

Predicting Graphical Passwords

Matthieu Devlin¹, Jason R.C. Nurse^{1†}, Duncan Hodges²,
Michael Goldsmith¹, and Sadie Creese¹

¹ Cyber Security Centre, Department of Computer Science,
University of Oxford, Oxford, UK

² Centre for Cyber Security and Information Systems, Cranfield University, UK
[†]jason.nurse@cs.ox.ac.uk

Abstract. Over the last decade, the popularity of graphical passwords has increased tremendously. They can now be found on various devices and systems, including platforms such as the Windows 8 and Android operating systems. In this paper, we focus on the PassPoints graphical-password scheme and investigate the extent to which these passwords might be predicted based on knowledge of the individual (e.g., their age, gender, education, learning style). We are particularly interested in understanding whether graphical passwords may suffer the same weaknesses as textual passwords, which are often strongly correlated with an individual using memorable information (such as the individuals spouses, pets, preferred sports teams, children, and so on). This paper also introduces a novel metric for graphical-password strength to provide feedback to an individual without the requirement of knowing the image or having password statistics a priori.

Key words: Graphical passwords; PassPoints scheme; user characteristics; usable security; password-strength metric

1 Introduction

As cyberspace becomes more pervasive throughout society, being used by everything from commerce and banking to how we interact with others, the ability to login securely to a computer or service has become critical to our ability to reliably use and secure these services. This is particularly true as we increasingly ask those who are not digital natives to engage with government and local services through the Internet. The most common form of authentication mechanism is passwords, often called the ‘keys’ to our digital life. These are an important technique for protecting both enterprises and individuals alike.

In theory, the textual passwords that are widely used today can provide a good degree of security, but in practice they are often insufficient due to human factors such as our inability to memorise a number of ‘strong’ passwords, typically deemed as a mixture of upper-case and lower-case letters, numbers and punctuation [1]. Various studies have highlighted that users often use the same (or very similar), usually ‘weak’, passwords across multiple accounts [2]. This

clearly does not permit the full security potential of textual passwords as an authentication mechanism and potentially poses a risk to the security of the system and the data it contains. Graphical passwords aim to exploit human cognitive processes in order to, arguably, provide a stronger authentication scheme which users find more memorable.

In this paper, we focus on one graphical password scheme, PassPoints, with the aim of investigating the extent to which these graphical passwords might be predictable. Rather than just focus on the choice of click location and image selection as in previous studies, our study considers the relationship between participants' PassPoints passwords and their individual characteristics, such as age, gender, ethnicity, education, risk appetite and learning style. Our work goes beyond current research articles which focus on heat-maps based on image clicks, to consider individuals' characteristics and the possibility of prediction in the same way that users tend to choose textual passwords based upon their spouses, pets, sports teams, children, and so on [3]. Any such relationship provides a good avenue of exploitation for an attacker.

This paper is structured as follows: in Section 2 we briefly reflect on graphical passwords, the PassPoints scheme, and related security concerns. Section 3 introduces our study and the user experiment conducted to investigate whether PassPoints passwords might be inferred based on a user's characteristics. Section 4 presents and discusses the key results of the study, before we conclude the paper in Section 5.

2 Background

Many graphical-password schemes have been proposed such as PassPoints, PassFaces and PassDoodle which claim to provide enhanced security, memorability and usability. The scheme examined in this paper is a cued-recall scheme called PassPoints which was developed by Wiedenbeck et al. [4]. A user creates a password by clicking (or touching on a touchscreen) an on-screen canvas a prescribed number of times (usually 5). The order and position of these clicks are recorded and establish the user's authentication token. To improve memorability, an image is used as the canvas. To authenticate, the user clicks the same points (within some tolerance) in the same order.

Much like textual passwords, the security of PassPoints passwords can be severely weakened by the predictable behaviour of users. For instance, a common issue with a PassPoints password is that users predominately choose the salient areas of an image as the password points. Golofit [5] showed that over 50% of user clicks in their study were in areas encompassing only 3% of the total image area. This demonstrates a huge reduction in the password size space and decreases the password's resilience to 'brute-force' attacks.

Dictionary attacks, where pre-computed lists of likely passwords are used to guess the password in question, have also been proven dangerously effective against PassPoints-style graphical passwords [6, 7]. Van Oorschott and

Thorpe [6] presented two dictionary-based attacks. The first was a ‘human-seeded’ attack, which uses experimental data to predict hot-spots of the images (areas with a large number of clicks). The second attack combined dictionaries with click-order patterns to create a first-order Markov model-based dictionary. This dictionary found 7–10% of passwords within 3 guesses.

