
Breaking a Visual CAPTCHA

1 **Anonymous Author(s)**

2 Affiliation

3 Address

4 *email*

5

6 **Abstract**

7 Visual CAPTCHAs have recently become a practical mainstream online
8 Turing test method to counter malicious or unauthorized access by
9 automated scripts. However, naïve implementations of this system can
10 create vulnerabilities that can be easily exploited with minimal
11 computational cost. In this paper, I will demonstrate a simple attack against
12 a popular open source web blog plugin, PHP CAPTCHA.

13

14 **1 Introduction**

15 Web sites offering services to a large public userbase are prime targets for spammers who
16 advertise products and scams in an effort to entrap users. To increase posting efficiency,
17 rather than manually post the content by hand, many spammers employ automated computer
18 scripts to crawl through the web and post content on multiple sites en masse. Consequently,
19 many site administrators have chosen to employ a Turing test approach to differentiate
20 between legitimate users and bots [1].

21 The most popular form of the CAPTCHA, a *Completely Automated Public Turing test to tell*
22 *Computer and Humans Apart*, is a visual image containing warped text characters with
23 cluttering, distortion, and noise. The premise for this approach assumes that humans are
24 more capable in decoding the text than the automated scripts.

25



26 Figure 1: Google CAPTCHA example

27

28 Modern CAPTCHAs, such those used for Google, are implemented specifically to thwart
29 straightforward optical character recognition (OCR) techniques employed by typical
30 document scanning software. Indeed, these systems are so designed in attempts to dissuade
31 even legitimate software developers from bypassing their systems. However, there are no
32 strict guidelines for designing CAPTCHA systems, which leads to several vulnerable
33 implementations. In this paper, I will investigate one such implementation, PHP Captcha
34 [2], and show how a straightforward combination of techniques from computer vision,
35 graphics, and machine learning can be used to attack such a system.

36

37 **2 Related Work**

38

39 A recent security research paper [3] surveyed a variety of modern private CAPTCHA
40 implementations and presented an anti-CAPTCHA implementation that effectively defeated
41 most simple CAPTCHA ranging from eBay, Digg, and CNN. The paper outlined the
42 weaknesses and common vulnerabilities of these systems and provided suggestions for
43 improvement. My paper will attempt to verify some of the techniques described in this
44 paper against a modern open-source CAPTCHA implementation. My implementation will
45 also differ by using a computer graphics approach to segmentation, and a low-cost feature
46 learning technique with K-means [9] for use in classification.

47

48 **2 PHP Captcha**

49 The CAPTCHA implementation I will attack in this paper is a popular open-source web page
50 add-on running on PHP [2]. This software appears frequently as a plugin in web blogs to
51 protect comment sections of pages from automated spammers. Prior to the rise in popularity
52 in ReCaptcha [4], this CAPTCHA has seen much popularity in WordPress blogs.

53

54 The feature set of PHP Captcha allows site administrators to customize the distortion and
55 clutter intensity of the CAPTCHA, the background, font, the color of the text and line, as
56 well as the thickness and number of the lines. The software also provides an audio
57 CAPTCHA, which has its own share of exploits [5], but this paper will focus on attacking
58 the visual CAPTCHA.

59

60 **3 Preprocessing**

61 Images have to be preprocessed to remove unwanted clutter from the characters prior to
62 segmentation. This clutter can interfere with processes further down in the pipeline, as they
63 introduce extraneous artifacts that can then be mistaken as part of the character body. This
64 section describes a few PHPCaptcha setups that require additional work.

65

66 **4.1 Foreground/Background Segmentation**

67



68

Figure 2: Input image (left) and two K-means color segments (right)


69

70 Many websites tend to customize the CAPTCHA with a background image of their own.
71 Some may even have a variety to improve the aesthetics surrounding such a mundane setup.
72 However, most backgrounds tend to be fixed and reused. This allows modeling the
73 background for background subtraction. Simple backgrounds with pixel colors significantly
74 different from the text in the foreground also can be segmented via K-means color
75 segmentation with merely two or three clusters. Figure 2 shows segmentation the result
76 with 2-means clustering.

