

The Plateau: Imitation Attack Resistance of Gait Biometrics

Bendik B. Mjaaland

Accenture Technology Consulting - Security, Norway
bendik.mjaaland@accenture.com

Abstract. Constituting a new branch within biometrics, gait biometrics needs to be extensively tested and analyzed to determine its level of fraud resistance. Previous results examining the attack resistance testing of gait authentication systems show that imitation, or mimicking of gait, is a venerable challenge.

This paper presents an experiment where participants are extensively trained to become skilled gait mimickers. Results show that our physiological characteristics tend to work against us when we try to change something as fundamental as the way we walk. Simple gait details can be adopted, but if the imitator changes several characteristics at once, the walk is likely to become uneven and mechanical. The participants showed few indications of learning, and the results of most attackers even worsened over time, showing that training did nothing to help them succeed.

With extensive training an impostor's performance can change, but this change seems to meet a natural boundary, a limit. This paper introduces the **plateau**, a physiologically predetermined limit to performance, forcing imitators back whenever they attempt to improve further. The location of this plateau determines the outcome of an attack; for success it has to lie below the acceptance threshold corresponding to the Equal Error Rate (EER).

1 Introduction

Biometric technology has applications ranging from accessing ones computer to obtaining visa for international travel. The deployment of large-scale biometric systems in both commercial (e.g. Disney World [7], airports [2]) and government (e.g. US-VISIT [15]) applications has served to increase the public awareness of this technology. This rapid growth in biometric system deployment has clearly highlighted the challenges associated in designing and integrating these systems.

Biometric features are divided into two categories: physiological and behavioral biometrics. Physiological biometrics are characteristics that you cannot alter easily, they are stable parts or properties of your body. Examples are fingerprints, DNA, iris and retina. Behavioral characteristics can be altered and learned, such as gait, signature and keystroke dynamics. Within computer science, biometrics is often considered within the field of pattern recognition, and creating automated system approaches to such has been challenging [9].

Recent advances in technology have shown us that we can automatically recognize individuals based on the way they walk. However, gait biometrics is a rather new area within biometrics, so the security of such technology needs to be challenged. Possibly the most intuitive threat towards gait biometrics is the mimicking attack, which is the topic of this paper. The reader should keep in mind that this is only one out of several attack points against biometric systems, a good reference here is [16], where Ratha et al. identified eight such attack points - imitation corresponding to number one in this research.

Contribution: Mjaaland et al. published research on gait mimicking in [11], and this paper extends this research. New insights to the imitator's challenges are presented, and conclusions are drawn on people's ability to adopt the gait characteristics of others. The term **plateau** is introduced, representing the natural boundary to an impostors performance in gait mimicking - as the title of this paper suggests.

This paper is organized as follows. Section 2 presents gait as a biometric feature and explains how it can be processed and analyzed. A description of Mjaaland's gait mimicking experiment is also included here. Section 3 introduces the plateau, and shows how to derive it using regression of impostor performance data. Section 4 presents the results, showing that single plateaus were discovered for almost all impostors. Finally, Section 5 and 6 presents conclusions and further work, respectively.

2 Gait Biometrics

2.1 Gait Analysis

The gait of a person is a periodic activity with each gait **cycle** covering two strides - the left foot forward and the right foot forward. It can also be split up into repetitive phases and tasks. Gait recognition has intrigued researchers for some time, already in the 1970's there were experiments done on the human perception of gait [9]. The first effort towards automated gait recognition (using machine vision) was probably done by Niyogi and Adelson [14] in the early 1990's. Several methods are known today, and we can categorize all known gait capturing methods into three categories [4]: machine vision based (MV), floor sensor based (FS) and wearable sensor based (WS).

Using wearable motion recording sensors to collect gait data is a rather newly explored field within gait biometrics. The earliest description of the idea known to the author is found in Morris' [12] PhD thesis from Harvard University in 2004. Since then, the academic community at Gjøvik University College (HiG) has devoted much effort researching gait biometrics. Gafurov's PhD work covers a broad part of WS-based gait recognition [4], and several students have written their master's thesis on the same topic [1,8,13,17].