Dirik et al. [7] highlighted the importance of image choice on the vulnerability of a PassPoints password to a dictionary attacks. Using image processing techniques, they created a dictionary of the most likely click locations and showed an attack on two images. Their analysis of the two images demonstrated that images with a higher click-point entropy (i.e., an estimate of the number of different likely locations for each click point in the password) are much less vulnerable to dictionary attacks.

While there has been much more research in the area of PassPoints and hotspots generally, to our knowledge, there has been little emphasis on investigating the extent to which PassPoints might be predicted based on a user’s characteristics. Considering the possible impact of such an ability, this is the focus of this paper’s study.

3 Study and approach

To examine the predictability of PassPoints, we designed an online user study. A website was developed and deployed that would allow participants to (1) register for the study, (2) to create a PassPoints password based one of three pre-defined images presented below, and (3) to complete three short surveys, one on their demographics, one on learning styles (motivated by work in [8]) and the last one on their risk appetite (adapted from research by Weber et al. [9]). After the creation of their PassPoints password, each participant was asked to recall the PassPoints on the site on three separate occasions: once immediately after they completed the surveys, the next time three days after the password was created, and the final time seven days after initial completion.

As mentioned in Section 2, choice of images for a PassPoints password is extremely important; in the real-world, individuals can actually select their preferred images. To control our experiment however, we decided to present individuals with three images and allow them to select the one to use for their password in the study. The images that were used are shown in Figure 1, and were chosen because they have different themes (e.g., People¹, Landscape², and Animal³) that might appeal to different participants. Also, these images have multiple salient points. This was important to the study since images with few salient points are likely to result in similar passwords, making the images the limiting factor in creating similarities between passwords as opposed to the similar characteristics of the participants.

¹ ‘Marton Mere Swimming Pool’ by havenholidays (<https://flic.kr/p/4ycWeu>)

² ‘One of the Glens, Scotland’ by Chris Ford (<https://flic.kr/p/8BumLU>)

³ ‘Untitled’ by PollyDot (<http://pixabay.com/en/chameleon-lizard-multi-coloured-318649/>)



Fig. 1. Images which participants used to create their PassPoints passwords. All images were available under the Creative Commons Licence, and sized 620px by 413px.

Recruitment for the survey was carried out using social-media networks and posters around the university campus. For their participation, individuals were entered into a prize-draw with a chance to win a voucher. In total, 236 individuals registered for the study but only 150 individuals completed it; completion was based on whether the participant returned to complete the next phase as necessary, not whether they were (un)able to recall their password.

4 Results and discussion

Our presentation and discussion of results is structured into five main sections. Each section considers a separate part of the general question of whether passwords can be predicted.

4.1 Which characteristics do people who chose the same image have in common?

This section questioned whether there were any characteristics of participants that might be used to predict which of the three images they chose to create their password. To answer this question, we used a Fisher's exact test [10] on each attribute to test whether there were any attributes that led participants to choose a certain image. Fisher's exact test is a common statistical method used to determine whether there are associations between two categorical variables.

For the test, the participant attributes (e.g., **gender** or **education**), were gathered from the questionnaire data that define a participant's overall characteristics. This has been done since a single characteristic can encompass multiple attributes: for example there are 10 different attributes for the learning styles characteristic. Moreover, all continuous variables such as age or mean of responses to the risk-appetite questionnaire had to be discretised into categories, since the test uses nominal data.

From this analysis, we found that no significant results were obtained when applying the Fisher's exact test to each of the attributes on the three images and therefore, conclude that the choice of image by the user is not dependent of any single attribute.

Next, a multiclass classifier was fit to the attributes associated with the participants. Since the input was a vector of attributes and the output was one of 3 categories, corresponding to the 3 images, a multiclass classifier was an

appropriate technique. Backward stepwise regression was used to automatically choose which variables were most relevant to predicting a participant's image choice. At each 'step', the independent variable which has the least impact on how the model fits the data is removed. The remaining independent variables are those that have the most impact on fitting the model to the data.

To avoid overfitting in the model, the data was randomly shuffled and the first 80% of participants in the shuffled set were used to select and train the model. This process was repeated 50 times. If there existed an accurate model to predict the image that the participant would choose, then the same set of attributes (or a very similar set) would be selected in each model.

