

Remotely Telling Humans and Computers Apart: An Unsolved Problem

Carlos Javier Hernandez-Castro and Arturo Ribagorda

Security Group, Department of Computer Science, Carlos III University,
28911 Leganes, Madrid, Spain
{chernand,arturo}@inf.uc3m.es

Abstract. The ability to tell humans and computers apart is imperative to protect many services from misuse and abuse. For this purpose, tests called CAPTCHAs¹ or HIPs² have been designed and put into production. Recent history shows that most (if not all) can be broken given enough time and commercial interest: CAPTCHA design seems to be a much more difficult problem than previously thought. The assumption that difficult-AI problems can be easily converted into valid CAPTCHAs is misleading. There are also some extrinsic problems that do not help, especially the big number of in-house designs that are put into production without any prior public critique. In this paper we present a state-of-the-art survey of current HIPs, including proposals that are now into production. We classify them regarding their basic design ideas. We discuss current attacks as well as future attack paths, and we also present common errors in design, and how many implementation flaws can transform a not necessarily bad idea into a weak CAPTCHA. We present examples of these flaws, using specific well-known CAPTCHAs. In a more theoretical way, we discuss the threat model: confronted risks and countermeasures. Finally, we introduce and discuss some desirable properties that new HIPs should have, concluding with some proposals for future work, including methodologies for design, implementation and security assessment.

Keywords: HIP, CAPTCHA, design, implementation, flaw, methodologies, security assessment.

1 Introduction

Herein we discuss some aspects of current CAPTCHA design and implementation. We are interested in both CAPTCHAs in use at major web-sites, and other commercial and/or academic proposals, some of which are still not broadly deployed, but present interesting aspects. After that, we will show the publicly-known state-of-the-art in attacks. We will introduce some examples to discuss common pitfalls in their design and implementation. We will also discuss the

¹ Completely Automated Public Turing test to tell Computers and Humans Apart.

² Human Interactive Proof.

different threat models they assume. Then we will discuss CAPTCHAs in a theoretical way. We will try to address why they have failed so frequently - put it other way, why they have remained unbroken for so brief time. We will discuss the reasons that make their design a much harder problem than previously thought. Also, there are more ambitious proposals that have tried to cover different types of attacks, including human farming attacks. We will conclude recommending some good properties that new designs should meet; properties that can be used as a starting point for a design methodology and for their security assessment.

1.1 Motivation: Why to Study and Break CAPTCHAs

Today we do not fully understand if the problems that form the basis of the current CAPTCHAs designs are hard enough. Even worse, we do not have a way to tell if the precise method employed by a particular design uses the full strength of the underlying, supposedly hard, AI problem. We do not have a clear methodology to check if a design and its implementation are flaw-free. It is then interesting to try to break current CAPTCHAs and find pitfalls in their design to make the state-of-the-art advance and get to a point where better known and tested assumptions give rise to more secure CAPTCHAs.

1.2 CAPTCHAs' Evolution

CAPTCHAs are somewhat recent. Moni Naor was the first to propose theoretical methods of remotely telling apart computers from humans [19]. The first known use of a CAPTCHA on-line is credited to the AltaVista web search-engine in 1997 [22]. Andrei Broder developed a filter, randomly generating an image of printed text, and transforming it so that machine character recognition systems (OCRs) could not read it, but humans still could. In 2000, Udi Manber of Yahoo! described their "chat room problem" to researchers at Carnegie-Mellon University: 'bots' (automatic computer scripts) were joining on-line chat rooms, and pointing their users to advertising sites. Professors Manuel Blum, Luis A. von Ahn, and John Langford, from Carnegie-Mellon University, described some desirable properties that a test to exclude bots should have. These properties are mostly theoretical, and do not deal with specific design or implementation issues (i.e., nothing is said about whether challenge generation is better in the client or server side). The Carnegie-Mellon University team developed a 'hard' GIMPY, which picked English words at random, and rendered them as images of printed text, under a wide variety of shape deformations and image distortions[17], with an easier version put to work at Yahoo!.

Text-based CAPTCHAs. In the following years, OCR/text-based CAPTCHAs became increasingly popular. For awhile, attackers preferred those sites not protected by CAPTCHAs as still there were many services that did not use them, so advertising/spam/phishing/voting abuse etc. were still possible by avoiding CAPTCHAs. Even after graphic distortion and degradation, some approaches

have been able to "read" them automatically around 92% of the time [18], specially when it is possible to divide the rendered text image into its constituent letters or characters, a procedure called "segmentation". After some more attacks on OCR/text-based CAPTCHAs were made public, the outcome was not the creation of new paradigms, but the redesign of OCR/text-based proposals in a way to try to prevent those attacks [25]. Newer attacks appeared [32].



Fig. 1. Some current OCR/text based CAPTCHAs

Although now the commercial and academic proposals offer a wide variety of types to choose from, most deployed CAPTCHA belong to the OCR/text-based type of challenge. Among them, there are some interesting ones, as reCAPTCHA [13], which is an interesting mix of a normal OCR/text and database based CAPTCHA: it is based on extracts of texts that had not been correctly recognized by OCR software. That reason is why the test presents two words: one comes from the failed-to-OCR set, the other is one which reading is known: either an automatically generated and distorted one, or a failed-to-OCR one already read by humans. In this sense, reCAPTCHA is an *incremental CAPTCHA*. reCAPTCHA also includes an audio version, further explained in the audio section of this paper. The current Megaupload³ is interesting because uses a strong overlapping model (trying to make segmentation much harder). This approach is sound, as computers are even better than humans at recognizing isolated characters [33], so the most difficult problem for a machine is segmentation. BaffleText uses a somewhat related approach [28], using "occlusion or interference by random shapes", against which, they state, "image restoration cannot replace missing parts without prior knowledge of the occluding shapes", humans being "remarkably good at recognizing an entire shape or picture in spite of incomplete, sparse, or fragmentary information". Another atypical text-based scheme is the one proposed to protect the "Quantum Random Bit Generator Service"⁵.

