a user within the device itself rather than depending on outside sources that may or may not be available at the time authentication must take place.

Most of the studies in the table used a combination of inter-key latency and hold time, despite the conclusions of Karatzouni and Clarke in [25] that hold time is not a viable metric for mobile devices due to the unique tactile interface provided by a thumb-based keyboard, although they state that further study is necessary to conclusively prove this. The relatively high EER values for the studies seems to support this conclusion. Given that the three main types of mobile keyboard are distinct within themselves as well as when compared to laptop and desktop keyboards, it is unlikely that standard metrics used to date will be sufficient to make keystroke dynamics a viable authentication tool for mobile devices. New metrics such as finger pressure and unique combinations of keystroke latencies such as those suggested by Saevanee and Bhattarakosol [33] and Zahid *et al.* [34] must be created. Their study had the lowest EER at 9%, although it must be noted that it also had the smallest study size at just 10 people, which is likely a limiting factor as well.

The outcome of these studies shows that it is more difficult to uniquely identify a user who is typing on a mobile keyboard. The analysis of this is difficult because a mobile keyboard is smaller and forces the user to adopt a rigid set of movements in order to type, which makes their typing patterns variable and discontinuous [34]. Many of the researchers noted that no one authentication method should be used as the only method of verifying a user's identity [6], [21] since the current methods are, as yet, too error-prone and the typing patterns of mobile users are not distinct enough. If new methods can be innovated to uniquely identify a user, it is entirely likely that the promise of a transparent, user-friendly method of identifying users on mobile devices is within reach. Furthermore, combining this potentially powerful method of authentication with other methods such as voice pattern matching and use patterns will strengthen the proposed authentication method, especially when compared to current knowledge-based mechanisms.

IV. Recommendations

Examination of the relevant studies brings to light the following recommendations for future research in the area of keystroke dynamics, particularly as it applies to authentication of users on mobile devices.

1) Use the high-quality results seen with neural network pattern classification whenever possible. Statistical classifiers, while less computationally intensive, do not provide a strong enough level of pattern classification to support the needs of authentication systems.

2) At a minimum, calculate and report EER , FRR, and FAR for all studies in order to provide a metric for comparisons of past and future work in this area. Other error rates may also be reported for completeness or to show improvements in specific areas. In addition, always report study size.

3) Do not use the same study participants as both authorized and unauthorized users. This leads to significant differences in reported error rates, and does not adequately mimic a real-life unauthorized authentication attempt since in a real-life attempt it is unlikely the impostor's typing pattern will be known to the system.

TABLE II
CHRONOLOGICAL REVIEW OF LITERATURE IN MOBILE DEVICE KEYSTROKE DYNAMICS.

Study	Type	Keyboard Style	Metrics			Classifier	Study Size	Error Rates (%)		
			Inter-Key	Hold Time	Other			FAR	FRR	EER
Clarke *et. al.*, 2002 [21]	Static	Numeric	✔			Neural network	16	1.3%	0%	–
Clarke & Furnell, 2007 [19]	Static	Thumb	✔	✔		Neural network	32	–	–	12.8%
Karatzouni & Clarke, 2007 [25]	Static, pseudo-dynamic	Numeric	✔	✔		Neural network	50	–	–	12.2%
Buchoux & Clarke, 2008 [10]	Static	Numeric	✔			Neural network, statistical	20	0%	2.5%	–
Saevanee & Bhattarakosol, 2009 [33]	Static	Soft	✔	✔	✔	Neural network	10	–	–	9%
Zahid *et. al.* , 2009 [34]	Static, dynamic	Numeric	✔	✔	✔	Neural network, statistical	25	2.07%	0%	–
Campisi *et. al.*, 2009 [35]	Static	Numeric	✔	✔		Statistical	30	–	–	14.46%

4) Keep the processing on the device itself wherever possible in order to avoid having to offload processing to other systems that may or may not be available at the time authentication takes place, and to avoid having to secure both the device and the processing system. This will help avoid the potential privacy implications of duplicating user input on a system that may not be in the user's control.

5) Develop new keystroke characteristics such as finger pressure, unique combinations of digraph frequencies, and error correction rates to support the positive results seen with inter-key latency and hold time metrics. This is especially important in the case of mobile devices since hold time does not seem to provide enough distinction to be used as a metric in a mobile environment. Ensure that these metrics are truly new and do not simply rename existing metrics.

6) Consider the effects of predictive text seen on mobile devices on the uniqueness of keystroke patterns for mobile users.