The following attributes were in all 50 models: `auditoryNorm`; `tactileNorm`; `visualNorm`; `ethical`; `gambling`; `health`; `investment`; `recreational`; `social`. The `auditoryNorm`, `tactileNorm` and `visualNorm` attributes correspond to normalised scores for the learning styles questionnaire (i.e., auditory learners, tactical learners or visual learners respectively [8]) and `ethical`, `gambling`, `health`, `investment`, `recreational` and `social` correspond to the mean response to sets of questions on those topics in the risk-appetite questionnaire. Therefore a participant's responses to the `ethical` set of questions (for example) is possibly indicative of their choice of image in some way.

Although there were multiple attributes appearing in the model for each iteration, we found that the models always performed poorly when they were tested on the remaining 20% of the data in terms of accuracy, precision and recall. The repeated presence of a number of attributes in the model for each iteration provides some suggestion that a multiclass classifier could potentially be used to predict which image the participant chooses. However, the small sample size does not allow for an accurate model to be produced. A larger sample would be needed to explore this claim.

4.2 What characteristics do people who can recall their password have in common?

Another interesting question is whether there are any characteristics that act as an indicator for password recall, that is, those participants that are able to reliably recall the PassPoints and the order in which they are required. By fitting a logistic regression model to the data, a set of attributes that can be used to predict whether a participant will pass or fail can be extracted. As in the previous section, the desired model was initially unknown and therefore backwards stepwise logistic regression was used. The `ethical`, `investment` and `riskMeans` attributes appeared in the model on more than 10 of the 50 iterations. `riskMeans` refers to the mean of the user's responses to all risk appetite questions on the Likert (5-level) scale, and so a higher `riskMeans` value corresponds to a risk-taker and a lower value to a risk-avoider. Although this is far from unanimous, it may suggest that these attributes do have an influence on a participants' ability to recall their password after 3 days.

In terms of the accuracy, precision and recall performance of the model, these seemed promising (0.728, 0.753 and 0.957 respectively) but upon closer

inspection we found that there was clearly a fault with the classifiers; they are very good at predicting if a participant will recall but very poor at predicting if they will not. A classifier can achieve an artificially high level of accuracy if it predicts the most common output for all inputs and this is apparent here.

To determine whether there was a difference in where people who failed to recall their password after 3 days clicked and those that could recall clicked, a Fisher’s exact test was used. This involved first discretising the images into 20 px by 18 px boxes and counting the number of clicks in each box separately for those who passed and those who failed.

From our analysis of the images, only with the Animal image was a significant result ($p < 0.05$) found; i.e., the participant’s ability to recall their password is dependent on where they clicked. Using Fisher’s method to combine the p -values on each of the images yields $p = 0.008$ which is also significant. Therefore the null hypothesis that the participant’s ability to recall their password is independent of where they clicked is rejected. This may suggest that the users that fail to recall their PassPoints do so partly because of where they choose to click. There are many potential reasons for this, perhaps they tend or decide to choose points that are harder to remember, or fail to pick out locations that are easy to remember.

To visually compare the difference in where participants who passed and failed clicked, scatter plots were created with the points of participants who passed in green and those that failed in red. These are shown in Figure 2.




Fig. 2. Pass/Fail scatter plot for the People, Landscape and Animal images respectively; passes are shown in green and fails in red.

It is clear, particularly in the Landscape image in the middle of Figure 2, that the participants who failed to recall their password after 3 days chose different points to those who can recall. Often the differences in the positions of clicks are small but this may have a large impact on whether the participant can recall the password later on.

A Fisher’s exact test was used to determine whether there was a difference in where people who could recall after 3 days but not after 7, and people who could recall after 3 and 7 days clicked. When combined using Fisher’s method, $p > 0.05$ suggesting there is not sufficient evidence to reject the null hypothesis that they are independent, suggesting that there is no difference between where participants who could recall after 3 and 7 days clicked and participants who could recall after 3 days but not after 7.

4.3 Where do people click and what characteristics do people who click in the same place share?

To analyse where people clicked, saliency and heat maps were created for each image. Saliency maps provide a visual representation of how much each point of an input image stands out with respect to its neighbouring points. Heat maps show every participant's click-points as well as areas of the images that were popular amongst participants. Previous research [5, 6] has shown that users tend to use salient points as part of their password; this was tested as well.