WS-based gait collection has the advantage of being a rather unobtrusive way of collecting biometric data. It also has the immense advantage over MV of avoiding external noise factors such as camera placement and background or

lighting issues. Furthermore, MV and FS is an expensive solution in terms of camera and floor equipment, while the WS do not require any infrastructure in the surroundings, and it is mobile.

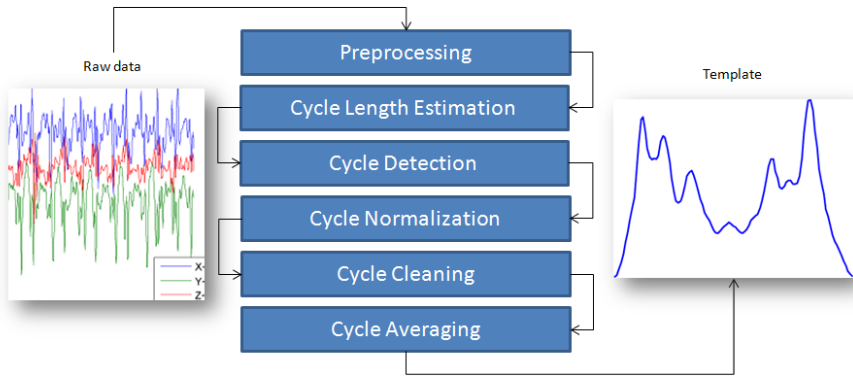


Fig. 1. Gait Analysis Overview

The best performance results have been achieved using WR technology, the best recorded to far is 2% [8], as opposed to 8% for MR and 65.8% for FS [5]. There are good reasons to choose WS, and this is the technology used in in the research presented in this paper. The analysis is based on the Cycle Length Method (CLM) developed by Gafurov [4] at Gjøvik University College. This method is essentially a framework on how to turn raw gait data into an averaged gait cycle of a fixed length. When the average cycle is found, this can be used as a user gait template.

The research presented in this paper is based on a self-developed software tool. The input to the tool is raw three-dimensional gait acceleration data, collected by wearable sensors attached to the hip. Data from the X, Y and Z direction constitute fragments of gait acceleration, which is combined into one resultant acceleration signal.

A high-level overview of the processing of raw gait data is shown in Figure 1. The first stage consists of preprocessing, where raw gait data is filtered and interpolated. The next three stages estimate the length of the cycles, detect their exact locations and normalizes them to the same length, respectively. The set of cycles is then cleaned (i.e. by removing outliers), and finally the average gait cycle is computed using simple averaging techniques.

The input to this algorithm typically consists of 10 – 20 seconds of normal walk, and the final output is an average gait cycle, consisting of two strides - one for each foot. This average can be compared to other averages by well-known statistical distance metrics. Holien studied distance metric performance for gait biometrics in [8], and the superior metric, Dynamic Time Warping (DTW), is also used in the research presented here. The DTW method disposes the

naturally occurring changes in walking speed and is able to compare signals of different lengths, and signals where the x-axis is shifted [8]. DTW can be used for various purposes, like for time-independent averaging, but in this research it was used simply as a distance metric. Algorithm details and MATLAB code used in this research can be found in [10].

2.2 The Mimicking Experiment

The research presented here is based on a mimicking experiment involving 50 participants, each enrolled with 10 templates, and 7 of these participated also as attackers.

The experiment is divided into three scenarios. The first is **the friendly scenario**, consisting of regular gait data from a test group of 50 participants. **The short-term and long-term hostile scenarios** are the two other scenarios, where a group of seven attackers attempted to imitate one specific victim. After each imitation attempt, the attacker was able to review his or her own gait on a video recording, and further improve the mimicking by comparing it to a video of the victim. A personal mimicking coach, or trainer, continuously assisted the attacker in the feedback process. Statistical results were also available. One session took about an hour, and five sessions were held for six "short-term" attackers, over a period of two weeks. The last attacker was the "long-term" attacker, having 20 sessions over six weeks.

The attackers and the victim in the two hostile scenarios were all selected from the friendly test group, and the attackers are referred to by numbering: A01, A03, A04, A18, A21, A38 and A41. The attackers were chosen such that their initial distances to the victim (i.e. the average distance between the victim's and the attacker's normal gait) consisted of both high and low values. Also, a reasonably stable gait was a requirement. In this regard the attackers represented "normal" people, with no specific advantages or disadvantages on average. Some anonymous non-sensitive details about the participants can be found in [10], but will not be included in this paper.