³ <http://www.megaupload.com/>. Their previous CAPTCHA has reportedly been broken during January, 2009⁴. <http://userscripts.org/scripts/show/38736>

⁴ Megaupload auto-fill captcha script.

⁵ <http://random.irb.hr/>, Center for Informatics and Computing of the Ruđer Bošković Institute (Zagreb, Croatia).

They ask to solve or compute a not-so-simple math expression to prove that you are human (and good enough in Math). To prevent automatic OCR reading, they do not rely on transformations, but on the low quality of the depicted text, using small fonts and antialiasing.

Another interesting and atypical OCR/text CAPTCHA is Teabag⁶, currently in production, which depicts the text in 3D, also transforming in 3D space. Other proposals based on the OCR/text reading challenge are the commercial Captcha2⁷, providing an interesting mix of an OCR/text-based and an image-based: the user must read the character presented after the message "click on the letter", and select the same letter by clicking close to it in the box displayed below the caption. There are also some OCR/text CAPTCHAs that rely on a broad range of parameters for text depicting, thus looking as if they were composite versions of simple ones, as Jcaptcha⁸, deployed in many sites⁹, Securimage¹⁰, an open-source free PHP CAPTCHA script, or BotDetect¹¹, a commercial proposal. There is also a proposal [23,10] that tries to mimic the human handwriting in a synthetic form for creating challenges that are more difficult for OCRs.

Picture-based CAPTCHAs. Many text-based CAPTCHAs were recently broken [18,30,31,35,36,32,33], so there is increasing concern about their overall strength and their accessibility for humans (as they became harder, for countering these attacks). Possibly influenced by the idea that the general vision problem seems to be a harder problem than character recognition, more designs lately focus on using pictures instead. Those CAPTCHAs do not really rely on a "general vision recognition problem", but in a downsized version of it, typically consisting in labeling images, finding their center, etc. Chew and Tygar[20] were the first to use a set of labeled images to produce CAPTCHA challenges, using Google Image Search. Google relates a picture to its description, and to its surroundings in the document that contains it. Ahn and Dabbish [21] proposed a new way to label images, embedding the task as a game, thus creating the PIX CAPTCHA database. HotCaptcha.com proposed a new way to use a large-scale human labeled database provided by the HotOrNot.com web-site¹², a site that invites users to post photos of themselves and rate others' in a sex appeal numerical scale. Oli Warner came with the idea of using photos of kittens to tell computers and humans apart [41]. KittenAuth features nine pictures of cute little animals, only three of which are feline. The database of pictures is small enough (< 100) to manually classify them, and this limitation seems troublesome even if we apply new methods involving image distortion. Another

⁶ OCR Research Team, Kharkov (Ukraine), at <http://ocr-research.org.ua/> and <http://barafanra.com/game-register.php>

⁷ <http://captcha2.com/>

⁸ <http://jcaptcha.sourceforge.net/>

⁹ Including a major airline web-site, (<http://www.iberia.com>)

¹⁰ <http://www.phpcaptcha.org/captcha-gallery/>

¹¹ <http://captcha.biz/captcha.html>

¹² No longer active, as of January 2009.

proposal, known as the HumanAuth CAPTCHA (fig. 2(a)), asks the user to distinguish between pictures depicting Nature or an artificial thing (like a clock, or a vase, for instance). In addition, the CAPTCHA has also an alternative version for the visually challenged, as they can switch to text-based descriptions of the pictures.



Fig. 2. Some current picture based CAPTCHAs

The main problem with HumanAuth is the picture database size, which is too small, even though the authors include an algorithm to mix it with a logo for watermarking and avoiding hash-function-based image classification (that will not affect text descriptions). ASIRRA [16] (fig. 2(b)) uses a similar approach, but using a database of more than 3 million photos from Petfinder.com. ASIRRA asks for identifying images under two categories, cats and dogs. VidoopSECURE¹³ is a company that provides security solutions, and gives access free of charge to their VidoopCAPTCHA, composed of 12 images, each one labeled with a letter, asking the user to enter the letters in those images that refer to three given categories. IMAGINATION [26] is a CAPTCHA divided in two stages: the system shows the user an image composed of 8 images tiled and asks the user to click "near" the geometric center of one of the images, then that image is transformed and presented again, and this time the user has to choose the word (from a set) most related to the image being showed. 3D CAPTCHA¹⁴ does not exactly fall into the picture database-based category. Instead, it uses pre-generated 3D models that are rendered in a randomized scene in 2D. For each part of each model, there is a textual description. The computer selects random objects from the set of available ones and sets them up in a scene, rendered also from a random point of view and with random light sources, and finally attaches letters to the different parts of the different models, asking the user to enter them in the sequence specified by a list of given parts of models (figure 2(c)). The idea of generating 2D CAPTCHAs from 3D models is not new, as [27] proposes a simpler scheme of 2D challenge generation from 3D object models.

