

An Empirical Study & Evaluation of Modern CAPTCHAS

Andrew Searles
UC Irvine

Yoshimichi Nakatsuka*
ETH Zürich

Ercan Ozturk
UC Irvine

Andrew Pavard
Microsoft

Gene Tsudik
UC Irvine

Ai Enkoji†
LLNL

Abstract

For nearly two decades, CAPTCHAS have been widely used as a means of protection against bots. Throughout the years, as their use grew, techniques to defeat or bypass CAPTCHAS have continued to improve. Meanwhile, CAPTCHAS have also evolved in terms of sophistication and diversity, becoming increasingly difficult to solve for both bots (machines) and humans. Given this long-standing and still-ongoing arms race, it is critical to investigate how long it takes legitimate users to solve modern CAPTCHAS, and how they are perceived by those users.

In this work, we explore CAPTCHAS *in the wild* by evaluating users’ solving performance and perceptions of *unmodified currently-deployed* CAPTCHAS. We obtain this data through manual inspection of popular websites and user studies in which 1,400 participants collectively solved 14,000 CAPTCHAS. Results show significant differences between the most popular types of CAPTCHAS: surprisingly, solving time and user perception are not always correlated. We performed a comparative study to investigate the effect of *experimental context* – specifically the difference between solving CAPTCHAS directly versus solving them as part of a more natural task, such as account creation. Whilst there were several potential confounding factors, our results show that experimental context could have an impact on this task, and must be taken into account in future CAPTCHA studies. Finally, we investigate CAPTCHA-induced user task *abandonment* by analyzing participants who start and do not complete the task.

1 Introduction

Automated bots pose a significant challenge for, and danger to, many website operators and providers. Masquerading as legitimate human users, these bots are often programmed to scrape content, create accounts, post fake comments or reviews, consume scarce resources, or generally (ab)use other

website functionality intended for human use [31, 46]. If left unchecked, bots can perform these nefarious actions at scale. CAPTCHAS are a widely-deployed defense mechanism that aims to prevent bots from interacting with websites by forcing each user to perform a task, such as solving a challenge [5]. Ideally, the task should be straightforward for humans, yet difficult for machines [68].

The earliest CAPTCHAS asked users to transcribe random distorted text from an image. However, advances in computer vision and machine learning have dramatically increased the ability of bots to recognize distorted text [35, 41, 74], and by 2014, automated tools achieved over 99% accuracy [39, 62]. Alternatively, bots often outsource solving to CAPTCHA *farms* – sweatshop-like operations where humans are paid to solve CAPTCHAS [54]. In light of this, CAPTCHAS have changed and evolved significantly over the years. Popular CAPTCHA tasks currently include object recognition (e.g., “select squares with...”), parsing distorted text, puzzle solving (e.g., “slide the block...”), and user behavior analysis [39, 62]. It is therefore critical to understand and quantify how long it takes legitimate users to solve current CAPTCHAS, and how these CAPTCHAS are perceived by users.

Several prior research efforts have explored CAPTCHA solving times, e.g., [24, 27, 33, 37, 58, 67]. For example, over a decade ago, Bursztein et al. [27] performed a large-scale user study, using over 1,100 unique participants from Amazon Mechanical Turk (MTurk) [3] as well as CAPTCHA farms. Their results showed that CAPTCHAS were often more difficult or took longer to solve than was expected. There was a loose correlation between time-to-annoyance and abandonment, with higher abandonment rates observed for CAPTCHAS that took longer to solve. The same study also showed several demographic trends, e.g., users outside the US typically took longer to solve English-language CAPTCHA schemes. However, since this study, the CAPTCHA ecosystem has changed substantially: new CAPTCHA types emerged, input methods evolved, and Web use boomed.

More recently, Feng et al. [33] used a similar methodology, with 202 participants, to study the usability of their newly-

*:† Work done while at UC Irvine.

proposed senCAPTCHA in comparison to text, audio, image, and video-based CAPTCHAS. They found that senCAPTCHA outperformed the alternatives, both in terms of solving time and user preference. They used Securimage [55], a free open-source PHP script, to generate text and audio CAPTCHAS, and they implemented their own image and video CAPTCHAS.

Building upon and complementing prior work, this paper evaluates CAPTCHAS *in the wild* – specifically, the solving times and user perceptions of *unmodified* (i.e., not re-implemented) *currently-deployed* CAPTCHA types. We first performed a manual inspection of 200 popular websites, based on the Alexa Top websites list [2], to ascertain: (1) *how many* websites use CAPTCHAS, and (2) *what types* of CAPTCHAS they use. Next, we conducted a 1,000-participant user study using Amazon MTurk, wherein each participant was required to solve 10 different types of CAPTCHAS. We collected information about participants’ CAPTCHA solving times, relative preferences for CAPTCHA types, types of devices used, and various demographic information.

One notable aspect of our user study is that we attempted to measure the impact of experimental context on participants’ CAPTCHA solving times. Half of the participants were directly asked to solve CAPTCHAS, whilst the other half were asked to create accounts, which involved solving CAPTCHAS as part of the task. The latter setting was designed to measure CAPTCHA solving times *in the context* of a typical web activity.

One inherent limitation of any user study, especially when using MTurk, is that we cannot ensure that all participants who begin the study will complete it. All of our results should therefore be interpreted as referring to *users who are willing to solve* CAPTCHAS, rather than users in general.

Indeed, having noted that some participants began but did not complete our main study, we conducted a secondary MTurk study specifically designed to quantify how many users abandon their intended web activity when confronted with different types of CAPTCHAS. We believe that CAPTCHA-induced *user abandonment* is an important – yet understudied – consideration, since every abandoned task (e.g., purchase, account creation) represents a potential loss for the website.

To facilitate reproducibility and enable further analysis, we provide the entire anonymized data-set collected during our user studies, along with our analysis code.²

2 Research Questions & Main Findings

We now present our research questions and summarize our main findings. Table 1 shows how our findings relate to prior work at a high level, with detailed comparisons in Section 7.

RQ1: How long do human users take to solve different types of CAPTCHAS? Specifically, we aimed to measure solving times for CAPTCHAS that users are likely to encounter (e.g., those used on popular websites). Our results align with

previous findings [24, 27, 33] in showing that there are significant differences in mean solving times between CAPTCHA types. For comparison, we also identified the current fastest attacks on each type of CAPTCHA (Table 3).

RQ2: What CAPTCHA types do users prefer? In order to understand users’ relative preference for various types of CAPTCHAS, we asked participants to rate all CAPTCHA types on a Likert scale of 1 – 5, from least to most enjoyable. Our results show that there are marked differences in participants’ preferences, with average preference scores ranging from 2.76 to 3.94. Our results also show that average solving time is *not fully correlated* with participants’ preferences, which means that other factors, beyond the amount of time required to solve a CAPTCHA, influence participants’ preferences. Our analysis of data from prior studies [33, 48, 66] shows that their data supports this finding (even if they do not discuss it explicitly).

