24. CAPTCHAs

24.1. Using CAPTCHAs

Consider the following scenario: you've created a website that allows users to upload a picture. Your server will scan the picture for text using a compute-intensive algorithm and return the text to the user. An adversary wants to mount a denial-of-service (DoS) attack on your website by uploading lots of bogus images, forcing your server to run the expensive algorithm on all the bogus images.

Consider another scenario: Your website has a login page. An adversary wants to steal a legitimate user's account, so they try to brute-force the user's password by submitting login requests with every possible password.

Generally, when we're building websites, we'd like to build websites for people: we don't want robots. CAPTCHAs are a test that ask the fundamental question: *Is this a human?* Consequently, when we design CAPTCHAs, we want to choose problems that are easy for humans, but difficult for computers.

CAPTCHAs are primarily focused on machine vision problems, which are traditionally difficult for computers to solve. Historically, CAPTCHAs consist of a series of distorted letters or words. There are a wide variety of CAPTCHAs: some with color, some with low contrasts, some with merged-together letters, etc. A more recent example you may be familiar with is Google's reCAPTCHA algorithm, which shows you some images and asks you to identify the objects in the pictures (e.g. "Select all images with boats.")

24.2. Issues with CAPTCHAs

There's an inherent arms race present here: as solving algorithms get better, our defense deteriorates. The reason why CAPTCHAs have gotten so much harder over the last decade is because individuals have spent time creating much better solving algorithms - and we're reaching a point where it's becoming more and more difficult for humans to solve CAPTCHAs quickly.

Of course, those implementing CAPTCHAs often miss the original motivation behind their development. The original CAPTCHA paper included the subtitle "How Lazy Cryptographers do AI" as the intent was to force attackers to solve harder problems in machine vision. Now modern

CAPTCHAs such as Google ReCAPTCHA are focused on getting humans to provide training data for AI systems which means the CAPTCHAs are inherently self defeating for those deploying the CAPTCHA.

In some cases, it's necessary to provide an alternative, accessible CAPTCHA method, such as an audio-based spoken phrase that a human is required to transcribe. In this case, we've unintentionally opened up a new attack vector: attackers may now target the audio-based CAPTCHA, which may be easier to solve than the traditional image-based CAPTCHA.

If you search "crack CAPTCHA" on Google, you'll likely find many CAPTCHA solving services for as low as $0.10 cents per CAPTCHA. These services use humans to do the actual work. These days, a CAPTCHA no longer asks the question of "Is this a human or a bot?" Instead, it says "Is this a human, or a bot willing to spend a fraction of a penny?"

The takeaway: if something is worth $0.10 or more to an attacker, CAPTCHAs do not work.