¹³ <http://drupal.org/project/vidoopcaptcha>

¹⁴ <http://spamfizzle.com/CAPTCHA.aspx>

Other visual types. We propose in [2] some video CAPTCHAs schemes based in real videos, analyzed and then transformed. The user has to answer whether them had been transformed or not, and the way. The videos need to have a former analysis to avoid easy-to-solve specimens. With some extra care, it is also possible to use on-line video repositories as the input source for some of our CAPTCHA schemes. Using animation, [3] proposes a general way of preventing relay and human-solving farms attacks: they include a time component in their CAPTCHAs, so that the same challenge has a time dependant answer, and could not be easily display exactly as-is in a remote computer. A long chimera of many proposals is to be almost transparent, not noticeable by the user. One surreptitious CAPTCHAs that does not alter navigation is [29], which consists of replacing many of the links of the web-site for images with captions to tell where to click. Metadata¹⁵ is supposed to be another type of this scheme, but no information is public to date.

Audio CAPTCHAs. Visual CAPTCHAs put in production need a counterpart to provide access to the visually impaired[15]. Audio CAPTCHAs can be hard to solve for humans [4]. The GMail and Microsoft Live! audio CAPTCHAs are based on voice samples of people reading decimal digits (0 to 9), with some background noise. They have already been broken with simple algorithms [40,39]. A different type of audio CAPTCHA is the new¹⁶ one provided by reCAPTCHA. Using the same idea on which their visual CAPTCHA is based, they provide sound recordings (from radio broadcasts) that have been partially recognized by their software. The user has to correctly recognize at least the part that the software audio recognizes. Note that this scheme is weaker than the visual one, as here the test is either a previously human-solved test, or one to be solved (and not a mix of both). If it is one to be solved, then the current best software for audio recognition will be able to break the CAPTCHA.

2 Attacks

To date, publicly known attacks on CAPTCHAs are mostly the work of amateur programmers or researchers. Some schemes have been known to change without a public notice of its breakdown, so there is a chance that some attacks have been and will remain unpublished.

2.1 Text Recognition

OCR/text-based CAPTCHAs have been the first type to be deployed, and are still the most widely used type. It is logical that they have attracted so far most of the attention. The first published attack against an OCR/text CAPTCHA is devoted to Mori and Malik [18]. They use shape-matching (in a two steps

¹⁵ (http://www.forbes.com/2008/11/25/captcha-pramana-bots-tech-identity08-cx_ag_1125captcha.html)

¹⁶ Introduced in 2009 to replace their old, digit based, model.

approach) to find candidate letters, and then filter the candidate words. In another approach they try to find the whole words directly. Another attack [30], based on distortion estimation, succeeds 99% of the time. For simpler OCR/text CAPTCHAs, it is possible to use simpler algorithms (as pixel counting [31], letter derotation [35], etc.). There are open source algorithms able to decode the most simple OCR/text based [36]. More recent attacks [32] allow to break most OCR/text based CAPTCHAS (including Yahoo, GMail, etc.). It has been shown [33] that typical text transforms do not increase security, as computers recognize isolated characters better than humans.

These examples are just current attack possibilities, but there are many others to explore. During our preliminary research, we have found that an attack which focusses on regular shape detection (using adapted Hough transforms for segments, ellipses, etc.) is of use against Baffletext, and also Megaupload. Captcha2's background is not a major problem; it follows some clear color-shift patterns. Also, the transformations used to depict the letters are fixed, and uses the same set of transformed letters for depicting the character that has to be clicked on in the box. Thus, a simple pattern matching algorithm would be able to solve it. Teabag has many possible analysis, notably those based on the non-uniformity of the challenge (see later in 2.7).

2.2 Audio Recognition

Typical audio CAPTCHAs are very basic, and excluding very few exceptions, most current implementations can be considered defeated. Current proposals are so simple that no complex analysis is required, just a slight filtering and pattern matching [40,39]. More sophisticated attacks have been studied [12], rendering even better success ratios. The new version (as of 2009) of the reCAPTCHA audio CAPTCHA is one of the few that remains unsolved, but as we stated in the previous section, it is not as strong as the visual one, and we are exploring the possibility that unsolved challenges can be used to bypass it.

2.3 Side-Channel Attacks

These type of attacks are based on deviations from randomness that allow for a correlation among the challenges and their answers. In this section we present a detailed example (against HumanAuth), and we briefly describe other attack against ASIRRA [5]. The ASIRRA Public Corpus, that the developers of ASIRRA have created to help researchers, is composed of around 25.000 images classified as cats or dogs, a half in each class. We analyzed all of the files using the ENT¹⁷ tool, producing a formatted output (in ARFF format, so to be used with Weka [6]). These were later processed by a classifier which was able to distinguish cats and dogs pictures with a nearly 60% accuracy, without using any kind of image recognition technique. The simplest decision tree [7], based on only the **size** of a jpeg file, is able of distinguishing between cats and dogs with

¹⁷ <http://www.fourmilab.ch/random/>

an accuracy over 58%, significantly better than random. Following an identical approach, we used the ENT tool with the images in the HumanAuth source code, producing an ARFF file. The best classifier (in this case RandomForest) was able to show an accuracy rate of 77.8761%, significantly better than the $\frac{68}{68+45} = 60.177\%$ that a trivial classifier (that always predicts the larger class) will do. To prevent easy image library indexing, the authors of HumanAuth decided to merge a PNG image with the random JPG image taken from the library: they put the watermark PNG in a random position into the JPG canvas and merge both using a level of transparency. This suggests that the initial small set of images, when used with the scheme proposed by the HumanAuth authors of merging with a watermark, may not be of use against this type of attack, even though might be enough to prevent hash-function indexing. Choosing a different watermark that alters more the original image could be beneficial, but that would be also at the expense of human recognition.