Fig. 3. Saliency map (left) and heat maps (middle and right) for the People image



Fig. 4. Saliency map (left) and heat maps (middle and right) for the Landscape image




Fig. 5. Saliency map (left) and heat maps (middle and right) for the Animal image

Figures 3, 4 and 5 show the saliency maps and heat maps of the People, Landscape and Animal images respectively. There are some similarities between the maps particularly for the People (Figure 3) and Animal (Figure 5) images, agreeing with the previous research that some users pick salient points as part of their passwords. On the People image, it can be seen that the children's faces in the top left corner are both salient and a popular region for participants to click, as well as the chameleon's eye, nose and feet in the Animal picture.

The saliency map and heat map for the Landscape image in Figure 4 are noticeably different. The saliency map focuses on the rocks and stream in the

bottom right quadrant of the image, whereas the heat map shows the most popular region as the central mountain. It is not surprising that the mountains were used by many participants as part of their password as they are memorable points; it is perhaps more surprising that the mountain peaks were not picked up as salient points in the saliency maps. This may be due to the Itti-Koch-Neibur algorithm [11] that was used to implement the saliency map. An alternative algorithm might have produced a slightly different saliency map that was more consistent with naively expected salient regions.

Fisher’s exact tests were performed on 43 categorical attributes of the participants (obtained from the survey data) against the ‘boxes’ in which they click on the images (as defined using the same discretisation technique as in the previous section). Of the 43 tests, the only significant value at the $\alpha = 0.05$ level was for the `GamblingCats2` attribute when combining the p -values for the three images using Fisher’s method, $p = 0.003$. `GamblingCats2` corresponds to a participant’s responses to questions in the gambling section of the risk appetite survey and thus, their attitude towards gambling may in some way affect where they click.

Fisher exact tests were also done on combinations of attributes, such as being married and having children, and position of clicks; but all results were non-significant. In summary, many participants choose similar points, that are often salient, but the attributes that they possess, at least the ones in this study, seem to have little effect on the locations where they will click.

4.4 Do people share the same pattern of clicks (i.e., the order) and if so which characteristics do they share?

Each participant’s click pattern was analysed and classified (similar to Ref [6]) with patterns including: Left-to-right (LR), Right-to-left (RL), Top-to-Bottom (TB), Bottom-to-top (BT), Bottom-left-to-top-right (LR_BT), Clockwise (CW), Anticlockwise (ACW) and None (NONE). From our analysis, we found that in terms of frequency, None (i.e., no pattern) was the most common, followed by LR and TB, occurring 71, 37 and 23 times respectively.

A Fisher’s exact test was used to test whether there were statistical differences in the proportion of two patterns; LR versus RL, TB versus BT, CW versus ACW and any pattern versus no pattern (NONE). The findings highlighted significant result for the LR versus RL patterns. This is not surprising, as the vast majority of participants were from countries where reading and writing is performed from left-to-right, so it is to be expected that this pattern would be transferred into other activities such as creating a graphical password.

To examine whether there is a difference in the pattern a participant adopts based on their attributes, a Fisher’s test was performed on each attribute and pattern. Three combinations were significant at the $\alpha = 0.05$ level without correction for multiple comparisons: `gender` with BT pattern ($p = 0.016$), `children` with TB pattern ($p = 0.014$) and `riskCats3` with BT pattern ($p = 0.021$).

The first of these significant results shows that men are more likely than women to use a BT pattern (proportions: $men = 0.16$, $women = 0.03$). The second shows that participants who do not have children are more likely to

use a TB pattern than those with children (proportions: *without children* = 0.19, *with children* = 0.00). The final result shows that participants with a medium risk average are less likely to use a BT pattern than those with a high or low average and that those with a high average are most likely to use the BT pattern (proportions: *low* = 0.20, *medium* = 0.08, *high* = 0.36). It is difficult to qualify these results; it is conceivable that there would be a difference between genders but the other two results are harder to explain.

The final analysis looks at whether the image can influence the participant to create a password with a certain pattern. A Fisher’s exact test on counts of the number of participants that used the pattern in question and those who did not against the 3 images, was used.