Two other mimicking experiments have been conducted, by Gafurov [6] and Stang [17]. However, these experiments are not extensive enough to provide valid indicators. Gafurov himself calls this part of his research a minimal-effort experiment, and in [10] Mjaaland presents a critique on both. In these experiments the participants do not conduct any structured training, and the search for *learning* is not properly conducted in either experiment.

The experiment described in this paper distinguishes itself from others by focusing on the attacker's mimicking skills, using sources of feedback like video capture and personal coaching to improve it. This investment of effort in impostor training has not been attempted in experiments before.

3 The Plateau: A Limit to Learning

As the experiment was conducted, attackers felt that despite that they sometimes improved performance, they found it very hard to improve beyond a certain

point. This limit, or plateau, is described in this section. The actual results are presented in Section 4.

3.1 Plateau Characteristics

A plateau can be defined as "a state or period of little or no change following a period of activity or progress" [3], so on a learning curve it would correspond to the curve flattening out. Hence, observations concentrate around some value on the Y-axis, illustrated in the left part of Figure 2.

If exactly one plateau exists for each individual, then the success of an attacker is **predetermined** - the plateau has to lie below or near the acceptance threshold for an impostor to ever be able to succeed. How near it has to be depends on the variance in the data, as deviations potentially can be successful attacks.

The data in this experiment is not sufficient to make final conclusions on how the participants would be affected by an even more extensive training program. The uncertainty of the future is one of the reasons why the name "plateau" was chosen. If a temporary plateau is reached, and the performance later increases due to an extended training period, the term still makes sense. In this case one can imagine several plateaus belonging to the same performance plot, as illustrated in the right part of Figure 2. This situation should generate interesting questions, such as how to break through the different plateaus.




Fig. 2. Plateaus, conceptual illustration. A single plateau to the left, two to the right.

3.2 Analysis

Intuitively plateaus can be identified by looking at points of resistance, average values and converging curves. Coefficients from fitted "trend lines" can also be put to use for this purpose. Still, the most scientific way to find the plateau would be to look for a mathematical *limit*. The limit of a function tells us exactly where it is heading when x goes to infinity.

What separates the plateau from a mathematical limit should be addressed. While the limit is a purely mathematical concept that may or may not properly illustrate learning as we know it, the plateau opens for more human-like function

behavior. The main difference is that if a function exhibits a limit, only one such limit can exist. If only one plateau exists for an attacker, then the plateau and limit are identical. However, in the above it was suggested that an observed set of data points could exhibit several plateaus, and it would be valuable to study how each one could be "broken" if this was the case.

If several plateaus exist for an attacker, and the attacker is improving over time, then the lowest one will equal the limit. It is important to note that when statistical distance metrics are used, lower score means less separability, or difference, and higher mimicking performance. Hence, if the attacker is worsening his performance, the highest plateau equals the limit. In general we can refer to this as the final plateau.

The main objective of this research is to identify learning - in other words a systematic change in the performance of the attackers over time. A regression analysis was conducted for the purpose of finding a learning curve.

The regression model chosen for the analysis is based on an exponential curve:

$$Y(X) = \beta_1 + \beta_2 e^{\frac{\beta_3}{X}}, \quad (1)$$

where β_1, β_2 and β_3 are constants, the regression parameters, and $Y(X)$ is the estimate of observation X . This model removes natural fluctuations in performance, and also converges. The latter property is desirable because it corresponds to the fact that a learning process normally has a limit to its effect. It will also simplify the process of identifying plateaus.

Final (and single) plateaus are identified by taking the mathematical limit of the learning curve:

$$\rho_{nm} = \lim_{x \rightarrow \infty} Y_{nm}(X), \quad (2)$$

where where ρ_{nm} is the plateau of participant nm (e.g. 01 for A01), and Y is the learning curve of that participant.

