Use Your Illusion: Secure Authentication Usable Anywhere

Eiji Hayashi
CMU/CyLab Japan
ehayashi@cmu.edu

Nicolas Christin
CMU/CyLab Japan
nicolasc@cmu.edu

Rachna Dhamija
Harvard University
rachna@deas.harvard.edu

Adrian Perrig
CMU/CyLab
perrig@cmu.edu

ABSTRACT

In this paper, we propose and evaluate Use Your Illusion, a novel mechanism for user authentication that is secure and usable regardless of the size of the device on which it is used. Our system relies on the human ability to recognize a degraded version of a previously seen image. We illustrate how distorted images can be used to maintain the usability of graphical password schemes while making them more resilient to social engineering or observation attacks. Because it is difficult to mentally "revert" a degraded image, without knowledge of the original image, our scheme provides a strong line of defense against impostor access, while preserving the desirable memorability properties of graphical password schemes.

Using low-fidelity tests to aid in the design, we implement prototypes of Use Your Illusion as i) an Ajax-based web service and ii) on Nokia N70 cellular phones. We conduct a between-subjects usability study of the cellular phone prototype with a total of 99 participants in two experiments. We demonstrate that, regardless of their age or gender, users are very skilled at recognizing degraded versions of self-chosen images, even on small displays and after time periods of one month. Our results indicate that graphical passwords with distorted images can achieve equivalent error rates to those using traditional images, but only when the original image is known.

Categories and Subject Descriptors

K.4.4 [**Electronic Commerce**]: Security; K.6.5 [**Security and Protection**]: Unauthorized access

General Terms

Design, Human factors, Security

Keywords

Graphical passwords, Social engineering, Distortion

Symposium on Usable Privacy and Security (SOUPS) 2008, July 23–25, 2008, Pittsburgh, PA USA
.

Figure 1: Example of the relationship between original and distorted pictures. It is difficult to infer the meaning of this distorted image without having seen the original image before (see Figure 2).

1. INTRODUCTION

Mobile phones are increasingly integrating security-sensitive services, such as electronic wallets (e.g., [24]) and banking transactions [2]. Secure use relies on the crucial assumption that the person using the device is the legitimate owner. In fact, in the majority of cases, a lost or stolen phone delegates all rights to its new "owner," putting the burden on the user to provide the security needed to avoid such situations.

In addition, many online services are now offering secure access to both web users and mobile phone users, usually through differing interfaces. However, the emergence and success of products like Apple's iPhone evidences that users appreciate a consistent experience, regardless of the device they use, be it a laptop, PDA, cell phone, or desktop computer.

In short, we not only need to secure user authentication on mobile devices, but we must aim for secure user authentication systems that work equally well on miniaturized devices as they would on large size displays.

Traditional authentication systems that rely on textual password entry, despite their simplicity, are not adequate for small portable devices. Indeed, the compact form factor severely limits the amount of data that can be displayed and makes data entry challenging. Because it is impractical to enter long strings of text, users will tend to pick short passwords, which are predictable and insecure.

This paper asks the following question: *Is it possible to devise an authentication scheme that works independently of the hardware size, is usable by a wide range of people, and yet is more secure than PIN or password-based authentication?*

Before we delve into the specifics of our proposed approach, consider Figure 1, which is a distortion of a photograph. Try to guess what the original photo represents, and then compare your guess to Figure 2, on the next page.

Figure 2: Example of the relationship between original and distorted pictures. This figure represents the original picture, before being distorted to obtain Figure 1.

Once you have seen both pictures, it is clear that Figure 1 represents the object in Figure 2, and that seeing Figure 1 evokes the original object. However, without having been exposed to Figure 2, Figure 1 does not carry much meaning.

People are able to capture the salient features from images they see, and can easily recognize these features when a considerably degraded version of the image is presented to them. However, it is difficult to mentally "revert" from the degraded image to the original image, without having been exposed to the original picture.

As a mathematical analogy, image degradation can be viewed as a one-way function. By associating the original picture and the degraded version of that picture, the legitimate user can mentally revert back from the degraded version even when the original picture is not shown. However, this transform is difficult to perform without knowledge of the original image. A legitimate user can "see" something that is not there, which gives her an advantage over unauthorized parties.

This cognitive process allows us to develop an authentication scheme, called *Use Your Illusion*, that reconciles usability with strong security. Image degradation provides an effective line of defense against impostors, while being largely unaffected by a low graphical resolution. These properties make Use Your Illusion applicable to a wide range of contexts, such as computer systems, ATMs and cellular phones.

Use Your Illusion share some of the desirable properties of traditional graphical passwords schemes, such as DéjàVu [8] or Passfaces [21], that require users to recognize images rather than text. In particular, humans are better skilled at remembering pictures they have previously seen than they are at remembering complex text [23], which allows for stronger authentication tokens.

However, different from existing graphical password schemes, Use Your Illusion avoids the trade-off between memorability and vulnerability to social engineering that existing graphical password systems face. Indeed, schemes like Passfaces, where users can choose their own image portfolios, have shown to lead users to pick very predictable images that can be easily guessed by an attacker [20]. Use Your Illusion fully exploits the human ability to recognize degraded images, which considerably improves resilience to guessing (or social-engineering) attacks and observation attacks. At the same time, by allowing users to select pictures of their choosing, Use Your Illusion maintains high memorability properties, superior to that of schemes like DéjàVu, which rely on assigned, random pictures. In short, Use Your Illusion, Use Your Illusion addresses what has been perceived as the Achilles' heel of graphical password schemes, by removing the perceived trade-off between portfolio security and memorability.

By describing the design, implementation, and evaluation of Use Your Illusion, this paper makes several contributions. First, we propose a novel user authentication scheme based on the human ability to recognize degraded images. Second, we demonstrate the portability of the scheme by developing two prototypes – an Ajax-based web service and a Nokia N70 cellular phone implementation. Finally, we describe the design and results of a usability study that illustrates the viability of the scheme.

The rest of this paper is organized as follows. We next present a problem definition in the form of security requirements, and describe how related proposals have attempted to meet them. In the following sections, we show how we build Use Your Illusion, from the cognitive process to the design of the actual prototypes. We then describe a usability study, and provide empirical measurements of the effectiveness of our scheme. We conclude by listing some of the challenges that we still face, and summarizing our contributions.

2. REQUIREMENTS AND RELATED WORK

An ideal authentication system should provide strong security while maintaining high usability – it should be usable everywhere, by everyone, without the need for any specific training. In this section, we formalize our desired system requirements. We then discuss existing authentication proposals, in the context of our requirements.

2.1 Requirements

No unauthorized access. An attacker that attempts to impersonate a legitimate user should not have a higher probability of choosing authentication tokens than he would have if he chose them at random. For example, an authentication system that required a four-digit PIN authentication system and allowed users three attempts to enter their PIN would satisfy this property only if the probability an impostor gains access to the system does not exceed 3/10,000.