RQ3: Does experimental context affect solving time? Specifically, we aimed to quantify the difference in solving times between the setting where participants are directly tasked with solving CAPTCHAS versus the setting in which participants solve CAPTCHAS as part of a typical web activity, such as user account creation. We therefore ran two separate versions of our main user study: *direct* and *contextualized*, which we describe in detail in Section 4.2. Whilst there were several potential confounding factors in our study, our results show that experimental context could have an impact on CAPTCHA user studies, with the difference in mean solving times as high as 57.5% in our study.

RQ4: Do demographics affect solving time? We analyzed different self-reported metrics including age, gender, country of residence, education, Internet usage, device type and input method. In line with prior results [27], we found that all types of CAPTCHAS take longer for older participants. Specifically, [27] reported an increase in solving time for text-based CAPTCHAS of 0.03 seconds per year of participant age. Our results show an even stronger dependence with an average increase across all CAPTCHA types of 0.09 seconds per year. Additionally, [27] showed that participants with a PhD solved CAPTCHAS faster than all other educational groups. In contrast, our results show that our participants’ self-reported level of education does not correlate with their solving times.

RQ5: Does experimental context influence abandonment? Specifically, we aimed to quantify the extent to which abandonment within a CAPTCHA user study is influenced by i) experimental context, and ii) the amount of compensation offered. For different combinations of the above variables, we found that between 18% and 45% of participants abandoned the study after the presentation of the first CAPTCHA. Only one prior CAPTCHA user study [27] disclosed their observed rate of abandonment, which is similar to that observed in our study. Overall, participants in the contextualized setting were 120% more likely to abandon than their peers in the direct setting. This connection between experimental context and user abandonment is a new finding.

²<https://github.com/sprout-uci/captcha-study>

Table 1: Summary of research questions and main findings.

	Findings supporting prior work	Findings contradicting prior work	New findings on CAPTCHAS
RQ1: How long does it take humans to solve different types of CAPTCHAS?	Solving time across CAPTCHA types has a large degree of variance. [24, 27, 33]		
RQ2: What CAPTCHA types do users prefer?	Solving time is not correlated with user preference. [33, 48, 66]		
RQ3: Does experimental context affect solving time?			Solving time is heavily influenced by experimental context, with differences in means up to 57.5%.
RQ4: Do demographics affect solving time?	Age has an effect on solving time. [27]	Self-reported education does not correlate with solving time. [27]	
RQ5: Does experimental context influence abandonment?	High abandonment rates observed in CAPTCHA user studies. [27]		Experimental context directly affects the rate of abandonment.

3 Website Inspection

To understand the landscape of modern CAPTCHAS and guide the design of the subsequent user study, we manually inspected the 200 most popular websites from the Alexa Top Website list [2]. Where applicable, we use the terminology from the taxonomy proposed by Guerar et al. [40].

Our goal was to imitate a normal user’s web experience and trigger CAPTCHAS in a natural setting. Although CAPTCHAS can be used to protect any section or action on a website, they are often encountered during user account creation to prevent bots creating accounts. Thus, for each website, we investigated the process of creating an account (wherever available). Of the inspected websites, 185 had some type of account creation process, and we could successfully create accounts on 142 websites. Distinct domains operated by the same organization (e.g., amazon.com and amazon.co.jp) were counted separately. We visited each website twice: once with Google Chrome in incognito mode, and once with the Tor browser over the Tor network [17]. We used incognito mode to avoid websites changing their behavior based on cookies presented by our browser. We used Tor since anecdotal evidence suggests Tor users are asked to solve CAPTCHAS more frequently and with greater difficulty than non-Tor users. If no CAPTCHAS were displayed, we searched the page source for the string “CAPTCHA” (case insensitive).

Ethical considerations: Based on the Guidelines for Internet Measurement Activities [28], we did not engage in malicious behavior which may trigger additional CAPTCHAS. We used only manual analysis to avoid various challenges that arise from automated website crawling.

3.1 Results and analysis

Figure 1 shows the distribution of CAPTCHA types we observed during our inspection. The most prevalent types were:

reCAPTCHA [11, 14, 15] was the most prevalent, appearing on 68 websites (34% of the inspected websites). This

is a Google-owned and operated service that presents users with “click” tasks, which include behavioral analytics and may potentially result in an image challenge. reCAPTCHA allows website operators to select a difficulty level, ranging from “easiest for users” to “most secure”.

Slider-based CAPTCHAS appeared on 14 websites (7%). These typically ask users to slide a puzzle piece into a corresponding empty spot using a drag interaction. The timing and accuracy is checked for bot-like behavior.

Distorted Text CAPTCHAS appeared on 14 websites (7%). We observed differences in terms of text type, color, length, masking, spacing, movement, and background. Text type varied in several ways: 2D or 3D, solid or hollow, font, and level of distortion. Certain CAPTCHAS used masking, i.e., lines or shapes obscured parts of the letters.

Game-based CAPTCHAS appeared on 9 websites (4.5%). These present users with dynamic games and compute a risk profile from the results. For example, users are asked to rotate an image or select the correctly oriented image.

hCAPTCHA [9] appeared on 1 website. This is a service provided by Intuition Machines, Inc. that was recently adopted by Cloudflare [57] and is gaining popularity.

Invisible CAPTCHAS were found on 12 websites (6%). These websites did not display any visible CAPTCHAS, but contained the string “CAPTCHA” in the page source.

Other CAPTCHAS found during our inspection included: a CAPTCHA resembling a scratch-off lottery ticket; a CAPTCHA asking users to locate Chinese characters within an image; and a proprietary CAPTCHA service called “NuCaptcha” [13].

3.2 Potential limitations

Choice of website list: There are several lists of “popular” websites that could be used for this type of study, including the Alexa Top Website list [2], Cisco Umbrella [6], Majestic [16], TRANCO [56], Cloudflare Radar [7], and SecRank TopDomain [71]. These lists vary because of the differences in the methodology used to identify and rank websites. Following




Figure 1: Discrete distribution of discovered CAPTCHAs (full data available in the accompanying dataset).

the work of Bursztein et al. [27] and the recommendation of Scheitle et al. [60], we used the Alexa list.

Number of inspected websites: Since our website inspection was a manual process, we could only inspect the top 200 websites. This may also introduce a degree of systemic bias towards the types of CAPTCHAs used on the most popular websites. However, we specifically chose these websites because they are visited by large numbers of users.

Lower bound: Since we did not exercise all possible functionality of every website, it is possible that we might not have encountered all CAPTCHAs. Therefore, our results represent a lower bound, while the actual number of deployed CAPTCHAs may be higher. Nevertheless, we believe that we identified the most prevalent CAPTCHA types across all inspected websites.

Timing: Web page rankings change on the daily basis and CAPTCHAs shown by the same service may change. Given that our inspection was performed at a particular point in time, the precise results will likely change if the inspection were repeated at a different point in time. However, as explained above, we believe that the identified set of CAPTCHA types is representative of currently-deployed CAPTCHAs.