2.4 Feature-Based Attacks

In Philippe Golle's work [8], he presents the strongest attack against ASIRRA as of April-2009. This attack is based in image processing, as it divides the photographs into $N \times N$ cells of color and texture (gray-scale) information, and use that to feed two support-vector machine (SVM) classifiers that, when used together, are capable of classifying with around a 83% accuracy, thus allowing them to solve the 12-photos challenge with a 10.3% probability. They found that color-presence features are more accurate for classifying cats and dogs than color-histogram features. With this scheme, the authors reach a 77.1% accurate classification rate. With texture processing, they reach an 80.4% successful classification. Combining both methods, they reach an 82.7% accurate classification rate.

2.5 Attacks against Database-Based CAPTCHAs

When a CAPTCHA is based on a database of knowledge (i.e., labeled pictures), and especially if that database is public, there are some possible attacks against its database that could thwart its reliability. Basically, those are:

1. Database poisoning attacks: if the database is public, and not protected, we can upload information in a way that assures that when later confronted with challenges created using that uploaded information we will solve those challenges.
2. Database indexing attacks: if the database is small enough and/or the bandwidth is high enough, we can download (maybe partially) the database contents and get enough information to solve or greatly simplify the CAPTCHA.

2.6 Implementation Flaws

Some CAPTCHA protection systems can be bypassed without using OCR techniques at all, simply by reusing the session ID of a known CAPTCHA image

[44]. That is the result of a very bad implementation, but it was not uncommon some years ago. A common error, even today¹⁸, is that some codify the answer to the challenges in the URL address or in the value of a form field. Knowing this, we can request as many challenges as we like with the same solution. We can calculate medium values of those challenges, and thus launch a mean attack [45]. Other implementations use a hash (such as an MD5 hash) of the solution as a key passed to the client. If the number of possible answers is finite, or their distribution is not uniform, solutions hashes could be learned in a proportion enough to solve it. Another straightforward implementation flaw consists on using only a small fixed pool of challenges. For example, HumanAuth uses less than one hundred images, and even masking them with logos, we cannot prevent them from being indexed or characterized [5]. As another example, QRBGs challenges are not created on-demand, but repeated [43]. If the answer to the challenges could be any (i.e., floats, in a broad range) and it followed an uniform distribution, repetition of challenges would not be so deleterious, but in nearly all other cases is fatal. Some additional design flaws in HumanAuth (commented later) imply there is a big chance than the challenge answer is 0, or, alternatively, a small integer. This makes possible to use another, very successful, kind of attack: if answering 0 fails, then we will answer with a succession of integers that will run through N , starting with the smallest absolute values. As a last and typical implementation flaw, we have to note that some algorithms use a quite unwise way of communicating with the CAPTCHA server, that is easy to attack [14].

2.7 Design Flaws

In this section we introduce some common design flaws that have rendered some known CAPTCHAs much less secure than intended.

Biased answer distribution. One obvious mistake, but quite common to find, is selecting a clearly non-uniformly distributed subset of the possible answers as correct answers. One example is QRBGs (MathCAPTCHA), whose designers use one-digit figures in all their arithmetic operations. That makes it likely that the answers will be small integers. Additionally, they use derivatives of trigonometric functions when x has a typical value of $\frac{\pi}{2}$ or π and they do not expect floats as answers (possibly to avoid precision problems), so low-integers and specially zero are very good candidates for successful blind answers (up to 93% success rate, depending on problem subtype) [11].

Another example is phase one of IMAGINATION, in which the user has to click on the "center" of any of the images. Clearly this center distribution is not a uniform distribution. We find other examples in OCR/text CAPTCHAs that do not use the full set of possible characters. For example, the Megupload CAPTCHA avoids values O, I, J, 0 to prevent user confusion. Worse, it always uses the scheme of three letters followed by a digit. That makes it more

¹⁸ See, for example, the CAPTCHAs currently used by Fotolog, Uol, Conduit, MetroFlog and others.

user friendly, but also much weaker. Teabag uses only three characters, and are strongly not uniformly distributed [34], possibly to avoid characters that are difficult to distinguish in the 3D projections they use. In a sample of 100 challenges, characters 'S', 'Z', '3', 'P', 'b', 'w', 'M', 't' and 'd' appeared more than 3% of the time (peak: 4.3%), while a large set of other 34 characters appeared none, including '1', '0' (perhaps to avoid coincidence with 'I' and 'O'), etc. The chi-square value for this distribution with 61 degrees of freedom gives us 262.08, corresponding to a two-tailed p value of less than 0.0001. As the challenge answer does not differentiate among upper and lowercase¹⁹, the situation is even worse. In this case, the most common character is letter 'm' (6%), and there are 4 characters that appear more than a 3% of the times, while other 18 do not appear at all. That means that just blindly answering the string 'mmm' will be a good guess. Depending on the scheme, uniformity in the answer distribution can be considered a theoretical aim, but we should always know how far we are from it, as being too far can render our CAPTCHA useless.