Universal use. The authentication system should produce consistent results regardless of the physical environment in which it is used. For example, the system should not be affected by lighting conditions or surrounding noise.

Limited scale. The authentication system should be usable on typical cellular phones or devices of similar size and capabilities. Specifically, the system should work with small color displays of relatively low resolution (e.g., 350x400), and at most a dozen of keys (e.g., numeric pad and function keys). The processor and memory requirements should not exceed the resources available in medium-priced cellular phones.

Limited training. The authentication system should not require users to undergo a long or intensive training session prior to first use. Ideally, the system should be usable in a few minutes by a new user with minimal external assistance.

Rapid setup. Creation of a new authentication tokens should be fast, so that users can create accounts and reset their authentication tokens with ease.

Rapid authentication. A legitimate user should be able to authenticate herself to the system in seconds.

Low error rates. The number of occurrences where the authentication system denies access to a legitimate user should be held at low levels (e.g., comparable with PIN-based authentication systems where users have the ability to select PINs at their discretion). Thus, the system itself must have a negligible (or null) error rate, and legitimate users should make mistakes infrequently.

High memorability. As a corollary to the above requirement, we need the system to use authentication tokens that are easily memorable. This requirement also implies that the number of authentication tokens to remember must be held small. Using the authentication system in several different contexts (e.g., online banking, instant messaging, ...), should not require significant additional efforts from the user compared to using the authentication scheme in

a single context. Further, authentication tokens should be memorable enough to be easily identified even in cases of rare use, i.e., when the interval between two distinct authentication sessions is of several weeks.

2.2 Related Work

Authentication schemes are often characterized as "something you know" (e.g., passwords and PINs), "something you have" (e.g., physical authentication tokens), or "something you are" (e.g., biometrics).

Biometric authentication may satisfy the *limited training*, *rapid setup* and *fast authentication process* properties. However, existing technology can not accommodate the *universal use*, *low error rate* and *no unauthorized access* properties simultaneously. For instance, voice authentication has significant error rates in noisy environments, facial recognition schemes are sensitive to variations in lighting conditions, and fingerprint readers can be defeated by fake fingerprints [18]. Physical authentication tokens satisfy the *limited training*, *low error rate* and *high memorability* properties. However, they are highly vulnerable to theft, and therefore have trouble satisfying the *no unauthorized access* property: as soon as a physical authentication token is stolen, an attacker is able to impersonate its legitimate user. Hence, in the rest of this section, we focus on knowledge-based authentication systems, which seem to offer the most promise to fulfill all of the properties we set out to achieve.

Knowledge-based authentication systems usually satisfy *limited training*, *fast authentication process* and *universal use*, with some exceptions noted below. The question is how good are the proposed schemes at avoiding unauthorized accesses and limiting mistaken access denials. Furthermore, how do the schemes work across a variety of device types?

Human memory imposes several limitations on the effectiveness of knowledge-based authentication systems. Traditional passwords and PINs depend on *recall*, the ability to remember items from memory without help. Even if it were possible to easily enter long strings of random symbols into cellular phones, people cannot remember them easily.

A number of graphical authentication schemes have been proposed (e.g., [8, 16, 31]), based on the observation that humans are considerably better at remembering images than they are at remembering text [23]. The schemes attempt to address the limits of human memory by relying on different cognitive processes.

Some graphical password schemes rely on recall. For example the Draw-A-Secret scheme requires users to draw an image on a grid in order to authenticate [16]. This is a recall task, because the user has to recollect the drawing from memory during each authentication. Thorpe et al. show that the actual Draw-A-Secret password space is smaller than theoretically possible, because users tend to choose symmetric passwords with a small number of strokes, which are easier to recall [27, 28].

Thus, graphical passwords based on recall may have many of same problems of textual passwords, and they do not satisfy our *no unauthorized access* requirement. Furthermore, even though some smartphones have a stylus, most cell phones do not have pointing devices, so this scheme does not satisfy our *limited scale* requirement.

Other authentication systems rely on *cued recall*, where clues are provided to the user to aide the recollection task. For example, some graphical password schemes use textual or mnemonic hints to help users remember their passwords [19]. Many graphical password schemes require users to authenticate by clicking on points, or selecting regions, of an image that were previously chosen by the user [4, 31, 32]. In this case, the image itself serves as a clue to the regions of the image that a user must recall. Because users have to recall the points they have chosen, there tend to be "hot spots," or regions that are often selected because they are most memorable or most obvious, thus leading to predictable passwords [29, 32]. Another disadvantage is that click-based schemes may need a relatively large display to provide a large enough key space, and they require a pointing device. As such, they do not fulfill our *limited scale* requirement.

Recognition, the ability to remember items we have seen before, is an easier memory task than recall. In particular, humans have an impressive ability to recognize pictures they have seen before, and they recognize pictures better than they recognize words or sentences [23]. Research has shown that individuals can distinguish large sets of previously seen pictures from new "distractor" pictures at high levels of accuracy [23, 25]. Many graphical password schemes rely on this skill by asking users to select previously chosen images from a larger subset of images [6, 8].

Research shows that people remember images more accurately when they are semantically meaningful and when the images are generated by people themselves [17]. Therefore, graphical password images self-chosen by users may be more memorable. However, given some knowledge about the user, self-chosen images are also easier for an attacker to predict [20], which defeats our *no unauthorized access* requirement. One countermeasure is to assign pictures to users rather than allow them to choose images themselves. Another option is to use abstract images, such as Random Art, which are less predictable than real images [8]. However, users tend to prefer meaningful images [8], and research has shown that image scenes that are coherent and semantically meaningful are stored more accurately than incoherent or abstract images [5, 10].

Another alternative, closer to our proposal, but used to combat shoulder-surfing, is to conceal as much information as possible from an image by obscuring it with noise [13]. While adding noise is effective at preventing shoulder surfing, legitimate users have difficulty recognizing the obscured pictures, and authentication failure rates could be high.

One graphical authentication scheme is based on computing a "path" of known pictures [30]. This approach seemed promising, with error rates in the order of 5%, and no unauthorized accesses, but has recently been shown to be vulnerable to an eavesdropping attack [11]. In addition, this scheme requires a large screen, long authentication times (in the order of minutes), and users need to sustain significant training prior to use.

In summary, all previously proposed schemes are not able to simultaneously meet the requirements of no authorized access, universal use, limited scale, limited training, rapid setup, rapid authentication and low error rates.

3. USE YOUR ILLUSION

3.1 Overview

Use Your Illusion is a graphical password scheme that allows the use of images that are self-chosen and familiar to the user, yet that are not easily predictable by an attacker.

In our system, we allow users to select their own graphical password images. Psychology studies show that images that are self-generated are recognized better than those that are not [17]. In addition, users enjoy the ability to select and personalize their image portfolios, and they tend to choose images that are semantically meaningful to them [8, 9, 20]. There are several option for allowing users to generate images (e.g., by selecting them from an existing database of images or by uploading pictures from the user's com-

puter). In the context of mobile devices, the user can capture their own images with a camera embedded to the mobile device.