Other types of CAPTCHAs: We only inspected mainstream websites (i.e., those that would appear on a top websites list). This means that there could be CAPTCHAs that are prevalent on other types of websites (e.g., on the dark web) but are not included in our study. However, studying these *special-purpose* CAPTCHAs might require recruiting participants who have prior experience solving them, which was beyond the scope of our study.

Impact of limitations: The above limitations could have had an impact on the set of CAPTCHA types we identified and subsequently used in our user study. However, we have high confidence that the CAPTCHA types we identified are a realistic sample of those a real user would encounter during typical web browsing. For instance, BuiltWith [5] has analyzed a dataset of 673 million websites and identified 15.2

million websites that use CAPTCHAs. reCAPTCHA accounts for 97.3% and hCAPTCHA for a further 1.4%. The CAPTCHA types used in our study therefore account for over 98% of CAPTCHAs in this large-scale dataset.

4 User Study

Having identified the relevant CAPTCHA types, we conducted a 1,000 participant online user study to evaluate real users’ solving times and preferences for these types of CAPTCHAs. Our study was run using Amazon MTurk and can be summarized into the following four phases:

1. Introduction: Participants were first given an overview of the study and details of the tasks to complete.

2. Pre-study questions: All participants were then asked to provide demographic information by answering the pre-study questions shown in Table 11 in Appendix B.

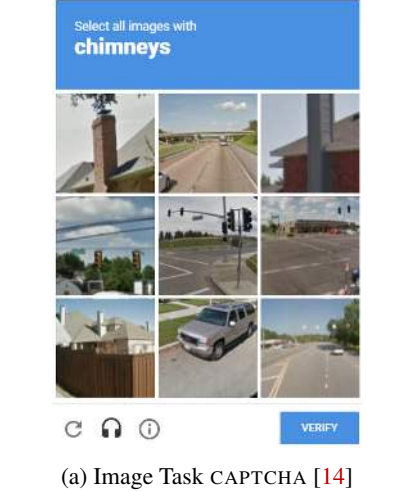
3. Tasks: Participants were asked to complete tasks, which included solving exactly ten CAPTCHAs, presented in random order. Unless otherwise stated, each CAPTCHA was *unique* (i.e., freshly generated per participant). Participants had to solve each CAPTCHA in order to progress to the next step, thus preventing them from speeding through the study.

4. Post-study question Finally, participants were asked questions about the CAPTCHAs they had just solved. The exact questions and possible answers are shown in Table 11 in Appendix B.


4.1 Choice of CAPTCHAS

Based on our website inspection (Section 3), we selected the following ten types of CAPTCHAs:

- Two reCAPTCHA v2 CAPTCHAs: one with the setting *easiest for users* and the other with *most secure*. Note that we do not have control over whether the user is shown an image-based (Figure 2a) challenge in addition to the click-based (Figure 2b) task.
- Two game-based CAPTCHAs from Arkose Labs [4]: one required using arrows to rotate an object (Figure 3a) and the other required selecting the upright object (Figure 3b).
- Two hCAPTCHAs [9]: one with easy and one with difficult settings (Figure 5).
- One slider-based CAPTCHA from Geetest [8]: we selected Geetest because it was used on several of the inspected websites and offers a convenient API (Figure 4).
- Three types of distorted text CAPTCHAs (Figure 6): (a) the *simple* version had four unobscured characters, (b) the *masked* version had five characters and included some masking effects, and (c) the *moving* version contained moving characters.

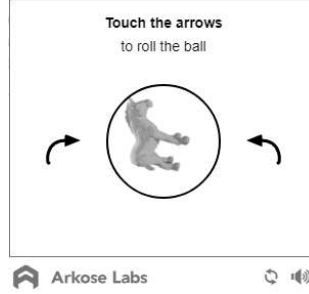


(a) Image Task CAPTCHA [14]




(b) v2 checkbox CAPTCHA [11]

Figure 2: reCAPTCHA [11, 14, 15]



(a) Rotation CAPTCHA



(b) Orientation selection

Figure 3: Arkose Labs [4]



Figure 4: Geetest [8]




Figure 5: hCAPTCHA [9]




Figure 6: Distorted text CAPTCHAS

These form a representative sample of CAPTCHAS we encountered in our website inspection. Although hCAPTCHA only appeared once, we included it since it is an emerging image-based approach, which claims to be the largest independent CAPTCHA service [10].

4.2 Direct vs. contextualized settings

We initially hypothesized that we would observe a difference in behavior depending on experimental context. In order to evaluate this, we designed two settings of the study: 500 participants completed the *direct setting*, whilst the other 500 completed the *contextualized setting*. In both settings, each participant solved exactly ten CAPTCHAS in random order.

Direct setting: This setting was designed to match previous CAPTCHA user studies, in which participants are directly asked to solve CAPTCHAS. The MTurk study title was “CAPTCHA User Study” and the instructions in the first phase informed users that their task was to solve CAPTCHAS. In the second phase, in addition to the basic demographic information, participants were asked about their experience with and perception of CAPTCHAS; see Table 11 in Appendix B. In the third phase, participants were shown ten CAPTCHAS in random order. The fourth phase was the same for both settings.

Contextualized setting: This setting was designed to measure CAPTCHA solving behavior *in the context* of a typical web activity. We selected the task of user account creation, as this often includes solving a CAPTCHA. The MTurk study title was “Account Creation User Study” and the first and second phases did not mention CAPTCHAS. In the third phase, participants were asked to complete ten typical user account creation forms, each displaying a CAPTCHA *after* the participant clicked submit, as is often the case on real websites. This sequencing allowed us to precisely measure the CAPTCHA solving time in isolation from the rest of the account creation task. The account creation task was a basic web form asking for a randomized subset of: name, email address, phone number, password, and address. To avoid collecting personally identifiable information, participants were provided with

Table 2: Summary of demographic data for the 1,400 participants of the main user study.

Age	Residence	Education	Gender	Device Type	Input Method	Internet Use
30 - 39 (531)	USA (985)	Bachelors (822)	Male (832)	Computer (1301)	Keyboard (1261)	Work (860)
20 - 29 (403)	India (240)	Masters (243)	Female (557)	Phone (74)	Touch (125)	Web surf (397)
40 - 49 (271)	Brazil (50)	High school (210)	Non-Binary (11)	Tablet (25)	Other (14)	Education (87)
50 - 59 (106)	Italy (27)	Associate (98)				Gaming (30)
≥ 60 (58)	UK (24)	PhD (24)				Other (26)
18 - 19 (31)	Other (74)	No degree (3)				

synthetic information at each step. Each page also included a large banner clearly stating not to enter any personal information. The fact that we were specifically measuring CAPTCHA solving time was only revealed to participants after they completed the first three phases.

4.3 Timeline and compensation

The primary study ran for two months with a total of 1,000 distinct participants.³ Participants were initially paid \$0.30 for completing the direct version and \$0.75 for the contextualized version, as the latter involved a larger workload. After completing the study, we realized we may have unintentionally under-compensated participants,⁴ since the median HIT completion time was 4.4 and 11.5 minutes for direct and contextualized versions. We therefore retroactively doubled all participants’ compensation to \$0.60 and \$1.50, which equates to approximately \$7.80 – \$8.20 per hour.