Biased challenge distribution. Any bias from randomness in the characteristics of the challenges can allow for challenge analysis, which may lead to side-channel attacks (if they allow correlation with the answers), or simpler-than-intended challenge categorization and/or analysis. One example is Teabag. In it, the frontal borders of the characters can be selected by area size, because this distribution is far from uniform: so size has a meaning. Additionally, thanks to the non-uniformity of the image (for example, the areas at the borders), it is easy to tell the angles of the image: flood-filling the corners and studying the lines that they create, we can de-rotate and de-shear the image to a normal 2D perspective. There is also a correlation among pixels that allow for (probabilistic) back border detection, using very simple algorithms, like growing the background areas by pixel continuity. In a fair percentage of the challenges generated, the non-character part of the image can be completely or almost completely removed [34].




Fig. 3. Teabag, original and first phase of process

¹⁹ Teabag page does not specify this detail, but the implementation at <http://barafanra.com/game-register.php> does not make any difference.

Another example is the Megaupload CAPTCHA: the letters and the digit are always printed using the same font type, similar if not equal to Antique Olive (as identified by Identifont²⁰). They are rotated, but at a precise angle, either clockwise or counter-clockwise, being the first letter rotated clockwise, the second counter-clockwise, etc. Additionally, it prevents more than two characters from overlapping (this and the former allow for the creation of overlapping maps with all combinations, although they are not necessary to break the CAPTCHA) [9].

Correlation between Challenge and Answer. This is a worse scenario than the ones presented before: here, the challenge somehow provides (non intending, that is "leaked" or "side-channel") information on the answer. When the answer is binary, then the problem is similar to the previous biased challenge distribution problem. The already explained side-channel attack against HumanAuth and ASIRRA clearly exemplifies the dangers of this scenario.

Limiting the size of the set of correct answers. Is a straightforward weak point that takes place much more commonly than what could be thought. Clear examples are MegaUpload, asking always for 3 letters and a digit (in that order) and not distinguishing case (also avoids some characters (I, J, L, O, 0), possibly to avoid confusions), or Teabag, asking for only 3 characters (and not distinguishing case, either). Also, many OCR/text base CAPTCHAs limit the number of characters of the answer.

Categorization of the answer. There is not always the need to make it easy for a program to know if the guessed answer for the CAPTCHA was right or wrong. Interaction with the remote system (typically, web-site) can go on as usual, being the actual data results different, etc. If possible, it should be avoided the immediate knowledge of whether an answer to the challenge is correct or wrong, or any other way to know if it is close to being a correct answer. We can communicate this information using some intermediary communication mechanism (like email accounts, which will also need to be controlled not to use the same more than n times), or present that information to the user in a way difficult to distinguish automatically.

User dependance. Make the CAPTCHA, in any way, dependant on the challenger, is in general a very bad idea, specially worse if this dependance is, or could be, known. One example of this problem is ASIRRA that, to further increase the possibilities of adoptions of the pets showed in the CAPTCHA, will show the pets in Petfinder that are close to the challenger (using IP geolocation). This weakness is egregious, as it renders easier to many types of attacks, including database poisoning (the number of poisoned records to bypass the CAPTCHA is significantly lower), database indexing, etc.

²⁰ <http://www.identifont.com>

2.8 Human Solving

CAPTCHAs are designed to be solved by humans, but there is an ongoing market of CAPTCHA solving services [38]²¹ (typically located in an area where relatively cheap labor can be found) and relay attacks, that present the CAPTCHA challenge to a human user that receives some benefit from solving it [37]. This is part of the CAPTCHA economics. There is still place for larger abnormalities, i.e., one client willing to break a CAPTCHA can actually sell a CAPTCHA service to other sites, and use them as relay agents to solve it. We have already mentioned a CAPTCHA that tries to prevent human relay attacks.

3 The General CAPTCHA Design Problem

In this section we try to build upon the previous sections that showed different kinds of CAPTCHAs; the AI problems they try to be based on, some of their design problems and implementation flaws; and the currently known attacks against them; to discuss the general problem of CAPTCHA design, and finally, their current strength and limits.

3.1 Threat Model

The original purpose of a CAPTCHA was determining a human from a machine remotely. If that determination is possible, we can protect services from automatic/massive requests. With the evolution of the attacks on some web services, it has been more evident that the ability to tell a human from a machine may not be enough. Is a human in a CAPTCHA solving farm something we should avoid? Should a human surfing through another web site, or presented with another program GUI, not be eligible for solving our CAPTCHAs? Is a human assisted computer program still an automatic attack? There are different types of cryptanalytic techniques depending on the threat model we choose: chosen plaintext, known plaintext, known cyphertext, related key, side-channel, etc. Analogously, there are different threats that a CAPTCHA can choose or not to prevent, from the most basic to the most comprehensive:

1. be able to distinguish humans from computers, by measuring an "human" quality, ability or behavior.
2. be able to prevent 'magnifying/human-assisted' attacks: distinguish humans from human-assisted algorithms, i.e., algorithms that with some human intervention can solve the CAPTCHA many times. This assisted algorithms have a big ratio of CAPTCHAs solved per human intervention²².

²¹ Some examples of human-solving services are (<http://captchaocr.com/ad.html>) and (<http://decaptcher.com>).