Once the pictures have been selected, they are distorted using a non-photorealistic rendering algorithm that eliminates most details in the image, while preserving some features such as color and rough shapes. Because information is lost in the rendering algorithm, it is impossible to mathematically revert back to the original image from its distorted version. As such, the distortion function is analogous to one-way functions used in cryptography.

Next, we prime the user during a training session to associate the distorted image with the original image and the meaning of that image. During the priming phase, we display the original and distorted pictures side-by-side, and ask the user to practice selecting their images from a set of "distractor images."

To authenticate, the user must choose her own distorted images from a set of distractor images. We rely on the fact that human perception is affected by what we know. When the meaning of the image is known, our brains impose that knowledge on our perception and it becomes hard to interpret the picture in another way [12]. Furthermore, there is evidence that the ability to recognize objects in degraded images increases dramatically with familiarity of the subject in the image (this effect is very strong in the case of faces [7, 14]). During the authentication task, the user can recognize the original objects in the image by using color and shape cues and by remembering the semantic meaning that she previously associated with the image.

Next we discuss the overall architecture of Use Your Illusion and elaborate on the threat model and possible attacks against the scheme. We then provide details about our prototype implementation, where we pay particular attention to accommodating the limitations of a small device.

Figure 3: Image portfolio assignment. In this example, the user has been asked to take $p = 3$ pictures and is subsequently presented with the three pairs of original images-distorted images. The distorted images will be used as authentication tokens in the authentication phase.

3.2 System Architecture

During account setup, the user must select a personal *image portfolio* of p pictures. To authenticate, we present the user with a challenge set of n pictures, where $n > p$. The user has to correctly identify the p pictures within the challenge set that belongs to her portfolio.

Use Your Illusion consists of three phases: portfolio creation, practice and authentication.

Portfolio creation phase To create an image portfolio, the user first has to create a set of p images. To enhance memorability, users are encouraged to create their portfolio pictures, rather than using default pictures. Here, we asked users to capture p photos they want to use for authentication. The photos should be taken in as secure an environment as possible; in particular, these photos should not be transferred to the authentication device using an insecure channel. Ideally, the user should be able to use the authentication device itself to capture the images. For instance, when Use Your Illusion is deployed on a mobile phone, the mobile phone camera is the best option.[1]

Once the photos are taken and passed to the authentication device (e.g., the cell phone itself), they are distorted using a lossy filter. Using a lossy filter ensures that it is mathematically impossible to revert from the distorted image to the original image. The transform is performed on the authentication device.

Use Your Illusion does not mandate a specific type of filter. The selection of the "best" filter possible is an area that warrants further investigation, but cognitive research can give us some heuristics: the rough shape and colors of the original picture should be preserved to make the distorted picture more memorable. When considering small portable devices like cellular phones or PDAs, the constraints imposed by the size and quality of the display require that the distortion filter can work with low resolutions. In other words, a distortion filter that yields a high resolution output is probably unsuitable.

The resulting set of p distorted pictures is assigned to the user as her image portfolio. The original pictures and distorted pictures are shown simultaneously to the user, so that the user can mentally associate the distorted pictures with their original meaning. Figure 3 is a screen shot of the image portfolio assignment.

Concomitant with the portfolio creation phase, the authentication device selects $(n-p)$ pictures to be used as decoys during the authentication phase. The $(n-p)$ decoys *do not change* until the portfolio images are revoked and new ones are recreated. If decoys were changed from one authentication challenge to the next, an attacker could indeed infer the portfolio images by simply observing several authentication challenges in succession, and identifying as portfolio images the pictures that are constantly displayed across all challenges. This type of attacks is called intersection attacks, on which we elaborate in the next section.

There are two possible approaches are possible to generate decoy images. One approach is to generate "synthetic" decoys purely algorithmically, that is, without using any original photo. Another approach is to generate decoys by applying the lossy filter to a set of existing photos that were not taken by the user. Generating convincing synthetic decoys remains an open problem. Indeed, we found that synthetic decoys look very different from distorted images and could easily be detected, thereby immediately revealing the pictures chosen by the user. Hence, we suggest using distorted versions of photos the user did not take. In a cell phone, the manufacturer could store a number $N \gg n$ of decoys to secure memory at manufacturing time. As long as the decoys images are selected at random from a database of $N' \gg N$ images that evolves over time, decoy images should be hard to identify.

Training phase After the portfolio creation phase, the system conducts a short training phase to improve memorability of the portfolio images. In this phase, a challenge set of portfolio images and decoy images is presented to the user. The user can practice se-

[1]While there are still some mobile phone models that do not have an embedded camera, these models generally do not support applications that require strong user authentication.

lecting their image portfolio, and the system provides immediate feedback on whether the image is correct. The user can access original-distorted image pairs at any time. Because of this, the training phase should be conducted in as secure an environment as possible.

After the training phase is complete and the user is confident that she remembers her images, the original pictures used in the generation of the distorted pictures should be removed from the cell phone memory.

Authentication phase During the authentication phase, the user must correctly select her p portfolio pictures from the challenge set. The decoy images themselves are produced using original pictures, and the distortion levels are high enough that most details of the original pictures are obscured. Therefore, an outsider will have a very hard time identifying which pictures belong to the portfolio, even if she possesses information about the user's personal preferences. Contrary to the decoy database, the challenge set of n pictures (portfolio and chosen decoys) presented in the authentication phase does not need to be stored in secure memory.

3.3 Attacks and Countermeasures

We next investigate how Use Your Illusion addresses possible attacks aimed at impersonating a legitimate user.

Brute force attack The simplest attack is to try to randomly guess the correct portfolio. With a challenge set of n pictures, and a portfolio of p pictures, the probability that a single random guess succeeds is $1/\binom{n}{p}$.

Obviously, if we allow the impostor to try all possible $\binom{n}{p}$ combinations, then she will eventually manage to fraudulently authenticate. To prevent such an undesirable outcome, we use a reference counter that locks the device after t failed authentication attempts, similar to systems used in automated teller machines. A locked cell phone would need to be returned to the manufacturer or service provider for unlocking.

The probability that attacker can impersonate the user within t trials, using a succession of random guesses, is $t/\binom{n}{p}$. For example, for $(n,p,t) = (27,3,3)$, we get $3/\binom{27}{3} \approx 0.001$. We can adjust (n,p,t) according to the desired failure probability. For instance, increasing n to $n = 36$ yields a failure probability less than that of a random four-digit PIN-based system.

We note that the actual failure probability required is, in practice, going to depend on the application considered, and the adversarial model. For instance, a cellular phone used for online banking should require a much lower failure probability than an online video game access on a fixed terminal, considering both the effects of a failure (losing money vs. having to create a new character), and the potential failure modes (cell phone stolen or lost vs. active intrusion by third-party).

Educated guess In an educated guess attack, the impostor tries to guess the user's portfolio pictures based on previously obtained information about the user, e.g., through social engineering.