4.4 Ethical considerations

This user study was duly approved by the Institutional Review Board (IRB) of the primary authors’ organization. No sensitive or personally identifiable information was collected from participants. We used the pseudonymous MTurk worker IDs only to check that participants were unique.

Since the contextualized setting did not inform participants of the actual aim of the study beforehand, two additional documents were filed and approved by the IRB: (1) “*Use of deception/incomplete disclosure*” and (2) “*Waiver or Alteration of the Consent*”. After each participant completed the contextualized setting, we disclosed the study’s actual goal and asked whether they gave us permission to use their data. No data were collected from participants who declined.

4.5 User study implementation

The realization of the user study included a front-end webpage and a back-end server. The front-end was a single HTML page that implemented the four phases described above. To prevent any inconsistencies, participants were prevented from going

back to a previous phase or retrying a task once they had progressed. Timing events were captured with millisecond precision using the native JavaScript `Date` library. Timing events were recorded at several points for each CAPTCHA: request, serve, load, display, submit, and server response. We measured *solving time* as the time between a CAPTCHA being displayed and the participant submitting a solution, as is done in prior CAPTCHA user studies [23, 24, 27, 34, 37, 38, 43, 47, 48, 52, 58, 67, 75]. Depending on the type of CAPTCHA, this might include multiple rounds or attempts.

We used Amazon MTurk to recruit participants, host the front-end, and collect data. While most types of CAPTCHAs shown by the front-end were served from their respective providers, distorted text CAPTCHAs were not available from a third-party provider, as these are usually hosted by the websites themselves. We therefore set up our own back-end server to serve distorted text CAPTCHAs. Specifically, we downloaded a total of 1,000 unique distorted text CAPTCHAs of three different types, and stored these in a local MongoDB [19] database. We used a `Node.js` [20] server to retrieve and serve CAPTCHAs from the database. Every participant was served one text CAPTCHA of each type, and each unique text CAPTCHA was served to three different participants.

Table 2 shows the demographic information of the participants who completed the study. The demographics of the two subgroups who completed direct and contextualized studies are very similar to each other.

4.6 Potential limitations

Use of MTurk: Webb et al. [69] reported several potential concerns regarding the quality of data collected from MTurk. Of their six criteria, our study did not implement two: consent quiz (1) and examination of qualitative responses (2), which we acknowledge as a limitation. The remaining four criteria can be either evaluated through collected data or are not an issue for our study. Eligibility (3) and attention check (4) can be verified via the accuracy of text-based CAPTCHA responses, which confirm that nearly all of our participants were focused and provided correct data. Response time (5) was within our expected range. Study completion (6) was not an issue, since each participant had to complete every CAPTCHA to proceed.

Bots and farms: Similarly, Chmielewski et al. [30] reported a decrease in data quality, citing bot and farm activity.

³To the best of our knowledge, all participants were distinct. We configured Amazon MTurk to only allow unique accounts to participate.

⁴In terms of US federal minimum wage.

However, Moss and Litman [53] subsequently used several bot-detection measures to evaluate whether bots could be contaminating MTurk data, and found no evidence of bot activity. Every participant who completed our study solved ten modern CAPTCHAS, which although possible, would be more difficult for bots. Since we configured MTurk to only allow one completion per MTurk account, farm activity was also limited. Therefore, we are reasonably confident that our results are not influenced by bots or farms.

Choice of CAPTCHAS: One consequence of using the CAPTCHA types we identified in Section 3 is that our user study results are not directly comparable with those from prior CAPTCHA user studies. In general, it is difficult to directly compare such studies, as even if the same *types* of CAPTCHAS are studied, different implementations may be used e.g., reCAPTCHA and hCAPTCHA are both image-based CAPTCHAS, but could give different results.

Unmodified CAPTCHAS: In order to maximize the level of realism in our study, we used existing unmodified CAPTCHAS. We therefore did not have fine-grained control over the precise behavior of these CAPTCHAS, nor the ability to obtain more fine-grained measurements of participants' accuracy or performance beyond overall solving time. However, like previous studies, we consider overall solving time to be the most important measurable quantity.

Invalid inputs: Unfortunately, the input field for the CAPTCHA preference question in our post-study questionnaire was a free text field rather than a pull-down menu. This allowed some participants to provide preference scores outside the requested 1-5 range. We therefore excluded invalid preference scores from 163 participants.⁵

Abandonment: Since we did not record how many participants began our main study, we cannot precisely quantify the rate of abandonment. To investigate this further, we performed an additional abandonment-focused study (Section 6), where we observed a 30% abandonment rate. We can therefore assume a similar abandonment rate for our main study. Whilst the impact of this level of abandonment is unclear, it could potentially affect the ecological validity of our results, as the participants who were willing to complete the study may not be representative of all users.

Confounding factors: There were several differences between our direct and contextualized settings, some of which may be confounding factors when comparing these two groups. For example, participants in the contextualized setting had to do more work, so their attention or focus might have been reduced during CAPTCHA solving. Differences in compensation or participants' perceived benefit of completing the task (i.e., creating an account vs. solving a CAPTCHA) may have affected motivation or likeliness to abandon the task.

⁵However, we have high confidence that these participants did not provide incorrect or rushed responses during the rest of the study because their average accuracy in text-based CAPTCHAS was similar to the study-wide average. We therefore retained their measurements in other sections.

5 Results & Analysis

This section presents the user study results. Unless otherwise indicated, results are based on the full set of participants.

5.1 Solving times

This subsection addresses **RQ1: How long do human users take to solve different types of CAPTCHAS?** Figure 7 shows the distribution of solving times for each CAPTCHA type. We observed a small number of extreme outliers where the participant likely switched to another task before returning to the study. We therefore filtered out the highest 50 solving times per CAPTCHA type, out of 1,000 total.

For reCAPTCHA, the selection between image- or click-based tasks is made dynamically by Google. Whilst we know that 85% and 71% of participants (easy and hard setting) were shown a click-based CAPTCHA, the exact task-to-participant mapping is not revealed to website operators. We therefore assume that the slowest solving times correspond to image-based tasks. After disambiguation, click-based reCAPTCHA had the lowest median solving time at 3.7 seconds. Curiously, there was little difference between easy and difficult settings.

The next lowest median solving times were for distorted text CAPTCHAS. As expected, simple distorted text CAPTCHAS were solved the fastest. Masked and moving versions had very similar solving times. For hCAPTCHA, there is a clear distinction between easy and difficult settings. The latter consistently served either a harder image-based task or increased the number of rounds. However, for both hCAPTCHA settings, the fastest solving times are similar to those of reCAPTCHA and distorted text. Finally, the game-based and slider-based CAPTCHAS generally yielded higher median solving times, though some participants still solved these relatively quickly (e.g., < 10 seconds).