77

78 **4.2 Noise Removal**

79



80

Figure 3: Input image (left) and denoised output (right)

81

82 Another scheme to avoid segmentation is via the addition of random salt and pepper noise
83 with the same color as the text. This is vulnerable to an iterative removal of pixels that
84 contribute poorly to the average energy within its surrounding patch until there are no
85 further changes. This is essentially a crude execution of the Gibbs algorithm [6]. Removal
86 of larger segments that have been collected can be done with an edge-preserving median
87 filter. Note that very thin lines on the image would also be removed during this denoising.

88
89
90


4.2 Line Removal



91 Figure 4: Image vulnerable to line removal by color segmentation

92
93
94
95
96
97

Lines with thickness close to width of the character body are difficult to distinguish from the characters themselves. However, PHPCaptcha allows for the user to modify the color templates of the lines and text, which leads to a vulnerable case where the text, background, and line can be fully separated through once again by K-means clustering by color.



98 Figure 5: Vulnerable PNG input, line removal, and repair with inpainting

99
100
101
102
103
104
105
106
107
108
109
110

As of version 3.0 of the software, PHPCaptcha is also found to have an unpatched vulnerability due to the palette coding of PNG images generated by PHP graphics library. For vanilla CAPTCHA images presented in PNG format without a custom background, regardless of the amount of clutter by lines or noise or signature text, it is possible to completely separate the text and clutter by merely comparing the palette IDs in the image encoding. As in Figure 5, the lines can be removed entirely due to the palette ID over the pixels on the line (and over the text) being different from the palette ID of the text. A simple XOR operator over the values generates the center image.

Removing the lines creates gaps within the image that may impair segmentation and recognition pieces later in the pipeline. Consequently, it is necessary to fill in the missing pixels through inpainting [8].

111
112


4 Character Segmentation

113 Tightly spaced text create difficulties in segmentation due to difficulty distinguishing which
114 given pixels belong to a character, especially if the characters are actually touching. In
115 particular, it is difficult to segment by horizontal spacing alone. Here I introduce an
116 alternative to segmentation inspired from computer graphics: seam-carving [7].

117 A seam is a path of connected pixels traveling from the top to the bottom of the image.
118 Pixels must neighbor in location by exactly one pixel smoothly without ever becoming
119 horizontal or a sharp changing in direction. These requirements ensure seams travel from
120 the top of the image to the bottom as steeply and quickly as possible while spatially close.

121 Seams are produced differently from the seam-carving paper in that a penalty function is
122 used in place of an energy function: Seams that pass through a pixel of any character (a
123 black-colored pixel) accumulate penalty costs while passing through the background (a pixel
124 of any other color) does not. Consequently, the seam will prefer to go around a character
125 rather than through it when possible.

126



127 Figure 6: Clockwise from top left: input image, seam costs, seam carve, and binning

128

129 This property can be used for segmentation. A seam passing in between characters can act as
130 the segmentation boundary that separates two clusters: the left side and the right side. Given
131 that the background subtraction has already labeled that all characters have a black color, the
132 set intersection of these black pixels and the pixels split by the boundaries produce the
133 desired clustering.

134 The seam boundaries can then be taken as bins. Empty bins can be pruned to return only
135 those that contain characters. Notice that with sufficiently small bins this approach can be
136 used for CAPTCHAs that vary the number of characters. This approach also allows forced
137 segmenting between touching characters, which still have a smaller accumulated penalty
138 compared to seams crossing a full character body.

139

140 3 Crop and Resize

141 Pixels clustered via binning are then cropped and rescaled to a suitable template image of
142 40x60 pixels each. These letters are then grouped together for the training set. Since
143 CAPTCHA by design fails if even one character is misidentified, this problem can be
144 interpreted as standard OCR problem for individual characters after removal of clutter.

145



146 Figure 7: Typical characters samples extracted for a single letter.

147

148 PHPCaptcha by default installation settings do not distinguish between capital or lowercase
149 variations of the letters. Therefore, we can train these together as a single unit, and thus
150 reducing the loading of training, and the need to gather a sufficiently large training
151 collection for each individual letter. This also allows combing similar characters like S, s, B,
152 B, and W, w without additional costs in storage.