V. CONCLUSION

Keystroke dynamics has received a significant amount of attention in authentication research circles because it has the possibility of providing a transparent, acceptable method of authentication that can be implemented using existing hardware. The first studies on desktop keyboards were conducted by Spillane in 1975 [11], and since then researchers have continued to progress by performing studies on additional keyboard types, examining new pattern matching algorithms, and identifying new keystroke characteristics. Furthermore, researchers have adopted commonly used comparison methods such as False Accept Rate (FAR), False Reject Rate (FRR), and Equal Error Rate (EER) to quantitatively show improvements in keystroke dynamics methodology. From these studies we have learned that it is unlikely that keystroke dynamics alone will be robust enough to uniquely identify users, but it shows great promise as a part of a larger multimodal biometric authentication method.

ACKNOWLEDGMENT

The author would like to gratefully acknowledge the funding provided by the Scottish Informatics and Computer Science Alliance (SICSA), and editing guidance provided by Karen Renaud and Tim Storer.

REFERENCES

[1] A. Adams and M. A. Sasse, "Users Are Not the Enemy," *Communications of the ACM*, vol. 42, no. 12, pp. 40–46, December 1999.
[2] A. Herzberg, "Payments and Banking with Mobile Personal Devices," *Communications of the ACM*, vol. 46, no. 5, pp. 53–58, May 2003.
[3] N. Clarke, S. Furnell, P. Rodwell, and P. Reynolds, "Acceptance of Subscriber Authentication for Mobile Telephony Devices," *Computers & Security*, vol. 21, no. 3, pp. 220–228, 2001.
[4] R. E. Smith, *Authentication: From Passwords to Public Keys*. Addison-Wesley, 2002.
[5] S. Kung, M. Mak, and S. Lin, *Biometric Authentication A Machine Learning Approach*. Prentice-Hall, 2005.
[6] A. Jain, A. Ross, and S. Prabhakar, "An Introduction to Biometric Recognition," in *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 14(1), IEEE. IEEE, January 2004, pp. 4–20.
[7] N. Clarke, S. Karatzouni, and S. Furnell, *Emerging Challenges for Security, Privacy and Trust*, ser. IFIP Advances in Information and Communication Technology. Springer Boston, 2009, vol. 297/2009, ch. Flexible and Transparent User Authentication for Mobile Devices, pp. 1–12.
[8] M. Tanviruzzaman, S. I. Ahamed, C. S. Hasan, and C. O'Brien, "ePet: When Cellular Phone Learns to Recognize its Owner," in *Proceedings of ACM Workshop on Assurable & Usable Security Configuration (SafeConfig)*. Chicago, IL, USA: Collocated with the ACM Conference on Computer and Communications Security (CCS), November 2009, pp. 13–18.

[9] L. Araujo, L. S. Jr., M. Lizarraga, L. L. Ling, and J. B. Yabu-Uti, "User Authentication Through Typing Biometrics Features," *IEEE Transactions on Signal Processing*, vol. 53, Issue 2, Part 2, pp. 851–855, 2005.

[10] A. Buchoux and N. Clarke, "Deployment of Keystroke Analysis on a Smartphone," in *Proceedings of the 6th Australian Information Security Management Conference*. Perth, Western Australia: SECAU - Security Research Centre, 2008, pp. 40–47.

[11] R. Spillane, "Keyboard Apparatus for Personal Identification," IBM Technical Disclosure Bulletin, Tech. Rep. 17, 1975.

[12] M. Brown and S. Rogers, "Method and apparatus for verification of a computer user's identification, based on keystroke dynamics," U.S. Patent Number 5,557,686, September 17, 1996.

[13] J. Garcia, "Personal identification apparatus," U.S. Patent Number 4,621,334, November 4, 1986.

[14] J. Young and R. Hammond, "Method and apparatus for verifying an individual's identity," U.S. Patent Number 4,805,222, February 14, 1989.

[15] F. Bergadano, D. Gunetti, and C. Picardi, "User Authentication Through Keystroke Dynamics," *ACM Transactions on Information and System Security*, vol. 5, no. 4, pp. 367–397, November 2002.

[16] S. Bleha, C. Silvinsky, and B. Hussien, "Computer-Access Security Systems Using Keystroke Dynamics," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 12, pp. 1217–1222, December 1990.

[17] F. Monrose and A. Rubin, "Keystroke Dynamics as a Biometric for Authentication," *Future Generation Computer Systems*, vol. 16, pp. 351–359, 2000.