For instance, in an image-based authentication scheme that allows a user to take a picture by herself and to use the picture as is as part of her portfolio, the following scenario is possible. Assume the attacker has previously learned that the user owns a white dog, Fifi. If the attacker finds a picture of a white dog in the challenge set, the attacker can guess the picture is actually included in the user's portfolio.

In Use Your Illusion, similar to the above example, users create their own pictures. However, Use Your Illusion never uses the original pictures beyond the practice phase (and the original pictures are destroyed as soon as the training phase is completed). Because Use Your Illusion only utilizes distorted images resulting from a lossy filter transform, an educated guess attack is less likely to succeed. In the above example, the distorted picture of Fifi in the user's portfolio will resemble nothing more than a predominantly white blob. Even if the impostor knows of the existence of Fifi, figuring out that the white blob originally came from a picture of Fifi, and not from a picture of a chicken or a snowman, may be difficult. Thus, the educated guess attack is much more difficult to carry out in Use Your Illusion than in almost all image-based password authentication schemes. Of course, more difficult does not mean impossible: if familiar objects can be easily identified even after distortion, the attacker may receive clues that the image is part of the user's portfolio.

Observer attack In an observer attack, the attacker identifies the pictures in the user's portfolio by observing authentication procedures of the user. Observer attacks, or "shoulder surfing" attacks, are currently one of the most significant threats to user authentication. For example, there are many reports of ATMs equipped with rogue video cameras that record authentication sessions [3]. We propose two countermeasures to mitigate the threat of observer attacks.

First, we constantly change the respective positions of the decoys and portfolio pictures on the authentication screen(s), so that the authentication pattern cannot be inferred by observing which keys are pressed. This line of defense is particularly useful when combined with the use of optical filters that render the display difficult to observe from a distance.

Second, we avoid showing any hints regarding the picture selection. In particular, no feedback is given to the user when a picture is selected. In a cell phone environment, we have the added benefit that it is hard for an observer to observe which keys are pressed, given the small size of the keyboard or display. Without any correlation between the keys pressed and the output on the display, an observer cannot identify the user's portfolio.

While we do believe these defenses are likely to be adequate for small mobile devices, we elaborate on how Use Your Illusion can help in designing schemes with strong resilience to observation attacks in a separate study [22]. In short, we show in [22] that the combination of Use Your Illusion with a tactile device allows to achieve strong resilience against observation attacks while maintaining high usability.

Intersection attack In an intersection attack, the impostor identifies the legitimate user's portfolio pictures by observing multiple authentication session; the attacker can then take the intersection of all of the images that are observed to reveal the user's portfolio [8].

For simplicity, assume a challenge set of size $n = 3$, and a portfolio set of size $p = 1$, and denote the portfolio picture by X. On a first authentication attempt, the challenge set may consist of $\{A,B,X\}$. Now, if decoys are changed from one authentication session to the next, the challenge set proposed in a subsequent authentication may be $\{D,X,C\}$. An impostor only needs to look at both challenges to figure out the portfolio picture is $\{A,B,X\} \cap \{D,X,C\} = \{X\}$.

In Use Your Illusion, we resist the intersection attack by always maintaining identical decoys in each authentication challenge. In the above example, the challenge set would always be $\{A,B,X\}$. A and B would only be replaced when the user decides to change her portfolio image from X to Y. Therefore, each authentication challenge is always the same, and the intersection of many challenges reveals nothing about the portfolio.

In summary, Use Your Illusion appears resistant to brute force, guessing, prediction, and casual shoulder surfing. Use Your Illusion is more vulnerable to prolonged observation attacks and to spyware, even though some of the countermeasures used against

observer attacks increase the difficulty of carrying out such attacks. For instance, a simple keystroke logger on the device would not defeat the scheme, because the order of pictures in the challenge set changes from one authentication session to the next. In order to defeat the scheme, spyware must simultaneously record the contents of the display. Additional countermeasures against observation attacks using Use Your Illusion as a basis have been successfully devised [22].

A criticism of Use Your Illusion may be that it only provides marginally more resilience against observer attacks than PINs. However, the main objectives of Use Your Illusion is to thwart social engineering and guessing attacks, which are very easy to carry out with self-chosen passwords or PINs, while also addressing the human memory limitations and text entry problems encountered with PINs and passwords.

3.4 Prototype

We implemented both web-based and mobile device prototypes as a proof-of-concept and to conduct usability tests. The mobile device prototype is implemented in Java on Symbian OS, and has been tested on Nokia N70 cell phones. We selected the Nokia N70 handset because its core features (display, camera, keypad, memory) are fairly common and are of reasonably high quality. The web-based demonstration of the prototype is available at `http://arima.okoze.net/illusion/`, and attempts to closely emulate the Nokia N70 interface. Our online version has shown to work on the Apple iPhone as well.

The user can select three portfolio pictures ($p = 3$), derived from photos she has taken herself, as shown in Figure 3. In the authentication phase, the system present three sets of nine pictures to the user as a challenge set ($n = 27$). The user can make three attempts to login ($t = 3$). After three consecutive authentication failures, the phone is locked. The probability that a random guess results in a successful authentication is $3/\binom{27}{3} \approx 0.001$. As discussed before, we can lower that probability by increasing the number of portfolio pictures and/or decoys.

Figure 4 illustrates the challenge set for the user who had selected the portfolio shown in Figure 3. The challenge set is divided into three authentication screens containing nine pictures each. On a given authentication screen, an indicator at the top of the screen denotes the screen the user is currently seeing. As long as they are not authenticated, users can move forward or backward between the different authentication screens, using function keys present on the handset. The nine challenge set pictures in each screen are arranged in a 3x3 grid, and each of the nine pictures is associated with a numeric key. Some screens may only contain decoys, as is the case of the leftmost screen in this example, while some screens may contain more than one portfolio image, as in the case of the middle screen here. Finally, the "Clear" button allows users to reset their current selection and start the authentication process anew.

As discussed above, to avoid intersection attacks, the same challenge set will be presented as long as the portfolio remains unchanged. On the other hand, the order in which pictures are presented will differ from one authentication session to the next.

Image processing filter In our prototype, we select an *oil-painting filter* [15] as the lossy filter used to generate portfolio images from the original photos input by the user. A oil-painting filter blurs the edges and the colors of the processed picture while preserving the main shapes and colors of the original picture. The key parameter in the oil-painting filter is the size of the "brush" the filter uses to distort the original picture. Shortly stated, larger brush sizes reflect greater distortion.

Properly tuning the brush size is an important and non-trivial task: If the distortion level is too high, a legitimate user cannot easily recognize a distorted picture. If the distortion level is too low, an attacker may be able to guess the subject of a distorted picture, without knowing the original picture. To determine a suitable distortion level, we conducted a low-fidelity test which we discuss in the next section.

4. USABILITY EVALUATION

We conduct usability tests to help in the design and evaluate the effectiveness of the proposed scheme. We begin by conducting an informal, low fidelity test to calibrate the oil painting filter used in our prototype. We then conduct a formal usability experiment, using a between-subjects design. We complement this experiment with an additional test to determine possible age effects.