²² An example could be a variant of the Megaupload with multiple typefaces: solved with lines and shapes analysis (using the Hough transform) and human intervention (classifying the new schemes as corresponding letters).

3. be able to prevent relay attacks: distinguish humans presented with the CAPTCHA in the original CAPTCHA site, to those presented with the CAPTCHA in another site/interface²³ [3].
4. be able to prevent 'human farms' attacks: including methods to inhibit or make more difficult solving the CAPTCHA using 'farms' of solvers²⁴.

3.2 AI Hardness Not Transmitted

Most CAPTCHAs have been broken [18,8,5,9,11,14,42,40,39,32,31,30,34] due to one of these problems:

1. The "hard-AI" underlying problem they are based on is not the original one intended, but a much more specific and weaker one. For example, most OCR/text CAPTCHAs require to solve a problem far easier than the "general OCR" problem [32,9,34]. In this aspect, reCAPTCHA tries to approach the real problem more realistically.
2. Design and implementation flaws make them much more easy to bypass using a procedure that, analyzing the challenges in a non detailed manner, permits to guess an answer with a high percentage of success. Thus, these procedures can be called side-channel attacks, as they do not try to solve the original problem that the CAPTCHA designers intended attackers to solve, but one which is much easier [11,5]. That is the difference between the intended solving path and the actual solving path. The intended path is the one based on a hard problem (purportedly representing an AI-hard problem), the real one is the one based on any design and/or implementation flaw that avoids the harder way.

The difficulty of an AI-unsolved problem is thus not easily transmitted to a CAPTCHA design. One of the reasons is because we do not know how to categorize AI-hardness, so we do not know how to tell if a particular subset (generated by a CAPTCHA) is hard enough. Also, an AI-problem can be difficult to solve correctly more than, i.e., 15%, but we typically need a much smaller figure to be able to attack the CAPTCHA. Regarding the image recognition problem, used for many modern CAPTCHA designs, we face the same obstacles. "Image recognition/understanding" is not typically the same problem that the CAPTCHA designers finally implement. Human-solved database-based CAPTCHAs can approach the hardness of the original AI problem if we filter out easy specimens (which is what reCAPTCHA does), although they impose the limitations of a database-based CAPTCHA design, including the need of a constant input stream of unsolved challenges.

4 Related Problems

Designers face additional problems that have not to do with their resistance to attack. As many CAPTCHAs protect free services, or commercial services with

²³ Sites/programs that, upon solving the CAPTCHA, reward the user in some way.

²⁴ One possible way of doing this is with in-site "hidden" CAPTCHAs [29].

competence, CAPTCHAs have to be easy to solve and appealing. Ideally, if the service is worldwide, they have to be language-independent (i.e., reCAPTCHA is harder for non-English speakers, with a 97% vs. 92% success ratio). The problem of accessibility is also of greater concern, and few visual CAPTCHAs are fully already prepared for it (i.e., HumanAuth). If possible, they have to be also possible to operate through other channels (i.e., mobile phones). This requirements have made some researchers invent new forms of reusing known CAPTCHA schemes, as for example, converting text CAPTCHAs into clickable images [1]. An additional problem is the tendency to overlook the difficulty of their design, specially when designing new in-house developed CAPTCHAs, typically by groups of programmers that are not specialists, which designs are put right into production [11]²⁵.

5 Where to Go

Recapitulating on the previous sections, we present a brief corpus of conclusions, learned from the previously studied problems. We describe good properties that a CAPTCHA design should have, ways to theoretically study the assurance level of a design, and to empirically test for good properties regarding to randomness in different aspects of the CAPTCHA.

5.1 Good Properties CAPTCHAs Should Have

Here we present some good properties that new CAPTCHA designs should have:

1. Randomness, uniform distribution, in all parameters. For example, for a text CAPTCHA: variable number of characters, uniform numbers of pixels/areas/
lines/... with certain properties (color, group, size of group, etc.), different type faces used, size of images, etc.
2. No easier challenges: subtypes or alternatives should have the same strength (audio and visual CAPTCHAs).
3. Problem posed should be as broad, and as close as possible, to the AI problem that inspires the test.
4. Design should incorporate features to prevent relay attacks, and to detect automatic bypass easier.
5. Challenges should be independent and uniformly distributed (this excludes dependance from the user). Also the answers should be random, uniformly distributed. There should not be any statistical correlation among challenges and answers.
6. Make it difficult for a program to tell if its answer was or not correct.

As a good practice, we recommend, for any new CAPTCHA design, to be put into production in a test web-page, without other protections (to focus on the hardness of the CAPTCHA), for a period long enough to allow study.

²⁵ As an example, CAPTCHAs offered for Wordpress at (<http://wordpress.org/extend/plugins/tags/captcha>)

5.2 Assurance Level of CAPTCHAs

In this section we propose some basic methodologies to analyze the security offered by a CAPTCHA design. This analysis is theoretical, and studies the behavior of the CAPTCHA design under unfavorable assumptions. For knowing better the security offered by a CAPTCHA design, we propose to theoretically analyze the CAPTCHA under the following assumptions:

1. Answer repetition: if an attacker is able to collect a finite sample of challenges with the same answers, in any quantity, for any number of different answers she wants, confirm that she will not be able to create a better answer, for a challenge that pertains to one of those answers, than a random one (that is, that there is no better attack than trial and error).
2. Challenge repetition: if our CAPTCHA has only n different challenges, and we do not know their answer, there should be no better strategy to solve them than trial and error, which success ratio should be low. This should apply also when $n = 1$.
3. Non categorization: if our CAPTCHA is composed of different types of challenges, there should be no way to automatically distinguish them. There should also be no way of categorizing the difficulty of different challenges.