4 Results

4.1 Performance

The Decision Error Tradeoff (DET) curve shows system performance in terms of different thresholds, and the relationship between False Match Rate (FMR) and False Non-Match Rate (FNMR) in the system [8]. The curve is constructed by performing pairwise comparisons between all 500 templates enrolled in the system. The left part of Figure 3 shows the DET curve for the mimicking experiment. The Equal Error Rate (EER) is 6.2%, corresponding to an acceptance threshold of $T = 8.6449$. Since performance optimization was not the objective of this research, these results were indeed good enough to proceed. Also, by not lowering the EER further the attackers had an easier task, which is a scientific advantage in this case.

The right part of Figure 3 shows an example of mimicking results. As these results are measured using a statistical distance metric, downward sloping regression curves indicate improving results. For example, two identical gait samples




Fig. 3. Left: DET curve, friendly scenario, EER = 6.2%. Right: Long-term attacker A01 regression analysis.

yield a mimicking score of zero. The threshold of acceptance in this research is at $T = 8.6449$, so for the attacker to succeed he or she would have to exhibit a learning curve converging to a point below this value.

The example in Figure 3 is from the regression analysis of the long-term attacker A01. The dotted regression line is given by $Y_{01}(X) = 13.9962 + 0.2588e^{\frac{-19.8894}{X}}$, representing the learning curve. In this case the curve is rising, indicating worsening performance. Using confidence intervals for the parameters β_1 and β_2 , we can calculate a window of 95% certainty where the attacker is heading over time. The regression line of A01 converges to $\rho_{01} = 14.2550$, significantly higher than the acceptance threshold. The 95% confidence interval of this particular attacker is $[11.7599; 16.7501]$, so not even the most optimistic forecast yields a sufficiently low result for this attacker. More numeric results are found in [11].

Four out of seven attackers **worsened** their performance during training (A01, A04, A21 and A38), in other words - they were better off when they did not try at all. None of these were near the acceptance threshold, even according to the most optimistic forecasts. The analysis also shows that only two attackers (A03, A18) improved their performance, exhibiting downward sloping curves. The final attacker, A41, has an ill-defined regression curve with no obvious interpretation. However, all of these results are relevant and will be discussed in this paper.

4.2 Single Plateaus

As mentioned above two attackers improved during their training, A18 and A03, and their results are plotted in Figure 4 and 5, respectively. A03's performance was fluctuating around his plateau from the very beginning. Looking at his performance plot it can be seen that it is mostly his variance that changes. Very large fluctuations can be seen in the beginning, while at point A the results stabilize. A03 learned to focus on the particular characteristics of his gait that gave the best results, concentrating his result values around the plateau.

During the last 10 - 15 attempts, starting at point B in the figure, he decided to introduce more changes. This resulted in more fluctuations in both directions, but never produced stable improvement. A03 reported difficulty combining gait characteristics, and felt that he ended up walking in an unnatural way.

A18 will be used for numeric examples; his regression curve is defined as $Y_{18}(X) = 10.2662 - 2.0549e^{-\frac{22.3333}{X}}$. By also looking at the result plot, the reader can easily verify that this attacker is improving. His learning curve converges to $\rho_{18} = 8.2113$, which makes him the only successful impostor in the experiment. A03's curve converges to a limit above the acceptance threshold, $\rho_{03} = 10.5176$. For the complete regression analysis of A18, see Table 1.




Fig. 4. Attacker A03 learning curve superimposed on the data set. The straight line is the plateau, the attacker managed to stabilize his results here between point A and B.




Fig. 5. Attacker A18 learning curve superimposed on the data set. The straight line is the plateau, first reached at point D.

Although A18 did improve his performance, this improvement is mainly occurring between point C and D in Figure 5. After point D, the decreasing values seem to meet resistance, a performance boundary. The mathematical limit of the regression curve confirms diminishing improvement. A single plateau is found at the limit, 8.2113, with a 95% confidence interval of $[6.6288 < \rho_{18} < 9.7939]$.

When A18 realized that he had problems improving further, he made some radical changes to his gait. This made his walk mechanical and uneven, and eliminated some previously adopted gait characteristics. In the plot this can be seen, starting at point E with a significant increase, followed by high fluctuations and instability for the rest of the training. Such observations were made also for other attackers - significant changes of the gait seemed to neutralize previous improvement and acquired skill.

