Face Recognition CAPTCHA Made Difficult

Terence Sim
National University of Singapore
13 Computing Drive
Singapore 117417
tsim@comp.nus.edu.sg

Hossein Nejati
National University of Singapore
13 Computing Drive
Singapore 117417
nejati@nus.edu.sg

James Chua
National University of Singapore
13 Computing Drive
Singapore 117417
chuayams@comp.nus.edu.sg

ABSTRACT

A CAPTCHA is a Turing test to distinguish human users from automated scripts to defend against internet adversarial attacks. As text-based CAPTCHAs (TBC) have become increasingly difficult to solve, image-based CAPTCHAs, and particularly face recognition CAPTCHAs (FRC), offer a chance to overcome TBC limitations. In this paper, we systematically design and implement a practical FRC, informed by psychological findings. We use gray-scale and binary images, which are computationally inexpensive to generate and deploy. Furthermore, our FRC complies with CAPTCHA design guidelines, thereby ensuring its robustness.

Categories and Subject Descriptors

I.5.4 [**Pattern Recognition**]: Application—*Computer Vision*; H.5.3 [**Group and Organization Interfaces**]: Web-based Interaction; K.6.5 [**Security and Protection**]: Authentication

Keywords

CAPTCHA; face recognition; reverse Turing test

1. INTRODUCTION

A Completely Automated Public Turing test to tell Computers and Humans Apart, or CAPTCHA, is an automatic challenge-response test to distinguish humans from machines. The most popular type of CAPTCHA is text-based CAPTCHAs (TBC), which have been in use for over a decade [10]. However in order to stay ahead of OCR technologies, TBCs have become increasingly difficult even for human solvers. In addition, TBCs are inherently and unnecessarily restrictive: users need to be literate in the language of the CAPTCHA.

In recent years, image-based CAPTCHAs (IBCs), and in particular face-based CAPTCHAs, have been introduced to address the shortcomings of TBCs. IBC capitalizes on the human ability to understand images, such as recognizing an animal [10], understanding a scene (e.g. Pix[1]), or analyzing a face image [1]. Therefore, IBCs are usable by people from a wider age range and education level than TBCs. And if done right, they can be language-independent.

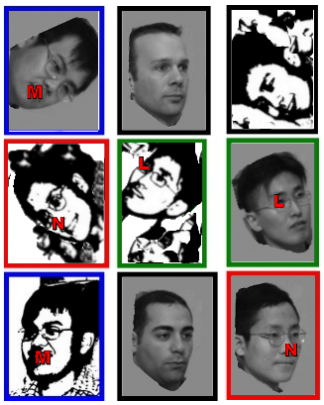

Figure 1: Our face recognition CAPTCHA: identify 3 pairs of faces from 9, by placing markers on corresponding facial parts (eyes, nose, mouth).

Among IBCs, face CAPTCHAs exploit a natural human ability — face detection and face recognition — to create tests that are easy for humans but difficult for machines. From the published literature (e.g. [7, 1, 4]), it is clear that most existing face CAPTCHAs feature face detection rather face recognition. Also, the few research work on face recognition CAPTCHAs (FRC) (e.g. [5, 2]) are fraught with problems: some are easily broken by random guessing or exhaustive attack, while others are solvable by current face recognition algorithms. These FRCs are clearly impractical and none of them comply with the CAPTCHA design guidelines [7], namely, being automatic and gradable, easy for humans, hard for machines, universal, resistant to random guessing and exhaustive attacks.

The purpose of this paper is to describe the first systematically designed face recognition CAPTCHA, based on sound design guidelines, psychological studies, and empirical experiments. We aim to make FRCs difficult, but only for machines. Our system uses only grayscale and binary images makes our FRC computationally inexpensive to generate and deploy. Furthermore, our FRC fulfills CAPTCHA design guidelines and is therefore robust against adversarial attacks. Results of testing our FRC on human users versus

[1]http://gs264.sp.cs.cmu.edu/cgi-bin/esp-pix

WWW'14 Companion, April 7–11, 2014, Seoul, Korea.
ACM 978-1-4503-2745-9/14/04.
http://dx.doi.org/10.1145/2567948.2577321.

state-of-the-art commercial face recognition software show that human users can solve the FRC with an 83.55% accuracy while machines fare only dismally.

2. THE HUMAN-MACHINE GAP

The key idea in our work is to exploit the performance gap between human and machine face recognition. By comparing psychological research with the literature on machine face recognition, we narrow our selection of face images and image operations to these four types. (1) Image crop: Humans can recognize a partially cropped face image [8]; (2) Facial expression: Although expressions introduce considerable deformation to facial appearance (and is thus difficult for machines), it has only a minor effect on a human's recognition ability [11]; (3) Head pose: Humans easily recognize a face even when the pose is changed [8]; and (4) Image binarization: Humans have the ability to compensate for poor image quality by gap-filling incomplete images [6]. Figure 1 shows how these 4 concepts are applied.