5.3 Security Assessment

In this section we propose a practical test that can be done over any CAPTCHA design and implementation, and that should be added to the ones aforementioned. This test measures the random distribution of answers to the challenges.

For this test, we create a large enough set of elements (T = test, A = answer) of tests with their correct answers. For the designers, creating this set should be straightforward. Then, using general randomness and statistical analysis tools[5], we search for non-uniformities in this distribution, that is:

1. Non-uniformities in the distribution of A (that would allow for a blind attack)
2. Correlations among T and A (that would allow for a side-channel attack)
3. Non-uniformities in the distribution of T (that could allow for type-of-challenge categorization, and challenge analysis)

These tests can be done for T as is (as a bit stream) or for some simple properties of T, which, depending on the type of challenge, can be color histogram, areas' sizes histogram, distances between similar areas, bit correlation with given vectors (i.e., trained by GA), maximum and minimum for a block of bytes, etc. This can be used as a very general analysis tool to realistically estimate the security parameters of any CAPTCHA proposal, and we believe it will be advisable to use it in the future before any similar systems are launched to have adequate, well-reasoned, and founded security parameters and realistic estimations. One of their main advantages is that it does not depend on the underlying format (image, sound, video, etc.) or problem, and that it could be useful for avoiding pitfalls such as the existence of some trivial and irrelevant parameter values

leaking too much information [5,11]. For the study of correlations and other predictive possibilities we, apart from the classical statistical tools, strongly recommend the use of Machine Learning algorithms such as those found on the free Weka [6] tool.

6 Conclusions and Further Work

It is possible to assume broader threat models and design CAPTCHAs that also prevent human farming, human relay attacks, and are able to detect when they are under computer attack, and when they have been bypassed. Hardness assumptions relying on the base AI problem underlying a particular CAPTCHA are often used as arguments, and given as facts. In reality, the AI problem proposed by the CAPTCHA is always a small subset, typically much easier to solve than the original. Worse, typically CAPTCHA designs have pitfalls that make them even easier to pass. We expect this work to be a good starting point: both for designers of new CAPTCHAs, to avoid some common design flaws that could render their ideas useless, and for the creation of additional methodologies for security assessment, and evaluations of the assurance level they provide.

References

1. Chow, R., Golle, P., Jakobsson, M., Wang, L., Wang, X.: Making CAPTCHAs Clickable. In: Proc. of HotMobile 2008 (2008)
2. Hernandez-Castro, C.J., Ribagorda, A.: VideoCAPTCHAs. In: Proceedings of the 5th International Conference on Security and Protection of Information, Brno (to be published, 2009)
3. Athanasopoulos, E., Antonatos, S.: Enhanced CAPTCHAs: Using Animation to Tell Humans and Computers Apart. In: IFIP International Federation for Information Processing (2006)
4. Bigham, J.P., Cavender, A.C.: Evaluating Existing Audio CAPTCHAs and an Interface Optimized for Non-Visual Use. In: CHI 2009 (2009)
5. Hernandez-Castro, C.J., Ribagorda, A., Saez, Y.: Side-channel attack on labeling CAPTCHAs (2009), <http://arxiv.org/abs/0908.1185>
6. Holmes, G., Donkin, A., Witten, I.H.: Weka: A machine learning workbench. In: Proceedings of the Second Australia and New Zealand Conference on Intelligent Information Systems, Brisbane, Australia (1994)
7. Quinlan, R.: Machine Learning. Morgan Kaufmann Pub., CA
8. Golle, P.: Machine Learning Attacks Against the Asirra CAPTCHA. In: ACM CCS 2008 (2008)
9. Hernandez-Castro, C.J., Ribagorda, A.: Preliminary analysis on the Megaupload CAPTCHA (Submitted for publication)
10. Achint, T., Venu, G.: Generation and Performance Evaluation of Synthetic Handwritten CAPTCHAs. In: ICFHR 2008 (2008)
11. Hernandez-Castro, C.J., Ribagorda, A.: Pitfalls in CAPTCHA design and implementation: the Math CAPTCHA, a case study (Submitted for publication)
12. Tam, J., Simsa, J., Hyde, S., Von Ahn, L.: Breaking Audio CAPTCHAs, http://www.captcha.net/Breaking_Audio_CAPTCHAs.pdf