Pattern	p -value	Proportion of Pattern on Image		
		People	Landscape	Animal
LR	0.911	0.222	0.243	0.273
RL	0.229	0.056	0.143	0.000
BT	0.790	0.138	0.100	0.091
TB	0.028	0.250	0.171	0.045
LR_BT	0.028	0.083	0.000	0.068
RL_BT	0.724	0.028	0.014	0.000
LR_TB	0.545	0.111	0.086	0.045
RL_TB	1.000	0.000	0.000	0.000
CW	0.020	0.000	0.057	0.159
ACW	0.896	0.083	0.129	0.114
NONE	0.587	0.528	0.429	0.500

Table 1. Difference in click-order patterns on the three images

There were significant results for the TB, LR_BT and CW patterns. Table 1 suggests that the image does have an influence on the patterns in a participant’s password. The TB pattern is an example of this with 25% of participants of the People image having the pattern whereas less than 5% of participants use it on the Animal image. Grouping the CW and ACW patterns, over 27% of participants with the Animal image used a ‘rotational’ pattern whereas only 8% and 19% used them on the People and Landscape images respectively. The LR pattern is prevalent in each of the images in similar proportions; 22%, 24% and 27% for the People, Landscape and Animal images respectively. This may be due to the majority of participants being nationals of countries with a left-to-right written language and this pattern creeps into other ‘observational’ activities.

4.5 Evaluating PassPoints password strength

Often when a user is creating an textual password on a website, they are provided with feedback on the strength of their password, usually on a scale from ‘weak’ to ‘strong’. It might be useful if the same feedback were given for graphical passwords, especially as most users are unfamiliar with them and may not understand what is a ‘strong’ password. To do this, a set of rules were created for scoring passwords based on information taken from saliency maps and click-order patterns. This allows the method to be used on any image and does not require any click-point data for that particular image. These rules were developed after the survey was completed and used the data collected from it, hence participants were not shown the strength of their password during the study.

Using the saliency maps generated for the images, the n^{th} percentile of saliency values for each image was calculated ($n = 90, 95, 97$). These values for n were chosen as Golofit [5] found that 50% of clicks occurred in only 3% of an image; the other two levels were included to increase the size of the salient region used, without making it too large. The regions of the saliency maps for each image were then filtered such that only the highest $(100-n)\%$ of saliency values remained; the set of these regions is denoted $\phi_{n,i}$ where $n = 90, 95, 97$ and $i = 0, 1, 2$; the image used.

The password weakness score was calculated as follows: (1) $\delta_{sal;n,i}$ = the number of the participant’s click-points in the filtered saliency map $\phi_{n,i}$; (2) $\delta_{pat;i}$ = 1 if the participant’s password exhibits a pattern described in the previous section, else 0; $\delta_{LR;i}$ = 1 if the participant’s password exhibits a left-to-right pattern as described previously, else 0. The password weakness score, $\Delta_{n,i}$, is then $\Delta_{n,i} = \delta_{sal;n,i} + \delta_{pat;i} + \delta_{LR;i}$, ($0 \leq \Delta_{n,i} \leq 7$) and therefore the lower the value of $\Delta_{n,i}$, the stronger the password. The left-to-right pattern was chosen as an indicator of password strength as over 20% of passwords had this pattern for each of the three images hence if an attacker prioritises for this common pattern, the password could possibly be obtained faster.

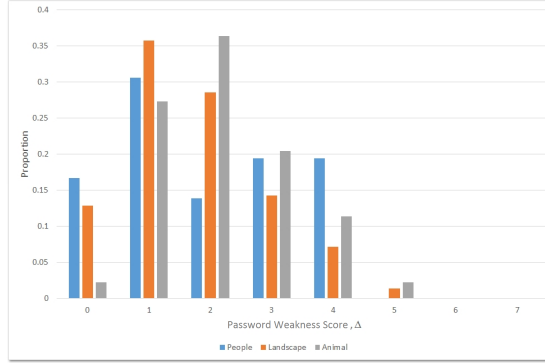


Fig. 6. Bar chart of proportions of password weakness score for each image

The password weakness score, $\Delta_{n,i}$, was calculated for each participant using $n = 90, 95, 97$ however $n = 90$ was chosen as the final percentile value for the saliency maps as it produced a larger distribution of password scores. Figure 6 shows the proportion of passwords with each password weakness score for each image. The Landscape and Animal images have approximately normal distributions of the proportions and the Landscape image has a positive skew. Therefore there are more ‘stronger’ passwords on the Landscape image than the Animal image. This may be because the Landscape image facilitates ‘stronger’ passwords than the Animal image or, by chance, the participants that chose the Landscape image pick ‘stronger’ passwords.

The People image has an approximately uniform distribution but with a larger value for $\Delta_{90,0} = 1$. Therefore, there are approximately equal proportions

of each password weakness score for $0 \leq \Delta_{90,i} \leq 4$. This may suggest that the People image promotes more ‘weaker’ (although also more ‘stronger’) passwords than the other two images. This may present a weakness as a user that creates a ‘weak’ password on the People image, may have chosen a ‘stronger’ password on the Landscape or Animal image because they discourage ‘weaker’ passwords.