4.1 Low-fidelity testing

We use a low-fidelity test with six participants, to decide on the optimal distortion level of the oil painting filter. We show each user a sheet of paper that contains an original picture and a set of distorted pictures, in which we vary the distortion level by changing the brush size.

We conduct the test in two phases. First we show distorted images to participants, when they have no knowledge of the original image. We progressively show images to participants, from the most distorted to the least distorted, to determine when they can recognize the subject in the image. Next, we show participants an original image and increasingly distorted versions of that image, to determine when the distorted image is no longer recognizable.

We discover that prior knowledge of the image increases the level of distortion that we can apply before the image becomes unrecognizable. Our findings lead us to determine that for 56x56 images, the optimal brush size should be comprised between 6 and 8 pixels. In this range, users who know the original picture can recognize its distorted version, while those who do not know the original picture cannot make sense out of the distorted image. We accordingly set the brush size to 8 pixels in the prototype.

4.2 Usability Testing

We next conduct a formal usability test to evaluate the Use Your Illusion prototype. We design our study to answer the following questions: How does distorting images affect the memorability of images in a graphical password scheme? How much does allowing the user to self-select an original image aid in the memorability of a distorted image? If we find that adding image distortion has no effect on the authentication success rate, then Use Your Illusion can reap the memorability benefits of previously proposed graphical password schemes [8,20], while being significantly more resilient to guessing attacks.

For this experiment, we recruit participants by posting flyers in four universities. 54 people participated in our study. Of these, 50 are students and 4 are university staff. Nine participants are female and 45 participants are male. All of the students major in electrical engineering, computer science or another scientific subject. Their ages range from 18 to 29 with a mean age of 23. None of these participants has taken part in our low-fidelity test. All usability tests, for all participants, are conducted by the same researcher.

We split participants in three groups. To avoid any bias due to factors such as memory, eyesight and/or familiarity with cell phones, we distribute participants evenly over the three groups. That is, we ensure that different groups present similar age, major, and origin (i.e., university) distributions.

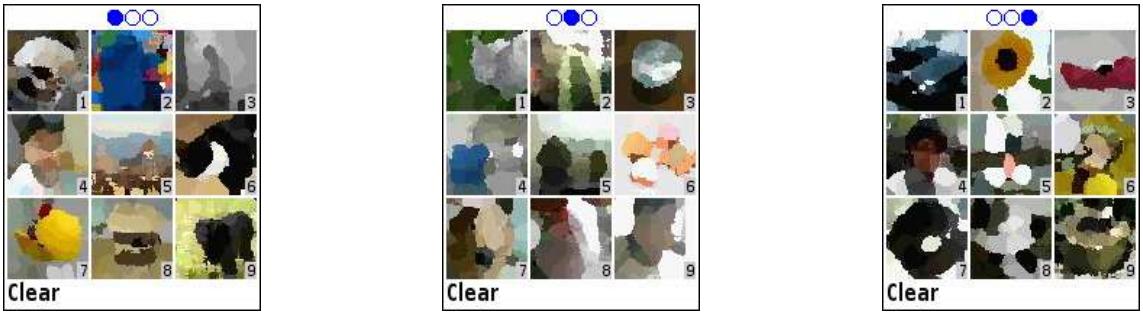

Figure 4: Example of a challenge set of images. This series of screenshots represents a challenge set as presented to the user on a Nokia N70 cellular phone. The user can navigate between different screens by using function keys on the handset.

To investigate how success rate and login time are affected by the type of images used, we assign each group a different portfolio type:

Self-selected, non-distorted: We ask participants in this group to take three pictures using a cell phone camera. We then assign the three pictures, *as is*, as their image portfolio. We also assign 24 decoy photographic images that are identical for each participant.

Self-selected, distorted: This is the "Use Your Illusion group." We ask participants in this group to choose three of their own pictures. They could select the three pictures from pictures they have already taken or they can capture new pictures for this usability test. We distort each of the pictures with an oil-painting filter. We then assign the distorted pictures as the image portfolio. We also assign 24 decoy photographic images that are distorted using the same filter settings. The decoys are the same for each participant.

Imposed, distorted: We assign participants in this group the same portfolio images as those imposed to participants in the "self-assigned, distorted" group. The only difference is that participants in this group do not take the pictures by themselves (they are imposed by the experimenter), and they do not see the original pictures. The decoy images are the same for each participant. The decoy images used in this group are identical to those used in the "self-selected, distorted" group.

Procedure Our usability test spans four weeks and consists of four sessions. All sessions are conducted in a university classroom. Each of the participants takes part in the test individually using the same phone and the same version of the prototype. Table 1 shows when each session occurs and the tasks that participants completes in each session. In the first session, we assign three portfolio images to participants according to the group to which they were imposed. Participants then complete the training phase described in Section 3, which lasts up to five minutes.

Session #	Date	Tasks
1	First day	Memorize portfolio Training Authenticate
2	Two days later	Authenticate
3	One week later	Authenticate
4	Four weeks later	Authenticate Fill out questionnaire

Table 1: Tasks completed in each session of the usability test.

After the training phase, we instruct participants to login by selecting their portfolio pictures from a challenge set. As discussed in Section 3, the prototype requires participants to authenticate by selecting their three portfolio images out of a set of 27 images.

In the second, third and fourth sessions, we ask participants to authenticate again by selecting their portfolio. At the end of the fourth session, we have participants fill out questionnaires for the sake of qualitative analysis.

Success rate We consider an authentication to be successful if the participant can correct identify their own portfolio images. We give participants a maximum of three attempts to login.

Table 2 shows success rates measured in the sessions. All participants in the two "self-selected" groups succeed in all sessions. On the other hand, after four weeks, the success rate of the "imposed, distorted" group decreases to 89%. Participants imposed to use distorted image portfolios, without knowledge of the original image, have lower success rates than the other two groups.

These results suggest that the graphical passwords with distorted image portfolios may achieve equivalent error rates to those that use traditional images, but only when the user knows the original image.

	The 1st day	2 days later	1 wk. later	4 wks. later
Self-selected, non-distorted	100%(18)	100%(18)	100%(18)	100%(18)
Self-selected, distorted	100%(18)	100%(18)	100%(18)	100%(18)
Imposed, distorted	100%(18)	89%(16)	94%(17)	89%(16)

Table 2: Percentage of successful logins. The number of participants with successful logins is included in parenthesis. All groups have 18 participants.

Table 3 shows the number of participants who successfully authenticate during the fourth testing session, which occurs four weeks after creating the portfolio. The table shows how many attempts it takes to succeed. This fourth testing session had the highest failure rates, so we omit tables corresponding to the first three sessions, since user performance was higher. If a participant cannot login within three attempts, we consider the authentication task a failure. All participants in the "self-selected, non-distorted" group are able to authenticate on the first attempt. In the "self-selected, distorted" group, 17 participants succeed in one attempt, and one succeeds in the second attempt. In the "imposed, distorted" group, 13 participants succeed in the first attempt, three participants succeed in the second attempt and two participants fail to authenticate.