13. von Ahn, L., Maurer, B., McMillen, C., Abraham, D., Blum, M.: reCAPTCHA: Human-Based Character Recognition via Web Security Measures. *Science Magazine* 321(5895), 1465–1468 (2008)
14. Caine, A., Hengartner, U.: The AI Hardness of CAPTCHAs does not imply Robust Network Security. In: *IFIP, Trust Management*, vol. 238, pp. 367–382
15. W3C Working Draft, Techniques for WCAG 2.0. G144: Ensuring Web Page contains another CAPTCHA serving same purpose using a different modality
16. Elson, J., Douceur, J.R., Howell, J., Saul, J.: Asirra: A CAPTCHA that Exploits Interest-Aligned Manual Image Categorization. In: *Proceedings of the 14th ACM CCS (2007)*
17. von Ahn, L., Blum, M., Hopper, N.J., Langford, J.: CAPTCHA: Using hard AI problems for security. In: *Biham, E. (ed.) EUROCRYPT 2003. LNCS*, vol. 2656, pp. 294–311. Springer, Heidelberg (2003)
18. Mori, G., Malik, J.: Recognizing objects in adversarial clutter: Breaking a visual CAPTCHA. In: *Proc. of CVPR 2003*, pp. 134–144. IEEE, Los Alamitos (2003)
19. Naor, M.: Verification of human in the loop or Identification via Turing Test, <http://www.wisdom.weizmann.ac.il/~naor/PAPERS/human.ps> (retrieved January 1, 2009)
20. Chew, M., Tygar, J.D.: Image recognition CAPTCHAs. In: *Proc. of ISC 2004*, pp. 268–279 (2004); A longer version as UC Berkeley Computer Science Division technical report UCB/CSD-04-1333
21. von Ahn, L., Dabbish, L.: Labeling Images with a Computer Game. In: *ACM Conference on Human Factors in Computing Systems, CHI 2004*, pp. 319–326 (2004)
22. US Patent no. 6,195,698, Method for selectively restricting access to computer systems, <http://www.freepatentsonline.com/6195698.html>
23. Rusu, A., Govindaraju, V.: Handwritten CAPTCHA: using the difference in the abilities of humans and machines in reading handwritten words. In: *Proceedings of the IWFHR-9, Tokyo, Japan*, pp. 226–231 (2004)
24. Stevanovic, R., et al.: Quantum Random Bit Generator Service for Monte Carlo and Other Stochastic Simulations. In: *Lirkov, I., Margenov, S., Waśniewski, J. (eds.) LSSC 2007. LNCS*, vol. 4818, pp. 508–515. Springer, Heidelberg (2008)
25. Baird, H.S., Riopka, T.: ScatterType: a Reading CAPTCHA Resistant to Segmentation Attack. In: *Proc. of the IS & T/SPIE Document Recognition & Retrieval Conference*, pp. 197–207 (2005)
26. Datta, R., et al.: IMAGINATION: A Robust Image-based CAPTCHA Generation System. In: *Proc. of ACM Multimedia Conf.*, pp. 331–334 (2005)
27. Hoque, M.E., Russomanno, D.J., Yeasin, M.: 2D Captchas from 3D Models. In: *Proceedings of the IEEE SoutheastCon 2006, Memphis (April 2006)*
28. Chew, M., Baird, H.S.: BaffleText: a Human Interactive Proof. In: *Proceedings of the 10th SPIE/IS&T Doc. Recog. Retr. Conf, DRR 2003 (2003)*
29. Baird, H.S., Bentley, J.L.: Implicit CAPTCHAs. In: *Proc. SPIE/IS&T Conf. on Document Recognition and Retrieval XII (DR & R 2005)*, pp. 191–196 (2005)
30. Moy, G., et al.: Distortion Estimation Techniques in Solving Visual CAPTCHAs. In: *Proc. of the CVPR 2004*, vol. 2 (2004)
31. Yan, J., Ahmad, A.S.E.: Breaking Visual CAPTCHAs with Naive Pattern Recognition Algorithms. In: *Proceedings of the ACSAC 2007 (2007)*
32. Chellapilla, K., Simard, P.Y.: Using Machine Learning to Break Visual HIPs. In: *Proc. of the Conf. on Neural Information Processing Systems, NIPS 2004 (2004)*

33. Chellapilla, K., Larson, K., Simard, P.Y., Czerwinski, M.: Computers beat Humans at Single Character Recognition in Reading based Human Interaction Proofs (HIPs). In: Procs. of CEAS (2005)
34. Hernandez-Castro, C.J., Ribagorda, A.: Analysis of the Teabag CAPTCHA version 1.2 (Submitted for publication)
35. DarkSEO programming: letter derotation,
<http://www.darkseoprogramming.com/2008/04/05/letter-derotation/>
36. PWNtcha, <http://caca.zoy.org/wiki/PWNtcha>
37. TROJ_CAPTCHAR.A Trojan horse to relay CAPTCHAs at TrendMicro,
<http://blog.trendmicro.com/captcha-wish-your-girlfriend-was-hot-like-me/>
38. Danchev, D.: Inside India's CAPTCHA solving economy,
<http://blogs.zdnet.com/security/?p=1835>
39. Van der Vorm, J.: Defeating audio (voice) CAPTCHAs,
<http://vorm.net/captchas>
40. Santamarta, R.: Breaking GMail's audio CAPTCHA,
<http://blog.wintercore.com/?p=11>
41. Kittenauth, O.W.: <http://www.thepcspy.com/kittenauth>
42. Techworld.com article on fallen CAPTCHAs,
<http://www.techworld.com/security/features/index.cfm?featureid=4168>
43. Hernandez-Castro, C.J., Ribagorda, A.: Pitfalls in CAPTCHA design and implementation: the Math CAPTCHA, a case study. *Computers & Security* (2009),
<http://dx.doi.org/10.1016/j.cose.2009.06.006>
44. Yeen, H.: Breaking CAPTCHAs without using OCR,
http://www.puremango.co.uk/2005/11/breaking_captcha_115/
45. W. Wieser. Captcha recognition via averaging,
<http://www.triplespark.net/misc/captcha/>