With the exception of reCAPTCHA (click) and distorted text, we observed that solving times for other types have a relatively high variance. Some variance is expected, especially since these results encompass all input modalities across both direct and contextualized settings. However, *relative differences in variances* indicate that, while some types of CAPTCHAS are consistently solved quickly, most have a range of solving times across the user population. The full statistical analysis of our solving time results is presented in Appendix C.

5.2 Preferences analysis

This subsection addresses **RQ2: What CAPTCHA types do users prefer?** Figure 8 shows participants' CAPTCHA preference responses after completing the solving tasks. The CAPTCHA types are sorted from most to least preferred by overall preference score, which is calculated by summing the numeric scores. Since easy and difficult settings




Figure 7: Solving times for various types of CAPTCHAS. Boxes show the middle 50% of participants, and whiskers show the filtered range. Black vertical lines show the median.

of hCAPTCHA are visually indistinguishable, we could only ask participants for one preference.

As expected, participants tend to prefer CAPTCHAS with lower solving times. For example, reCAPTCHA (click) has the lowest median solving time and the highest user preference. However, surprisingly, this trend does not seem to hold for game-based and slider-based CAPTCHAS, since these received some of the highest preference scores, even though they typically took longer than other types. This suggests that factors beyond solving time could be contributing to participants' preference scores. Notably, no single CAPTCHA type is either universally liked or disliked. For example, even the top-rated click-based reCAPTCHA, was rated 1 or 2 by 18.9% of participants. Similarly, over 31.0% rated hCAPTCHA 4 or 5, although it had the lowest overall preference score.

5.3 Direct vs. contextualized setting

This subsection addresses **RQ3: Does experimental context affect solving time?** Figure 9 shows histograms of CAPTCHA solving times for participants in the direct vs. contextualized settings. In every case except one, the mean solving time is lower in the direct setting. In most cases, the distribution from the contextualized setting has more participants with longer solving times, i.e., a longer tail.

The largest statistically significant difference is in reCAPTCHA (easy click), where the mean solving time grows by 1.8 seconds (57.5%). Second is Arkose (rotation), where it grows by 10 seconds (56.1%). Across all CAPTCHA types, the




Figure 8: Participant-reported preference scores for different types of CAPTCHAS, sorted from highest to lowest.

average increase from direct to contextualized is 26.7%. Similarly, the mean solving time for reCAPTCHA (easy image) increased by 63.6% in the contextualized setting showing the largest increase. However this was not statistically significant. This is likely due to the skew of participants in direct and contextualized versions receiving image-challenges, which is controlled by Google. Easy images were shown to 8.9% of contextualized and to 17.2% of direct setting participants, while hard images were shown to 25.5% and 30% respectively, resulting in different population sizes.

On the other hand, hCAPTCHA (difficult), which has the highest median solving time overall, showed no significant difference in mean solving time between direct and contextualized settings. This may be attributable to the difficulty of solving this type of CAPTCHA, regardless of the setting.

Results of Kruskal-Wallis tests confirm that there are statistically significant differences in mean solving times for all CAPTCHA types ($p < 0.001$) except Geetest, reCAPTCHA (image) and hCAPTCHA (difficult). While there were several potential confounding factors in our study, these results suggest that experimental context can have a significant impact on participants' CAPTCHA solving times, and must therefore be taken into account in the design of future user studies.

5.4 Effects of demographics

This subsection addresses **RQ4: Do demographics affect solving time?** We analyzed how demographic characteristics in our study correlate with CAPTCHA solving times. For some characteristics, such as education and gender, we did not observe large differences in CAPTCHA solving times (see Figures 13 and 14 in Appendix D).

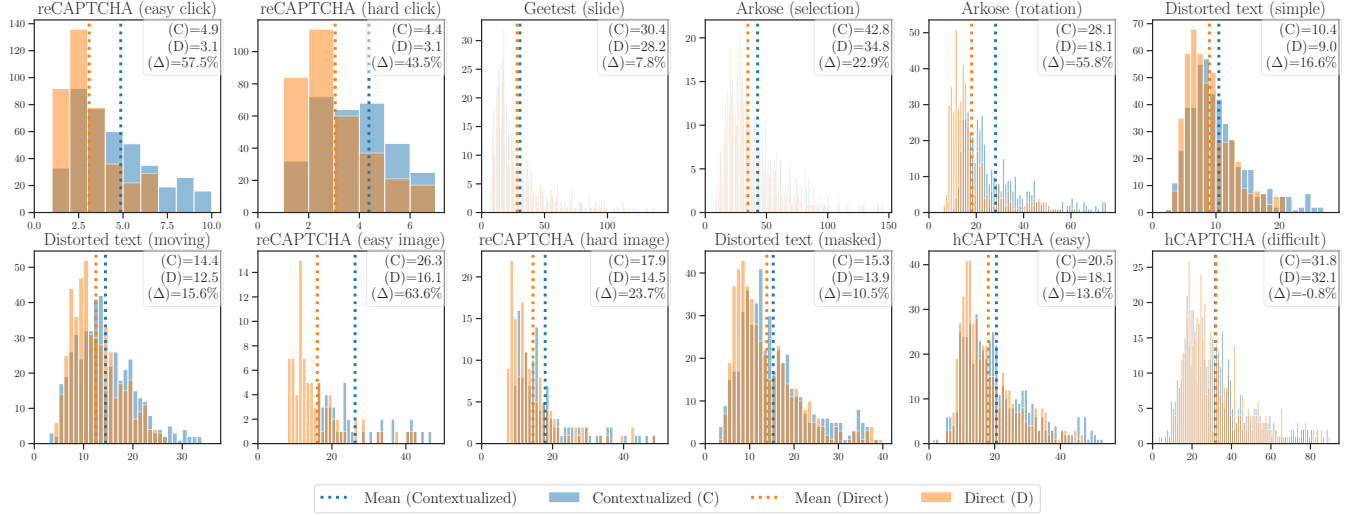


Figure 9: CAPTCHA solving times for direct (D) vs. contextualized (C) user study settings. The horizontal axis shows solving time in seconds, quantized into one-second buckets, and the vertical axis shows number of participants.