These results suggest that randomly imposed abstract images are difficult for some users to remember, even after a few attempts. However, users may be able to recover better from errors when the original image is known and when the image is self-selected.

	Attempt 1	Attempt 2	Attempt 3	Failed
Self-selected, non-distorted	18	-	-	-
Self-selected, distorted	17	1	-	-
Imposed, distorted	13	3	-	2

Table 3: Number of successful authentications after four weeks. The table presents the number of participants that succeed in each attempt during the fourth session (held four weeks after portfolio assignment).

	1st day	2 days later	1 wk. later	4 wks. later
Self-selected, non-distorted	11.5 (9.9)	12.3 (12.3)	12.7 (11.9)	12.5 (12.8)
Self-selected, distorted	12.4 (11.2)	16.4 (15.9)	14.3 (13.4)	17.9 (16.5)
Imposed, distorted	16.7 (14.1)	25.8 (19.0)	25.1 (17.6)	24.7 (16.7)

Table 4: Login time. The table gives the mean time, in seconds, users need to authenticate. Median login times are given in parentheses.

Login time Table 4 shows the means and medians of the login times in seconds. Time is measured cumulatively, meaning, the clock is *not* reset after a failed login attempt but instead runs until the user can successfully log in, or has failed to authenticate three consecutive times. Left and right values in the cells stand for mean and median respectively. We do not include the results for the participants who could not login at all. The difference in mean login times between groups is not statistically significant. However, we observed that participants in the "self-selected, distorted" group require a longer time to login that participants in the "self-selected, non-distorted" group, and they require less time to login than participants in the "imposed, distorted" group.

There is a considerable gap of 3 weeks between the third session and the fourth session. The login times for the "self-imposed, non-distorted" stay fairly consistent over the course of 4 weeks. The login time for the "self-imposed, distorted" group increases slightly after the course of a week and then rises by approximately 4 seconds after the three week gap. Compare this with the login time for the "imposed, distorted" group, which increases by more than 8 seconds after the first day and stays consistently at that level after 4 weeks.

Table 5 shows the mean number of times participants switch between pages in each test session. Recall that the 27 portfolio images are displayed on the cellular phone display 9 at a time; the participant must switch to the next page to see the next 9 images. When a participant can easily remember her portfolio, she can quickly choose her image and advance to the next page. As a result, the number of page advances would be at most two. On the other hand, when a participant can not recall her images well, she may have to switch back and forth several times. Obviously, page switching also increases the amount of time required to login. Our results show that on average, after 4 weeks, participants in the "self-selected, distorted" group require one more page switch than those in the "self-selected, non-distorted" group, while participants in the "imposed, distorted" group require approximately 3 more page switches than those in the "self-selected, non-distorted" group.

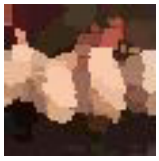
(a) People


(b) Shrimp dumplings


(c) Panda


(d) Battery

Figure 5: Examples of incorrect semantic meanings. Users unfamiliar with the original picture assign incorrect semantic meanings to the distorted pictures.

Overall, our results indicate that it is considerably more difficult to remember *imposed* distorted images than self-selected original images. On the other hand, self-selected original images are not made much more difficult to remember by adding distortion.

	The 1st day	2 days later	1 wk. later	4 wks. later
Self-selected, non-distorted	2.28	1.83	2.17	2.22
Self-selected, distorted	2.00	3.05	2.33	3.28
Imposed, distorted	2.94	4.89	4.39	5.11

Table 5: Mean number of page switches for all participants during the fourth session

Qualitative results In the post-experiment questionnaires, we ask all participants to evaluate how easily they can identify their portfolio images using a 5-point scale, where 1 is "very easy" and 5 is "very difficult". The average difficulty scores are 1.63 for the "self-selected, non-distorted" group, 1.59 for the "self-selected, distorted" group, and 2.17 for the "imposed, distorted" group.

We also ask the participants about the techniques they use to remember their portfolio images. 14 out of 18 participants using "self-selected, distorted" pictures reply that that they can "see" their original images within the distorted pictures and that they simply remember the meaning of the image. Four participants say that they memorize characteristic colors and shapes within the images. Interestingly, 12 our of 18 participants in the "imposed, distorted" group also indicate that they assign a semantic meaning to their portfolio images in order to memorize them, which mirrors observations from Stubblefield and Simon [26] on "inkblot authentication." Of these, seven participants incorrectly guess the meaning of the image. Four participants guess the subject correctly for one out of three portfolio images. Only one participant guesses correctly for all three of her portfolio images. As shown in Figures 5 and 6 incorrect guesses include mistaking shrimp dumplings for people and a battery for a panda. Some of the correct guesses include pictures of Winnie the Pooh and a wall clock.[2] In the next section, we discuss the factors that can influence the correctness

[2]Many other incorrect guesses, e.g., mistaking a duck for a dog, or correct guesses, e.g., a colorful T-shirt, are not represented here.

(a) Winnie the Pooh

(b) Wall Clock

Figure 6: Examples of correct semantic meanings assigned by participants in the "imposed, distorted" group. Some users are able to recognize Winnie the Pooh and the wall clock, even though they are not exposed to the original picture.

of a guess and its impact on security. Participants were asked to take pictures before the first test took place, and were given ample time to choose whichever pictures they saw fit. We find that 90% of the participants choose pictures whose subjects are things they see regularly in their daily life, such as pets, cars, and coffee mugs. Finally, many participants in the "self-selected, distorted" group indicate that they considered Use Your Illusion more as a game than as an authentication system in that it was fun and quite enjoyable to use.

4.3 Effects of age and distortion

While giving a good overview of the potential of Use Your Illusion, our main usability test presented above has a few shortcomings. Some of the differences (e.g., login times) between the different groups are not statistically significant, and the population sample is biased toward younger, technically-inclined subjects. Here, we complement the results obtained in our usability test with additional measurements gathered in the course of a separate usability study.

This complementary test concerns 45 participants (9 female, 36 male), none of which have taken part in the main usability study described above. All participants are over 18 years old. 5 participants are below 20 years old, 12 participants are between 20 and 29 years old, 18 participants and between 30 and 39, 6 participants are between 40 and 49, and 4 participants are over 50. Participants are a mix of undergraduate and graduate students, faculty, and (in majority) office workers in a public administration.

The experiment setup and procedure is similar to that used in the main usability test described in Section 4.2, with the two key differences that:

- There are only three sessions spanning, in total, a week (i.e., the fourth, "4-week after" session is not conducted here), and
- The "imposed, distorted" group is replaced by a "imposed, random" group, whose participants are assigned very highly distorted versions of pictures chosen at random from a set of about 35,000 pictures available on Flickr [1] under a Creative Commons License allowing commercial use and modifications, and marked as "camera-phone" pictures using the appropriate Flickr tag. The goal was to evaluate user performance when facing pictures to which a semantic meaning is difficult to assign.