[18] S.Cho, C. Han, D. Han, and H. Kim, "Web based Keystroke Dynamics Identity Verification Using Neural Network," *Journal of Organizational Computing and Electronic Commerce*, vol. 10, no. 4, pp. 295–307, 2000.

[19] N. Clarke and S. Furnell, "Authenticating Mobile Phone Users Using Keystroke Analysis," *International Journal of Information Security*, vol. 6, no. 1, pp. 1–14, January 2007.

[20] D. Polemi, "Biometric Techniques: Review and Evaluation of Biometric Techniques for Identificaiton and Authentication, Including an Appraisal of the Areas Where they are Most Applicable"," Institute of Communication and Computer Systems, National Technical University of Athens, Tech. Rep., April 1997.

[21] N. Clarke, S. Furnell, B. Lines, and P. Reynolds, "Subscriber Authentication for Mobile Phones Using Keystroke Dynamics," in *Proceedings of the Third International Network Conference (INC 2002)*, Plymouth, UK, 2002, pp. 347–355.

[22] K. L. Adair, S. T. Parthasaradhi, and J. Kennedy, "Real World Evaluation: Avoiding Pitfalls of Fingerprint System Deployments," in *Proceedings of Biometrics India Expo 2008*, Pragati Maidan, New Delhi, India, 2008.

[23] N. Clarke, S. Furnell, and P. Reynolds, "Biometric Authentication for Mobile Devices," in *Proceedings of the 3rd Australian Information Warfare and Security Conference 2002*, 2002, pp. 61–69.

[24] D. Gunetti and C. Picardi, "Keystroke Analysis of Free Text," *ACM Transactions on Information and System Security*, vol. 8, no. 3, pp. 312–347, August 2005.

[25] S.Karatzouni and N. Clarke, *New Approaches for Security, Privacy and Trust in Complex Systems*. Springer Boston, 2007, vol. 232/2007, ch. Keystroke Analysis for Thumb-based Keyboards on Mobile Devices, pp. 253–263.

[26] D. Umphress and G. Williams, "Identity Verification Through Keyboard Characteristics," *International Journal of Man-Machine Studies*, vol. 23, no. 3, pp. 263–273, September 1985.

[27] R. Joyce and G. Gupta, "Identity Authentication Based on Keystroke Latencies," *Communications of the ACM*, vol. 33, no. 2, pp. 168–176, February 1990.

[28] J. Leggett, G. Williams, M. Usnick, and M. Longnecker, "Dynamic Identity Verification via Keystroke Characteristics," *International Journal of Man-Machine Studies*, vol. 35, no. 6, pp. 859–870, December 1991.

[29] F. Monrose and A. Rubin, "Authentication via Keystroke Dynamics," in *Proceedings of the 4th ACM Conference on Computer and Communications Security*, ACM. Zurich, Switzerland: ACM, 1997, pp. 48–56.

[30] M. Obaidat and B. Sadoun, "Verification of Computer Users Using Keystroke Dynamics," *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, vol. 27, no. 2, pp. 261–269, 1997.

[31] S. Haidar, A. Abbas, and A. Zaidi, "A Multi-Technique Approach for User Identification through Keystroke Dynamics," in *Proceedings of the 2000 IEEE International Conference on Systems, Man, and Cybernetics*, vol. 2. Nashville, Tennessee, USA: IEEE, 2000, pp. 1336–1341.

[32] A. A. E. Ahmed, I. Traore, and A. Almulhem, "Digital Fingerprinting Based on Keystroke Dynamics," in *Proceedings of the Second International Symposium on Human Aspects of Information Security & Assurance (HAISA 2008)*, Plymouth, UK, July 2008, pp. 94–104.

[33] H. Saevanee and P. Bhattarakosol, "Authenticating User Using Keystroke Dynamics and Finger Pressure," in *Proceedings of the 6th IEEE Consumer Communications and Networking Conference*. Las Vegas, NV, USA: IEEE, 2009, pp. 1–2.

[34] S. Zahid, M. Shahzad, S. Khayam, and M. Farooq, *Recent Advances in Intrusion Detection*, ser. Lecture Notes in Computer Science. Springer Berlin / Heidelberg, 2009, vol. 5758, ch. Keystroke-Based User Identification on Smart Phones, pp. 224–243.

[35] P. Campisi, E. Maiorana, M. L. Bosco, and A. Neri, "User Authentication Using Keystroke Dynamics for Cellular Phones," *IET Signal Processing - Special Issue on Biometric Recognition*, vol. 3, no. 4, pp. 333–341, 2009.