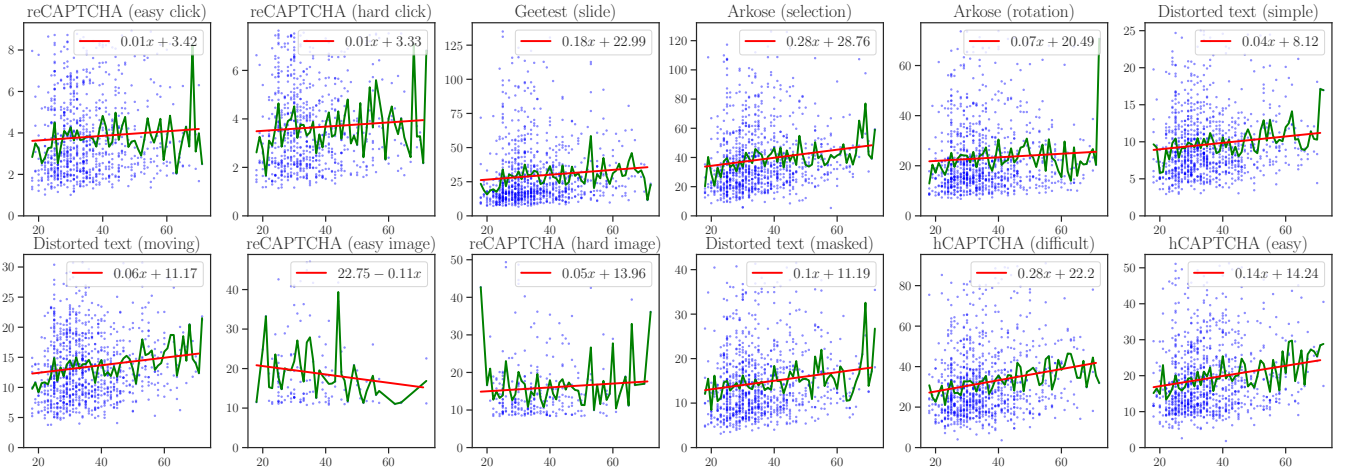


Figure 10: Effects of age in CAPTCHA solving time. The horizontal axis shows the age and the vertical axis shows the solving time. The red line shows the linear fit of the data points and the green line shows the average solving time per age.

5.4.1 Effects of age

Figure 10 shows the effect of participants' age on solving time. The green line is the average solving time for each age, and the red line is a linear fit minimizing mean square error. For all types, except reCAPTCHA (easy image), there is a trend of younger participants having lower average solving times. This agrees with prior results [27] and is especially noticeable in hCAPTCHA, Arkose (selection), and Geetest.

5.4.2 Effects of device type

Figure 11 shows the effect of device type. Although there are some differences in median between device types for certain CAPTCHA types, the Kruskal-Wallis test shows that

the differences in means are mostly not statistically significant. The only statistically significant differences are in distorted text CAPTCHAs ($p < 0.02$) and reCAPTCHA (hard click) ($p < 0.01$), where participants who used computers had a lower mean solving time compared to those using phones. Interestingly, we found a statistically significant difference between participants who used physical keyboards and those who used touch input for the simple and masked distorted text CAPTCHAs ($p < 0.02$), as well as reCAPTCHA (hard click) ($p < .001$), reCAPTCHA (easy click) ($p < .05$), and Arkose (selection) ($p < .003$). We found no statistically significant difference in mean solving times for moving distorted text.




Figure 11: Effects of device type.




Figure 12: Effects of typical Internet use.

5.4.3 Effects of typical Internet use

Figure 12 shows the relationship between participants' self-reported dominant Internet usage patterns and their CAPTCHA solving times. The Kruskal-Wallis test shows some initial evidence for statistically significant differences between participants who use the Internet primarily for work and those who use it for other purposes ($p < 0.05$). The former were typically slower than the latter in 8 out of 12 CAPTCHAs. However, some categories do not have a sufficient number of participants, thus further investigation is recommended.

5.5 Accuracy of CAPTCHAS

Table 3 contrasts our measured human solving times and accuracy against those of automated bots reported in the literature. Interestingly, these results suggest that bots *can* outperform humans, both in terms of solving time and accuracy, across all these CAPTCHA types. As mentioned in Section 4.6, our decision to use unmodified real-world CAPTCHAS means we only have accuracy results for a subset of CAPTCHA types (e.g., neither Geetest nor Arkose provide accuracy information). For the same reason, our accuracy results also include participants who only partially completed the study.

reCAPTCHA: The accuracy of image classification was 81% and 81.7% on the easy and hard settings respectively. Surprisingly, the difficulty appeared not to impact accuracy.

hCAPTCHA: The accuracy was 81.4% and 70.6% on the easy and hard settings respectively. This shows that, unlike reCAPTCHA, the difficulty has a direct impact on accuracy.

Distorted Text: We evaluated *agreement* among partici-

pants as a proxy for accuracy. As each individual CAPTCHA was served to three separate participants, we measured agreement between any two or more participants. We also observed that agreement increases dramatically (20% on average) if responses are treated as case insensitive, as shown in Table 4.

Table 3: Humans vs. bot solving time (seconds) and accuracy (percentage) for different CAPTCHA types.

CAPTCHA Type	Human		Bot	
	Time	Accuracy	Time	Accuracy
reCAPTCHA (click)	3.1-4.9	71-85%	1.4 [63]	100% [63]
Geetest	28-30	N/A	5.3 [70]	96% [70]
Arkose	18-42	N/A	N/A	N/A
Distorted Text	9-15.3	50-84%	<1 [77]	99.8% [39]
reCAPTCHA (image)	15-26	81%	17.5 [45]	85% [45]
hCAPTCHA	18-32	71-81%	14.9 [44]	98% [44]

Table 4: Agreement for distorted text CAPTCHAS.

	Average Agreement	Average Agreement (case insensitive)
Simple	84%	93%
Masked	50%	73%
Moving	62%	90%
Total	65%	85%

6 Measuring User Abandonment

This subsection addresses **RQ5: Does experimental context influence abandonment?** Upon completion, we observed that the number of CAPTCHAS solved during our study exceeded what would be expected based on the number of participants who completed the study. We hypothesized that this was due to participants starting but not completing the study. To measure this behavior, we conducted a second user study that collected timestamps between CAPTCHAS, regardless of whether the entire study was completed. We measured: (1) how many participants started the task; (2) how many abandoned the task when solving a CAPTCHA; and (3) if so, at which task and CAPTCHA.

This abandonment-focused study consisted of four groups, each with 100 unique participants. Two groups were presented with the direct setting and the other two with the contextualized setting (see Section 4.2). We hypothesized that the amount of compensation might also impact abandonment, so we doubled the compensation for one of the groups in each setting. The studies were run sequentially to avoid prospective participants simply picking the higher-paying study.

We summarize the key findings below, and present the full results in Tables 7, 8, 9, and 10 in Appendix A. Out of a total of 574 participants who started the study, 174 abandoned prior to completion (i.e., 30% abandonment rate). Several observations can be made: First, in the direct setting, 25% of the participants who ultimately abandoned the study did so before solving the first CAPTCHA, but this rose to nearly 50% in the contextualized setting. Second, doubling the pay halved the abandonment rate for the contextualized setting (as expected), but increased it by 50% in the direct setting. Third, participants in the contextualized setting were 120% more likely to abandon than those in the direct setting. Fourth, in the contextualized setting, participants at the higher compensation level solved CAPTCHAS faster than those at the lower compensation level (21.5% decrease in average solving time across all CAPTCHA types). Interestingly, in the direct setting, participants at the higher compensation level solved CAPTCHAS *slower* than those at the lower compensation level (27.4% *increase* in average solving time across all CAPTCHA types). Finally, some CAPTCHA types (e.g., Geetest) exhibited higher rates of abandonment than others.