A18 was the success story, as he learned enough to break the threshold. However, this was not the case for the other participants. The four participants that worsened their performance, A01, A04, A21 and A38, had their plateaus identified at 14.2550, 14.8051, 13.1676 and 13.3554, respectively. These results are significantly higher than the acceptance threshold. Figure 6 provides an illustration for A01 and A21.

Table 1. Regression analysis for attacker A18. The regression curve converges to a plateau at 8.2113, and the curve is verified using the regression of the residuals and a hypothesis test [10].

Regression model	$Y(X) = \beta_1 + \beta_2 e^{\frac{\beta_3}{X}}$
Regression curve	$Y_{18}(X) = 10.2662 - 2.0549e^{-\frac{22.3333}{X}}$
Limit / Plateau	8.2113
β_1 95% confidence interval	$9.6371 < \beta_1 < 10.8953$
β_2 95% confidence interval	$-3.0083 < \beta_2 < -1.1015$
Plateau 95% confidence interval	$6.6288 < \rho_{18} < 9.7939$
MSE	0.6503
Residual regression model	$Y(X) = \lambda_1 + \lambda_2 X$
Residual regression curve	$Y_{18}(X) = -0.1179 + 0.0034X$
λ_2 confidence interval 95%	$-0.0064 < \lambda_2 < 0.0133$
Residual MSE	0.6358
H₀	$\lambda_2 = 0$
H₀ P-value	0.4901 (Failed to reject H ₀)




Fig. 6. The discovered plateaus of A01 (left) and A21 (right)

The effects of the plateaus were striking. Attacker A01 found an unnatural walk that sometimes gave better results compared to what his plateau suggested. It turned out, even if that particular walk was a better way of mimicking, he could not stabilize it. On the figure this can be seen starting at point A. The variance increased and the attacker lost control of the previously improved gait characteristics, and the plateau kept forcing A01’s performance back.

A21 was the attacker with the least control of his gait. Most of the time he experienced problems when trying to change characteristics of his walking, and the reader may verify this looking at the high variation in his results. His worsening performance did sometimes stabilize around the plateau, around point B and C.

There were significant differences between the participants in terms of *how* they reached their plateau. The three attackers A03, A18 and A38 changed the most during the training, producing steeply sloped learning curves. Obviously this does not necessarily mean they increased their skills, just that their original score was far away from the plateau. Other attackers, like the long-term attacker A01, initially got results very close to their plateaus, and thus found it a lot harder to improve beyond that point.

Furthermore, one can ask whether or not a higher plateau, where a performance *decrease* hits a boundary, is less interesting than a plateau reached after an *increase* of performance. No matter what direction the learning curve takes, it reflects the ability of the attacker to change his walk into something different, and keep it stable. When it comes to learning we are interested mainly in how far away the plateau is from the initial few mimicking attempts.

The results show that it is possible to slightly improve the performance of gait mimicking using training and feedback. By experimenting with small and large changes of gait characteristics, two attackers did move somewhat closer to the victim. However, the performance increase shown is very limited. It was clear that the attackers met their natural boundaries and had huge problems moving on from that point. This indicates that even if you can train to adopt certain characteristics of your victim, the outcome of your attempt is predetermined by your plateau. If it lies too high, you will never be able to mimic that person. And most attackers actually experienced worsening performance, converging towards a plateau far above the acceptance threshold.

4.3 Breaking Through

It is useful to look at the possible existence of several plateaus simultaneously, challenging the assumption of a predetermined attack outcome. Looking at the plot for A41 in Figure 7, two plateaus are superimposed, illustrating a plateau break in point A. Notice the sharp decline of the regression curve after the break; this curve is converging to a negative value. We cannot calculate the second plateau directly due to this ill-defined regression result, hence the lines are not mathematically derived and are provided for illustration only.

We may also want to re-interpret A18's performance plot in Figure 5, and add a second plateau as shown in Figure 8. This attacker was initially producing very static results, indicating a plateau, but at point C the results began to improve. This improvement reached more resistance around point D. Reaching the new plateau represented a significant improvement, although A18's performance also became unstable. In fact it may seem like A18 returned to his original plateau in point E.