	The 1st day	2 days later	1 wk. later
Self-selected, non-distorted	100% (15)	100%(15)	100% (15)
Self-selected, distorted	100% (15)	100% (15)	100% (15)
Imposed, random	93% (14)	73% (11)	73% (11)

Table 6: Percentage of successful logins (complementary test). The number of participants with successful logins is included in parenthesis. All groups have 15 participants.

Table 6 shows the results obtained in this complementary usability test echo that gathered in the main usability test (see Table 2), which provides some insights that Use Your Illusion is usable regardless of age, or gender. Indeed, the differences observed between male and female subjects both in the main usability test and this complementary experiment are smaller than the differences between conditions. Likewise, participant background seems to have little, if any, incidence on overall performance.

Figure 7 plots the median authentication times against the age of our participants. For all authentication sessions, the average authentication time generally slightly increases with age for self-selected pictures. The effect is more marked when distortion is used, that is, for the group using Use Your Illusion. Further, the authentication times with self-selected pictures are considerably smaller than those obtained with imposed pictures.

Finally, the data obtained seems to indicate that the median authentication times in the groups using self-selected pictures is the same, regardless of whether distortion is used or not. To avoid making any assumption on the distribution of authentication times, we use a Mann-Whitney U test to test this hypothesis, and obtain $U_{\text{first day}} = 94$, $U_{\text{second day}} = 77$, $U_{\text{one week}} = 83$. With $n_1 = n_2 = 15$ participants in each group, we see that there is no difference in median authentication times between the "self-selected, non-distorted" group and the "self-selected, distorted" group with a significance level of 0.01. On the other hand, using the Mann-Whitney U test to compare the "self-selected, distorted" group and the "imposed, random" group shows that, at significance level 0.01, the mean authentication time of the "imposed, random" group is higher ($U_{\text{first day}} = 39$, $U_{\text{second day}} = 39$, $U_{\text{one week}} = 23$ here).

In short, *distortion does not seem to affect authentication times for self-selected pictures*, while assigning random pictures to individuals (unsurprisingly) prolongs the time needed to authenticate.

5. FUTURE WORK

Use Your Illusion poses a number of interesting questions, which warrant further investigation.

Recall that our scheme does not mandate a given image processing filter: our choice of an oil painting filter has been driven mostly by heuristic considerations.

Further experimentation is also needed to better evidence the resilience of the scheme to some types of attacks; the low fidelity test we conducted to determine optimal parameter selection, while highly encouraging, needs to be expanded to provide stronger statistical evidence that attackers are not easily able to "revert" a distorted image back to its original meaning.

While investigating how best to tune our filter during the course of our prototype design and implementation, we discovered that finding an optimal parameter set point for our lossy filter depends on the picture to be transformed. For instance, holding filtering

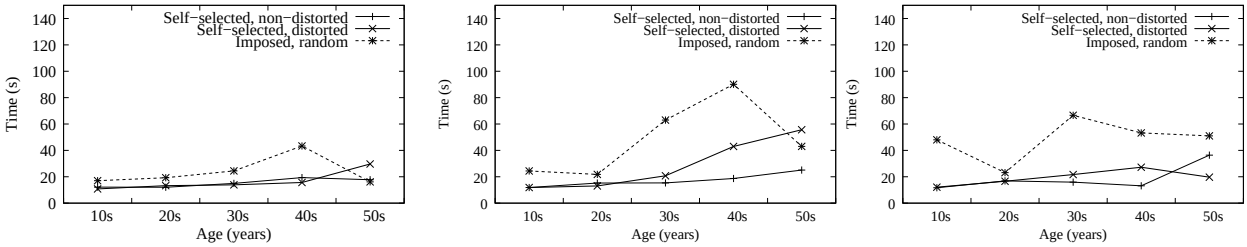

(a) First day (b) Two days later (c) One week later

Figure 7: Median authentication times vs. age (lower is better).

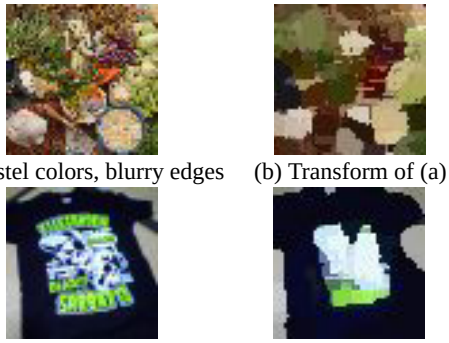

(a) Pastel colors, blurry edges (b) Transform of (a)

(c) Bright colors, sharp edges (d) Transform of (c)

Figure 8: Comparison of the filtering effects between different original images. Filtering parameters are held constant over both transformations. The transform (b) of an image (a) with pastel colors and slightly blurry shapes, conveys little information about the original picture. On the other hand, the transform (d) of an image (c) with bright colors and sharp edges, maintains most of the information of the original image.

parameters constants, distorting an image with bright colors and clear shapes preserves more information about the original image than distorting an image with blurrier shapes and less contrast.

Figure 8 shows two original-distorted image pairs, obtained using identical filtering parameters for both transforms. In Figure 8(a), the original image consists of pastel colors, and slightly blurry shapes. The resulting distorted image, Figure 8(b) does not convey much information from the original picture. Conversely, distorting an original image with bright colors and sharp edges, as in Figure 8(c) results in a picture very similar to the original image, as shown by Figure 8(d).

One can argue that neither Figure 8(a) or Figure 8(c) is an adequate picture for Use Your Illusion: (a) results in a distorted image that is not memorable, while (c) results in a distorted image that reveals too much information about the original picture. In the case of (a), a user can notice the problem and discard the picture during the training phase, so that we do not expect this problem to be a significant issue. The case illustrated in (c) is a little thornier in that it may be difficult to require users to voluntarily discard a portfolio image that is easy to remember. A related problem occurs when, after distortion, a picture may look similar to one of the chosen decoy.

Thus, it would be desirable to have a technical countermeasure to detect and adjust to such corner cases. For instance, if we had a metric that allows us to objectively measure the distortion level of a given picture, the filtering parameters could be adjusted for each original picture to obtain the maximum distortion level that still results in a memorable distorted picture. Finding such an objective "distortion metric" is an open problem.

Optimal configuration does not only concern filter parameters, but also authentication parameters. The near-perfect recall rate after four weeks tends to indicate that the choice of $p = 3$ portfolio pictures, $n = 27$ challenge set pictures may be a bit too conservative. One could want to increase either p or n to strengthen the resilience of the authentication scheme against brute force attacks. Further study is necessary to better characterize the relationship between n, p, and success rates.