Fig. 7. Weaker passwords: People image password (left) with $\Delta_{90,0} = 4$; Landscape image password (right) with $\Delta_{90,1} = 5$



Fig. 8. Strong passwords: Landscape image password (left) with $\Delta_{90,1} = 0$; Animal image password (right) with $\Delta_{90,2} = 0$

Figure 7 shows passwords that got high ($\Delta_{90,i} \geq 4$) password weakness scores and Figure 8 shows passwords that got low ($\Delta_{90,i} = 0$) scores. Both passwords in Figure 7 however, could be considered ‘weak’ by an observer suggesting that the rules defined can identify a ‘weak’ password. However, the Animal image-based password in Figure 8 is given a weakness score $\Delta_{90,2} = 0$ but does not appear to be as ‘strong’ as the score suggests, especially after comparison with the Landscape image to its left, which also has a score $\Delta_{90,1} = 0$.

It would seem that the only non-obvious point-choice is on the chameleon’s body, between its legs. The reason for this anomaly is that calculating the weakness score relies heavily on the saliency maps (5 of the possible 7 points are awarded from the position of clicks), and the saliency map for the Animal image does not recognise the chameleon’s nose as one of the most salient points, nor the front knee joint or either foot. However, each of these locations may be considered salient by a human observer. To enhance the accuracy of the password

strength metric, the algorithm implementing the saliency maps for the images would have to be improved so that the most ‘stand-out’ features are identified.

5 Conclusions

In this paper, we have investigated the extent to which the passwords under the PassPoints graphical-password scheme could be predicted based on knowledge of the password setter. In general, the findings have provided little statistically significant evidence that it is consistently possible to do so. While this is encouraging, a larger sample size is needed to confirm these results. This study has confirmed that participants tend to choose similar locations for their password points, thereby creating hotspots, especially around salient points in an image.

There is also some evidence that the click-point pattern that a user’s password exhibits is dependent on the image that they choose as their background, exposing a potential weakness in the scheme. However, this evidence was only available after processing the password data on the image, therefore an attacker would need prior knowledge of the image and have data about its usage in order to exploit this information. Finally, this paper has introduced a password strength metric that can provide feedback on a PassPoints password without the requirement of knowing the image or having password data for that image.

References

1. Nurse, J.R.C., Creese, S., Goldsmith, M., Lamberts, K.: Guidelines for usable cybersecurity: Past and present. In: Proceedings of the 3rd Cyberspace Safety and Security Workshop at the Network and System Security Conference, IEEE (2011)
2. Das, A., Bonneau, J., Caesar, M., Borisov, N., Wang, X.: The tangled web of password reuse. In: Proceedings of the Network and Distributed System Security Symposium. (2014)
3. Brown, A.S., Bracken, E., Zoccoli, S., Douglas, K.: Generating and remembering passwords. *Applied Cognitive Psychology* **18**(6) (2004) 641–651
4. Wiedenbeck, S., Waters, J., Birget, J.C., Brodskiy, A., Memon, N.: Authentication using graphical passwords: Basic results. In: Proceedings of HCI. (2005)
5. Golofit, K.: Click passwords under investigation. In: Proceedings of the 12th European Symposium on Research In Computer Security. (2007) 343–358
6. van Oorschot, P.C., Thorpe, J.: Exploiting predictability in click-based graphical passwords. *Journal of Computer Security* **19**(4) (2011) 669–702
7. Dirik, A.E., Memon, N., Birget, J.C.: Modeling user choice in the passpoints graphical password scheme. In: Proceedings of the 3rd Symposium on Usable Privacy and Security, ACM (2007) 20–28
8. Bixler, B.: Learning styles inventory. www.personal.psu.edu/bxb11/LSI/LSI.htm (n.d.) [Online; accessed 5-Jan-2015].
9. Weber, E.U., Blais, A.R., Betz, N.E.: A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors. *Journal of Behavioral Decision Making* **15**(4) (2002) 263–290
10. Field, A.: *Discovering Statistics Using SPSS*. 3 edn. Sage (2009)
11. Itti, L., Koch, C., Niebur, E.: A model of saliency-based visual attention for rapid scene analysis. *IEEE TPAMI* **20**(11) (1998) 1254–1259