Fig. 7. A41's results, showing a possible plateau break at point A, and a new plateau reached at point B.

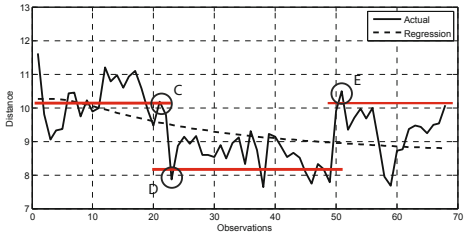


Fig. 8. A re-interpretation of A18's results, with a possible plateau break at point C, a new plateau at point D, and a return to the original plateau in point E.

The reader should note that the multi-plateau observations here are uncertain. Not enough data is available to support any claims on several plateaus, and the results have very high fluctuations in the areas of possible plateau breaks. A more plausible explanation would be that the "real" plateau is the second, or final plateau, and that the first is caused by noise from poor mimicking. Furthermore, the research also indicated that the assumed new plateau was very hard to reach, and that the degree of learning was inadequate to pose a real threat. Even with many plateaus, reaching the second requires extreme effort, and any proceeding plateaus are likely to be harder to reach than the one before.

5 Conclusions

This paper looked at fraud resistance of gait biometrics with respect to imitation, or mimicking attacks. The research showed that learning to walk like someone else is very difficult, and that there is a limit to the attacker's performance that causes the attack results to be predetermined. Either the attacker's normal walk, with a slight training adjustment, is close enough to that of the victim, or it is not. Extensive training, even with coaching and continuous feedback, does not seem to change this outcome.

The attackers hit a natural boundary that prevented them from improving their performance beyond a certain point. The effect of this phenomenon was striking, and given a name: a plateau.

We presented research on gait mimicking in [11], where an experiment was conducted with a specific objective in mind - increasing attackers imitation skills. This was attempted using various forms of training and feedback, tailored to see if the ability to mimic gait can change over time.

The main part of the experiment was "hostile", with attackers training to imitate the gait of the same victim. The participants did not show significant improvement or learning overall. A regression analysis was conducted in order to establish this fact - most learning curves were sloping upwards, indicating worsening performance.

With the research of this paper it was found that training has little or no effect on the plateau, it seems to be a physiologically predetermined boundary to every individual's mimicking performance. If only one such plateau exists, then it is mathematically the same as a limit, in essence - the value to which the learning curve converges. Natural fluctuations will be present, but the average results will approach the plateau over time.

A single plateau was identified for all attackers with an interpretable regression curve. However, it was speculated on whether or not individuals can exhibit more than one plateau. This is partly why the term plateau was chosen in the first place - it is not as strict as the mathematical definition of a limit.

To succeed in a mimicking attack with a single plateau, the plateau itself must lie below or close to the threshold of acceptance. That is, natural fluctuations may cause some successful attempts even if the plateau is above the threshold. This depends on the variance of the gait. Only one attacker (A18) posed a direct threat to the victim.

In order to succeed in a mimicking attack with several plateaus, the attacker has to reach a plateau that has the same properties as the one described above - close to or below the threshold. Little data suggested that multiple plateaus can exist simultaneously, but even if it could it is also likely that plateaus get harder to break the closer to the threshold they are.

It should be noted that the findings of this paper cannot necessarily be generalized to apply to other analysis methods. The results here applies to the combination of methods and configurations presented. However, a lot of the difficulties in gait mimicking are likely to be physiologically rooted, and thus it is reasonable to assume that the indicators are relevant in other contexts as well.

Summarized, the attackers had varying skills and results, and only one of them got to a plateau below the threshold. If we consider a future improvement in gait biometrics, that one attacker is a negligible threat. What is much more important - the attackers hardly learned at all. Their improvements were mostly insignificant, and they struggled hard to be free of their plateaus.

6 Further Work

There are various possibilities for future work on gait biometrics. For gait mimicking in particular - this research cannot provide final answers because the data set is not large enough, and not all the results can be generalized. More experiments are needed, with different methods, more test subjects and more extensive training. In particular, extended time frames and more training would be the key to get a definite conclusion on the single/multiple plateau issue.