Finally, while Use Your Illusion could be used for each single authentication instance, the 18 seconds taken on average to authenticate may be too high in some specific scenarios. Using mobile phones as railway passes, as can be seen in Japan, probably requires to authenticate in less than a second. To address such cases, one could envision a hierarchical authentication scheme, where Use Your Illusion is at the top of the hierarchy, and is used for all essential services. For services where, from the user's perspective, strong security is less of a stringent requirement than having a fast authentication process, one could use a simple PIN (that could only be changed after successfully authenticating in Use Your Illusion), or even no authentication at all, provided a successful authentication was performed within a reasonable timeframe.

6. CONCLUSION

In this paper, we propose a novel mechanism for user authentication that is secure and usable regardless of the size of the device on which it is implemented. Our system relies on the human ability to recognize a degraded version of a previously seen image. We illustrate how distorted images can be used to maintain the usability of graphical password schemes while making them more resilient to social engineering or guessing attacks.

We designed, implemented and tested a prototype authentication of Use Your Illusion for small, portable devices. A Web-based version of our prototype is available at `http://arima.okoze.net/illusion/`.

Our usability study provides evidence that users are extremely skilled at recognizing degraded versions of self-chosen images, even after long time periods of one month. Legitimate users who have been exposed to the original image can easily mentally revert a highly lossy, one-way transform on the image, even when it is not mathematically reversible and conveys limited information. Furthermore, many of our participants indicated that the authentication process is more enjoyable, and "game-like," compared to current alternatives.

While this paper focuses on addressing the critical need for strong user authentication on small portable devices, the results obtained with Use Your Illusion encourage us to consider a much wider

range of applications, such as ATMs and computing environments with large displays. We have, in fact, commenced researching possible applications of Use Your Illusion to observation-resilient authentication schemes [22].

More generally, we hope our paper can spur research in new methods to improve security systems by relying on human skills and strengths, rather than addressing human weaknesses.

7. REFERENCES

[1] Flickr. `http://www.flickr.com`.

[2] Phoney finance. The Economist. October 26, 2006. `http://www.economist.com/finance/displaystory.cfm?story_id=8089667`.

[3] R. Anderson. Why cryptosystems fail. In *Proc. ACM CCS*, pages 215–227, Nov. 1993.

[4] G. Blonder. United states patent, 1996. United States Patent 5559961.

[5] G. H. Bower, M. B. Karlin, and A. Dueck. Comprehension and memory for pictures. *Memory and Cognition*, 2:216–220, 1975.

[6] S. Brostoff and M. Sasse. Are passfaces more usable than passwords? A field trial investigation. In *Proceedings of HCI 2000*, pages 405–424, Sept. 2000.

[7] M. Burton, S. Wilson, M. Cowan, and V. Bruce. Face recognition in poor quality video: Evidence from security surveillance. *Psychological Science*, 10:243–248, 1999.

[8] R. Dhamija and A. Perrig. Déjà vu: A user study, using images for authentication. In *Proc. 9th USENIX Security Symp.*, Aug. 2000.

[9] R. Dhamija and J. D. Tygar. The battle against phishing: Dynamic security skins. In *Proc. 1st Symp. on Usable Privacy and Security*, 2005.

[10] A. Goldstein and J. E. Chance. Visual recognition memory for complex configurations. *Perception and Psychophysics*, 9:237–241, 1970.

[11] P. Golle and D. Wagner. Cryptanalysis of a cognitive authentication scheme. In *Proc. of the 2007 IEEE Symposium on Security and Privacy*, 2007.

[12] R. L. Gregory. *The Intelligent Eye*. 1970.

[13] A. Harada, T. Isarida, T. Mizuno, and M. Nishigaki. A user authentication system using schema of visual memory. In *Proc. BioADIT'06*, pages 338–345, Jan. 2006.

[14] Z. Henderson, V. Bruce, and M. Burton. Matching the faces of robbers captured on video. *Applied Cognitive Psychology*, 15:445–464, 2001.

[15] G. J. Holzmann. *Beyond Photography: The Digital Darkroom*. Prentice Hall, June 1988.

[16] I. Jermyn, A. Mayer, F. M. M. Reiter, and A. Rubin. The design and analysis of graphical passwords. In *Proc. 8th USENIX Security Symp.*, Aug. 1999.

[17] H. Kinjo and J. G. Snodgrass. Does the generation effect occur for pictures? *Amer. J. of Psych.*, 6:156–163, 2000.

[18] T. Matsumoto, H. Matsumoto, K. Yamada, and S. Hoshino. Impact of artificial gummy fingers on fingerprint systems. In *Proc. SPIE: Optical Security and Counterfeit Deterrence Techniques IV*, volume 4677, pages 275–289, Jan. 2002.

[19] W. Moncur and G. Leplâtre. Pictures at the ATM: exploring the usability of multiple graphical passwords. In *Proc. ACM CHI*, pages 887–894, Apr. 2007.

[20] F. Monrose, D. Davis, and M. Reiter. On user choice to graphical password schemes. In *Proc. of the 13th USENIX Security Symp.*, pages 151–164, San Diego, CA, Aug. 2004.

[21] Real User Corporation. The science behind Passfaces, 2001. `http://www.realusers.com`.

[22] H. Sasamoto, N. Christin, and E. Hayashi. Undercover: Authentication usable in front of prying eyes. In *Proceedings of 2008 ACM Symposium on Computer-Human Interaction (CHI'08)*, Florence, Italy, Apr. 2008. To appear.

[23] R. Shepard. Recognition memory for words, sentences and pictures. *J. Verbal Learning and Verbal Behavior*, 113(1):95–121, 1967.

[24] Sony Corporation. Overview of FeliCa. `http://www.sony.net/Products/felica/abt/dvs.html`.

[25] L. Standing, J. Conezio, and R. N. Haber. Perception and memory for pictures: single trial learning of 2,500 visual stimuli. *Psychonomic Sci.*, 19(2):73–74, 1970.

[26] A. Stubblefield and D. Simon. Inkblot authentication. Technical Report MSR-TR-2004-85, Aug. 2004.

[27] J. Thorpe and P. van Oorschot. Graphical dictionaries and the memorable space of graphical passwords. In *Proc. 13th USENIX Security Symp.*, Aug. 2004.

[28] J. Thorpe and P. van Oorschot. Towards secure design choices for implementing graphical passwords. In *Proc. 20th ACSAC*, Dec. 2004.

[29] J. Thorpe and P. van Oorschot. Human-seeded attacks and exploiting hot-spots in graphical passwords. In *Proc. 16th USENIX Security Symp.*, Aug. 2007.

[30] D. Weinshall. Cognitive authentication schemes safe against spyware. In *Proc. IEEE Symp. Sec. and Privacy*, May 2006.

[31] S. Wiedenbeck, J. Waters, J.-C. Birget, A. Brodskiy, and N. Memon. Authentication using graphical passwords: Basic results. In *HCI International*, July 2005.

[32] S. Wiedenbeck, J. Waters, J.-C. Birget, A. Brodskiy, and N. Memon. Authentication using graphical passwords: effects of tolerance and image choice. In *Proc. of the 1st Symp. Usable Privacy and Security*, pages 1–12, July 2005.