This initial investigation strongly motivates the need for further exploration of CAPTCHA-induced abandonment. Although we studied the impact of compensation and experimental context, there may be other reasons behind abandonment, such as: CAPTCHA type, CAPTCHA difficulty, and expected duration of study. Nevertheless, the trend of average users' unwillingness to solve a CAPTCHA during account creation (even for monetary compensation) is a relevant finding for websites that choose to protect account creation (and/or account access) using CAPTCHAS.

7 Related Work

CAPTCHAS are a well-studied topic, with several prior studies investigating both existing and novel CAPTCHA schemes.

7.1 Comparison of methodologies

Table 5 summarizes the key methodological aspects of prior CAPTCHA user studies, from which the following observations can be made:

- Most prior research has focussed on distorted text and newly-proposed CAPTCHA schemes.
- MTurk and proprietary websites have been the norm across CAPTCHA user studies (except DevilTyper [43]).
- Whilst almost all studies measured solving time, there is a bifurcation in terms of accuracy measurements: studies evaluating their own CAPTCHA schemes or reimplementing existing schemes typically have direct access to accuracy results, whereas those evaluating unmodified deployed CAPTCHAS can only measure quantities such as agreement.
- Most studies measured demographics and ratings or preferences. Some studies also measured workload, open response (perceptions), and perceived usability.

7.2 Detailed comparisons

We present detailed comparisons of our methodology and results with three representative prior CAPTCHA studies.

Bursztein et al. [27] presented the first large-scale study on human CAPTCHA solving performance. Focussing on distorted text and audio CAPTCHAS, they used both MTurk and an underground CAPTCHA-solving service to measure solving time and accuracy. In terms of solving times, they found that it took on average 9.8 and 28.4 seconds to solve distorted text and audio CAPTCHAS respectively. Although we did not evaluate audio CAPTCHAS (as we did not observe these in our website inspection), our results for distorted text CAPTCHAS broadly agree at 12.5 on average. Similarly to our study, they used *agreement* between participants as a proxy for accuracy. For distorted text, they observed 71% agreement, which is in line with our observation of 75% when averaging case sensitive and insensitive versions (see Table 4).

Feng et al. [33] presented senCAPTCHA, a new CAPTCHA type using orientation sensors designed specifically for mobile devices with small screens. They evaluated its security against brute-force and ML-based attacks, and its usability through two usability studies totalling 472 participants. The second user study compared senCAPTCHA against text-, audio-, image-, and video-based CAPTCHAS, some of which were reimplemented for the study. senCAPTCHA had the lowest median solving time (5.02 seconds), followed by image (9.6), video (10.08), text (11.93), and audio (47.07). With the exception of click-based reCAPTCHA, it can be extrapolated that

Table 5: Methodology and details of previous CAPTCHA-related user studies.

	CAPTCHA types	Delivery medium	Measurements	Survey methods	CAPTCHA source	Compensation (USD per # CAPTCHAS)
Ours	Text, Image, Game, Slider, Behavior	MTurk	Time, Agreement, Accuracy, Abandonment, Context	Demographics, Preference	Alexa	\$0.30-\$1.50 per 10
[27]	Text, Audio	MTurk, Website	Time, Agreement	Demographics	Alexa	\$0.02-\$0.50 per 24-39
[52]	DCG Captcha	MTurk	Time, Accuracy	Demographics, SUS	Newly proposed	\$0.50 per 4
[47]	reCAPGen Audio	MTurk	Time, Accuracy	Demographics, Rating/Preference	Newly proposed	\$4.00 per 60
[38]	3D/2D Text	MTurk	Time, Accuracy	Demographics, SUS	[51, 72, 73]	\$1.00 per 30
[43]	Text	MTurk, DevilTyper	Time, Accuracy, Abandonment	None	Major websites	\$0.03 per 15 (MTurk), 30.00 per 1.4 mil
[24]	Text, Audio, Interface	Website	Time, Accuracy	Demographics, Preference	Alexa	None
[34]	Text	Website	None	Demographics, Rating/Preference	Newly proposed	None
[37]	Jigsaw puzzle	Website	Time, Accuracy	Demographics, Preference	Newly proposed	None
[48]	Text, Game, NoBot	Website	Time	Workload, Perceptions, Preference	None	None
[33]	SenCAPTCHA, Text, Image, Audio, Video	MTurk	Time	Demographics, Preference, SUS	Newly proposed, [55]	\$1.25 per 9-15
[66]	Text, Behavior, Invisible, Game, Math	Unknown	Time	Demographics, Preference	None	None
[58]	Sketcha	MTurk	Time, Accuracy	Demographics	Newly proposed	\$0.05-\$0.30 per 10-12

senCAPTCHA would have a lower solving time than the other CAPTCHA types in our study. In terms of preferences, most participants in their study preferred senCAPTCHA. Out of the CAPTCHA types in our study, senCAPTCHA most closely resembles the game-based CAPTCHAS, which supports our finding that game-based CAPTCHAS are generally preferred over text and image-based CAPTCHAS (see Figure 8).

Tanthavech and Nimkoompai [66] performed a 40-participant user study, measuring solving time for five CAPTCHA types: click-based reCAPTCHA, text-, game-, math-based, and a newly-proposed invisible CAPTCHA, which is essentially a honeypot for bots. In terms of solving times, their distorted text measurement (12 seconds) is in the middle of our observed range (9-15 seconds), which is expected since it closely resembles our *masked* type of distorted text. Similarly, their click-based reCAPTCHA measurement (3.1 seconds) is on the boundary of our range (3.1-4.9), which suggests they may have configured the “easier for users” setting. Their game-based CAPTCHA appears to have a lower solving time than ours, but this is likely due to the type of game. We did not observe or evaluate any math-based CAPTCHAS. They also asked participants several post-study questions about the five CAPTCHA types. Interestingly, their participants “enjoyed” the game-based CAPTCHA more than reCAPTCHA (click), which is the inverse of our findings (see Figure 8), but may again be due to the different types of game.

Overall, where our study measured similar quantities to prior work, our findings broadly agree. However, there is still

a high degree of diversity in the sets of quantities measured in each study (e.g., types of CAPTCHAS, effect of experimental context), suggesting that a plurality of studies are needed to understand the full CAPTCHA landscape.

7.3 Summarized comparisons

In addition, Table 6 presents a summarized comparison of our results with those of other prior studies.

Solving Time: Overall, the average solving time in our study ranged from 3.6 to 42.7 seconds per CAPTCHA, which is a larger range than that observed by Bursztein et al. [27] in 2010 (9.8 – 28.4 seconds) but is similar to the 2019 study by Feng et al. [33] (medians ranging from 5.0 to 47.1 seconds). Although direct comparison of solving times is not always meaningful, even for the same CAPTCHA type (e.g., due to differing implementations or difficulty settings), we can identify a few trends. Firstly, our measured solving times for the three types of distorted text CAPTCHAS (9-15 seconds⁶) are within the range of observations from prior studies (6-20 seconds). We can therefore use this as a reference point for comparisons. Secondly, with the exception of behavior-based CAPTCHAS, we observed that all other CAPTCHA types took longer than distorted text. Without considering newly-proposed CAPTCHA types, this trend is consistent across most prior studies (with the exception of [33] and [66]). Thirdly, although we do not

⁶Unless otherwise stated, measurements refer to average solving time.