153

154 4 Feature Learning and Classifier Training

155 Applying techniques from Coate's [9], it is possible to generate basis or filters that represent
156 parts of the collected image patches over the training set. The features on the left have PCA
157 whitening applied rather than ZCA due to implementation restrictions of the library during
158 the time of coding. However, we can observe quantitatively the Gabor-like resemblance of
159 the features compared to the specific patches on the right image.

160




Figure 7: Dictionary features learned with K-means on the CAPTCHA training set with whitening PCA (left) and without (right).

161
162
163
164
165

5 Results

166 Using the code for PHP Captcha, 500 training samples and 200 test samples were generated.
167 For this paper, the samples were generated with straightforward background subtraction,
168 noise reduction, and line removal. Preprocessing and segmentation is applied for each
169 sample.

170 Unfortunately, the implementation of the system through Python and the Scikit-learn library
171 has not been fully successful due to the work-in-progress nature of the Coates' recent paper.
172 As a result, only the bases can be computed via K-means but without feeding it into a
173 convolution neural network or a simple linear SVM for training and classification.
174 However, performance as indicates in Coates' paper suggest relatively high performance for
175 classification to the extent of higher-end RBMs. We can expect similar performance here.

176
177

6 Limitations and Further Work

178 Line removal becomes particularly problematic when the colors cannot be segmented, and
179 the image format is not vulnerable, and the lines are as thick as the characters. In this case,
180 we can attempt to estimate the location of some of the lines by applying a Progressive
181 Probabilistic Hough Transform [10] to the image gradient. In the particular OpenCV
182 implementation, this returns endpoints of line segments each supported by the hypothesis
183 that points in between contribute to that straight line or curve. However, wavy lines cutting
184 across many characters horizontally prove difficult to remove without damaging the
185 characters, and many line segments can remain undetected due to the characters themselves
186 contributing as "noise" in Hough space.

187 In this scenario, the current approach would rely heavily on the seam carving segmentation
188 to make the least costly cut around or through these lines. However, it may be possible to
189 combine information from Hough estimations by reducing the cost of cutting through pixels
190 on or close to a suspected line segment.

191 References

- 192 [1] M. Naor. Verification of a human in the loop or Identification via the Turing test. Available
193 electronically: <http://www.wisdom.weizmann.ac.il/~naor/PAPERS/human.ps>, 1997.
- 194 [2] D. Phillips. SecurImage PHP Captcha 3.0. <http://www.phpcaptcha.org>, 2011.
- 195 [3] E. Bursztein, M. Martin, and J. Mitchell. Text-based CAPTCHA strengths and weaknesses. In
196 Proceedings of the 18th ACM conference on Computer and communications security (CCS '11). ACM,
197 New York, NY, USA, 125-138, 2011.

- 198 [4] Google Inc. reCAPTCHA. <http://www.google.com/recaptcha>
- 199 [5] A. Anonymous. Web App Security blog. <http://www.idontplaydarts.com/2011/05/exploit->
200 [phpcaptcha-securimage/](http://www.idontplaydarts.com/2011/05/exploit-), 2011.
- 201 [6] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions and the Bayesian restoration of
202 images*. Journal of Applied Statistics, 20(5):25–62, 1993.
- 203 [7] S. Avidan and A. Shamir, "Seam carving for content-aware image resizing," ACM Trans. Graph.,
204 vol. 26, no. 3, 2007.
- 205 [8] A. Telea, "An image inpainting technique based on the fast marching method," In Proceedings
206 of Journal of Graphics Tools, vol.9, no.1, pp.25–36, 2004.
- 207 [9] An Analysis of Single-Layer Networks in Unsupervised Feature Learning Adam Coates,
208 Honglak Lee and Andrew Ng. In NIPS*2010 Workshop on Deep Learning and Unsupervised
209 Feature Learning.
- 210 [10] J. Matas, C. Galambos, and J. Kittler, "Robust Detection of Lines Using the Progressive
211 Probabilistic Hough Transform", presented at Computer Vision and Image Understanding, pp.119-
212 137, 2000.