3. VALIDATION EXPERIMENTS

We fine-tuned our 4 choices above with extensive experiments conducted on both human volunteers (365 users from Amazon Mechanical Turk) and two commercial face recognition algorithms (PittPatt[2] and Luxand[3]). For our face images we used the popular Multi-PIE database [3]. Each facial identity has several original images exhibiting different facial expression, head pose, and illumination.

Through many experiments to test the effectiveness of different combinations of image cropping, pose and expression variation, and image binarization, we learned that (i) humans are better able to approximate the side-pose appearance of a face, after viewing the grayscale (instead of binary) image of the corresponding frontal face; (ii) tight cropping (no jaw-line) impairs human recognition more than loose cropping; (iii) cropping affects grayscale images more adversely than binary images. In addition, the presence of auxiliary information such as race, gender, or facial cues (e.g. glasses) dramatically improves human performance, as predicted by psychological studies [9]. On the other hand, machines cannot seem to exploit such ancillary information.

For the sake of brevity, we summarize our findings in Table 1. This compares the face recognition accuracy of humans versus machines on 2 tasks: Task (a) matches grayscale frontal with binary side facial images; while Task (b) matches grayscale side with side faces. It is clear that humans perform much better than current state-of-the-art algorithms. Their average accuracies are 78.47% and 83.55% for Tasks (a) and (b), respectively. In contrast, machines fare badly for both tasks.

From these validation experiments we build our FRC as a 3×3 grid of 9 face images (Figure 1): 4 grayscale, 5 binary, or vice versa. Among these 9 faces, there are 3 binary-grayscale side-side (Task (b)) pairs (and thus 3 non-matching faces). The task is to find the 3 matching pairs.

Can our FRC be cracked by random guessing? It is not difficult to calculate that the success rate of randomly pairing 6 out of 9 images is 1/240 ($= \frac{1\times3!}{{}^5C_1\times{}^4C_1^2\times{}^3C_1^2\times{}^2C_1}$). However, the task can be made harder by requiring the matching of facial features: the user also has to place 3 pairs of letters, M's, N's, and L's, on the mouths, noses and left eyes, respectively, on the matching faces. This simple addition drastically reduces the success of random guessing to 1 out of 716,636,160 ($= \frac{1\times3!}{{}^5C_1\times{}^4C_1^2\times12^2\times{}^3C_1^2\times{}^2C_1\times12^2}$). This is about 10 times less likely than getting killed by a falling asteroid (according to to CNET news[4]).

[2]http://www.pittpatt.com

[3]http://www.luxand.com/index.php

	Task (a)	Task (b)
Humans	78.47%	83.55%
Luxand	6.0%	0.0%
PittPatt	2.0%	0.0%

Table 1: Recognition accuracies of humans vs. machines on tasks with auxiliary information.

4. CONCLUSION

This paper describes how to design face recogntion CAPTCHAs that are difficult for machines, but easy for humans. Our work combines psychological insights into human recognition ability with knowledge of how current algorithms work, while fulfilling design principles for CAPTCHAs. We hope that our work enhances security for web-based services that need to distinguish humans from machines.

This research was partially carried out at the SeSaMe Centre. It is supported by the Singapore NRF under its IRC@SG Funding Initiative and administered by the IDMPO.

5. REFERENCES

[1] D. DSouza, P. C. Polina, and R. V. Yampolskiy. Avatar captcha: Telling computers and humans apart via face classification. In *EIT*, pages 1–6. IEEE, 2012.

[2] G. Goswami, R. Singh, M. Vatsa, B. Powell, and A. Noore. Face recognition captcha. In *BTAS*, pages 412–417, 2012.

[3] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker. Multi-pie. In *FG*, pages 1–8, 2008.

[4] J. Kim, S. Kim, J. Yang, J.-h. Ryu, and K. Wohn. Facecaptcha: a captcha that identifies the gender of face images unrecognized by existing gender classifiers. *Multimedia Tools and Applications*, pages 1–23, 2013.

[5] D. Misra and K. Gaj. Face recognition captchas. *AICT*, page 122, 2006.

[6] L. Pessoa, E. Thompson, and A. NoÃń. Finding out about filling-in: A guide to perceptual completion for visual science and the philosophy of perception. *Behavioral and Brain Sciences*, 21:723–748, 1998.

[7] Y. Rui and Z. Liu. Artifacial: Automated reverse turing test using facial features. *Multimedia Systems*, 9(6):493–502, 2004.

[8] P. Sinha, B. Balas, Y. Ostrovsky, and R. Russell. Face recognition by humans: Nineteen results all computer vision researchers should know about. *Proceedings of the IEEE*, 94(11):1948 –1962, nov. 2006.

[9] M. Unnikrishnan. How is the individuality of a face recognized? *Journal of Theoretical Biology*, 261(3):469 – 474, 2009.

[10] L. von Ahn, M. Blum, N. Hopper, and J. Langford. Captcha: Using hard ai problems for security. *EUROCRYPT 2003*, pages 646–646, 2003.

[11] J. S. Winston, R. N. A. Henson, M. R. Fine-Goulden, and R. J. Dolan. fmri-adaptation reveals dissociable neural representations of identity and expression in face perception. *J. Neurophysiol.*, 92:1830, 2004.

[4]http://tinyurl.com/a2yr54a