Different results might be achieved if the gait analysis scheme is significantly changed. For instance, if Fast Fourier Transformation (FFT) was used in order to look on gait in the frequency domain, we might see other trends and characteristics in the attackers performance. A large part of the observations are generalizable - for instance that physiological boundaries makes gait mimicking difficult, but we cannot claim that our observations hold under all circumstances. A future task could be to perform the same experiment with different analysis tools.

It would be interesting to get precise physiological explanations on *why* it is so hard to imitate other peoples walks. For this, anatomical and medical studies would have to come into play, and a successful report on the topic would definitely strengthen the documented security of gait biometrics.

Other aspects of mimicking could also be analyzed - like threats through cooperation. Cooperation in gait mimicking essentially means that two people try to walk like each other, and then maybe "meet in the middle". Hence, one person could enroll walking somewhat like a different person and, if successful, they could both authenticate with the same template.

It would be beneficial to try to identify so-called sheep and wolf characteristics within gait biometrics. Some people may be easier targets for imitation, "sheep", and some people may be better at impersonating than others, "wolves". Further, such (dis)advantages could be genetically determined, and these issues together can form entire new lines of research within gait mimicking.

On the field of gait biometrics in general there is a lot of work to do. The performance of gait recognition systems are not generally competitive to other biometrics at the time of writing, so the invention of new methods, and further development of the existing methods is necessary.

References

1. Buvarp, T.E.: Hip movement based authentication - how will imitation affect the results? Master's thesis, Gjøvik University College - Department of Computer Science and Media Technology (2006)
2. Clarke, R.: Biometrics in airports how to, and how not to, stop mahammed atta and friends (2003), <http://www.anu.edu.au/people/Roger.Clarke/DV/BioAirports.html>
3. Oxford Dictionaries. Compact Oxford English Dictionary of Current English, 3rd edn. (2005)
4. Gafurov, D.: Performance and Security Analysis of Gait-based User Authentication. PhD thesis, University of Oslo (2008)
5. Gafurov, D., Snekkenes, E., Bours, P.: Gait authentication and identification using wearable accelerometer sensor. In: Proceedings of the IEEE Workshop on Automatic Identification Advanced Technologies, AutoID (2007)
6. Gafurov, D., Snekkenes, E., Bours, P.: Spoof attacks on gait authentication system. Special Issue on Human Detection and Recognition (2007)
7. Harmel, K., Spadanuta, L.: Disney world scans fingerprint details of park visitors. The Boston Globe (September 3, 2006)
8. Holien, K.: Gait recognition under non-standard circumstances. Master's thesis, Gjøvik University College - Department of Computer Science and Media Technology (2008)
9. Jain, A.K., Flynn, P., Ross, A.A.: Handbook of Biometrics, vol. 556. Springer, US (2008)
10. Mjaaland, B.B.: Gait mimicking - attack resistance testing of gait authentication systems. Master's thesis, Norwegian University of Science and Technology, NTNU (2009)
11. Mjaaland, B.B., Bours, P., Gligoroski, D.: Gait mimicking - attack resistance testing of gait authentication systems. In: Proceedings of the 2nd Norwegian Information Security Conference, NISK 2009, NISNet, Tapir Akademiske Forlag, Trondheim, Norway (2009)
12. Morris, S.J.: A shoe-integrated sensor system for wireless gait analysis and real-time therapeutic feedback. PhD Thesis, Harvard University - MIT Division of Health Sciences and Technology (2004)
13. SØndrol, T.: Using the human gait for authentication. Master's thesis, Gjøvik University College - Department of Computer Science and Media Technology (2005)
14. Nixon, S.A., Adelson, E.H.: Analyzing gait with spatiotemporal surfaces. In: Proceedings of IEEE Workshop on Non-Rigid Motion (1994)
15. U.S. Department of State. Safety and security of u.s. borders/biometrics. State official online information (2008)
16. Ratha, N.K., Connell, J.H., Bolle, R.M.: An analysis of minutiae matching strength. IBM Thomas J. Watson Research Center (2001)
17. Stang, Ø.: Gait analysis: Is it easy to learn to walk like someone else? Master's thesis, Gjøvik University College - Department of Computer Science and Media Technology (2007)