Table 6: Comparison of results from prior user studies evaluating CAPTCHAS: audio (A), behavior (B), distorted text (DT), game (G), honeypot (HP), image (I), math (M), service (S), slider (SL), video (V) and newly-proposed (New). Some studies used non-unique (NU) participants or MTurk (MT). * denotes reimplemented CAPTCHA types.

	Unique users	CAPTCHAS solved	Average solving time (seconds)	Average accuracy
Ours	1,400 (MT)	14,000	9-15 (DT), 15-32 (I), 18-42 (G), 29 (SL), 3.1-4.9 (B)	50-84% (DT), 71-81% (I), 71-85% (B)
[27]	1,100-11,800 (MT)	318,000	9.8 (DT), 28.4 (A), 22.4 (S)	71% (DT), 31% (A), 93% (ebay DT)
[52]	120	480	8.5-16 (New), 17-47 (Attacks)	16-100% (New)
[47]	79	4,740	9.6 (New)	78.2% (New)
[38]	120	3,600	10 (3D-DT), 6.2-6.7 (DT)	84% (3D-DT), 92-96% (DT)
[43]	5,000 (NU), 44 (MT)	1.4 mil, 7,500	8.5-12 (DT)	79%-89% (DT)
[24]	162, 14 (Interface)	2,350	9.9 (DT), 50.9 (Blind DT), 22.8 (A)	80% (DT), 39-43% (A)
[34]	210	210	None	None
[37]	100	300	4.9-6.4 (New)	78%-87.5% (New)
[48]	87	261	20 (DT), 29 (G), 70 (NoBot)	None
[33]	436	4,920	12 (DT), 47 (A), 9.6 (I*), 5 (New), 12 (V*)	None
[66]	40	200	12 (DT), 0 (HP), 3.1 (B), 8.2 (G [76]), 4.1 (M [42])	None
[58]	558 (NU)	14,302	35 (New)	42%-88% (New)

evaluate any newly-proposed CAPTCHA types, the times reported for these by other studies are typically faster than most of the CAPTCHA types in our study, suggesting that there is scope for developing new CAPTCHA types with lower solving times. Finally, even in comparison to newly-proposed schemes, the behavior-based CAPTCHAS (e.g. reCAPTCHA click) appear to have the lowest solving times overall.

Accuracy: For the case-sensitive setting, we observed a relatively broad range of accuracy (i.e., agreement) measurements for distorted text (50-84%). However, in the case-insensitive setting, our accuracy range narrows to 73-93%, which more closely aligns with prior studies, which have reported distorted text accuracies in the range 71-96%. This suggests that both participants and prior studies have focussed on the case-insensitive setting. In terms of deployed CAPTCHAS, [27] reported an accuracy of 93% for distorted text CAPTCHAS used by EBay in 2010. This is higher than for the image-based CAPTCHAS we measured (71-81%), suggesting that the latter may have increased in difficulty.

Security: Table 3 shows a comparison of our results to prior security analyses. Automated attacks on various CAPTCHA schemes have been quite successful [21, 22, 25, 26, 29, 32, 36, 39, 44, 45, 49, 50, 59, 61, 63–65, 70, 77]. The bots’ accuracy ranges from 85-100%, with the majority above 96%. This substantially exceeds the human accuracy range we observed (50-85%). Furthermore the bots’ solving times are significantly lower in all cases, except reCAPTCHA (image), where human solving time (18 seconds) is similar to the bots’ (17.5 seconds). However, in the contextualized setting, human solving time rises to 22 seconds, indicating that in this more natural setting, humans are slightly slower than bots.

8 Summary & Future Work

This paper explores currently-deployed CAPTCHAS via inspection of 200 popular websites and a series of user studies totalling 1,400-participants. For the research questions we posed at the outset, our results:

RQ1: show that there are significant differences in mean solving times between CAPTCHA types.

RQ2: show that users’ preference is not fully correlated with CAPTCHA solving time.

RQ3: show that experimental context significantly influences CAPTCHA solving times.

RQ4: confirm the previously-reported effects of age on solving time.

RQ5: confirm the high rates of abandonment due to CAPTCHA-related tasks and identify that experimental context impacts abandonment.

We anticipate several directions for future work, including obtaining detailed measurements through a controlled user study, and further investigating the causes of abandonment.

9 Acknowledgements

We thank the anonymous reviewers for their valuable comments, and we are especially grateful to the shepherd for guiding us through several revisions. The work of UCI authors was supported in part by: NSF Award #:1840197, NSF Award #:1956393, and NCAE-C CCR 2020 Award #:H98230-20-1-0345. Yoshimichi Nakatsuka was supported in part by The Nakajima Foundation. Andrew Paverd was supported in part by a US-UK Fulbright Cyber Security Scholar Award.

References

- [1] 360.cn. <https://passport.360.cn/>.
- [2] Alexa Top Sites. <https://www.alexa.com/topsites>.
- [3] Amazon Mechanical Turk. <https://www.mturk.com/>.
- [4] Arkose Labs. <https://www.arkoselabs.com/about-us/>.
- [5] CAPTCHA Usage Distribution on the Entire Internet. <https://trends.builtwith.com/widgets/captcha/traffic/Entire-Internet>.
- [6] Cisco Umbrella 1 Million. <https://umbrella.cisco.com/blog/cisco-umbrella-1-million>.
- [7] Cloudflare Radar - Domain Rankings. <https://radar.cloudflare.com/domains>.
- [8] Geetest CAPTCHA. <https://www.geetest.com/en/Captcha>.
- [9] hCaptcha. <https://www.hcaptcha.com/>.
- [10] hCaptcha Is Now The Largest Independent CAPTCHA Service, Runs on 15% Of The Internet. <https://www.hcaptcha.com/post/hcaptcha-now-the-largest-independent-captcha-service>.
- [11] Invisible reCAPTCHA. <https://developers.google.com/recaptcha/docs/invisible>.
- [12] jrj.com. <https://sso.jrj.com/>.
- [13] NuData Security. <https://nudatasecurity.com/>.
- [14] reCAPTCHA. <https://www.google.com/recaptcha/about/>.
- [15] reCAPTCHA v3. <https://developers.google.com/recaptcha/docs/v3>.
- [16] The Majestic Million. <https://majestic.com/reports/majestic-million>.
- [17] The Tor Project: Privacy & Freedom Online. <https://www.torproject.org/>.
- [18] Xinhuanet. <https://mail.xinhuanet.com>.
- [19] MongoDB. <https://www.mongodb.com/>, 2021.
- [20] Node.js. <https://nodejs.org/>, 2021.
- [21] W. Aiken and H. Kim. POSTER: DeepCRACK: Using Deep Learning to Automatically CRack Audio CAPTCHAs. In *Proceedings of the 2018 on Asia Conference on Computer and Communications Security*, ASIACCS '18, page 797–799, New York, NY, USA, 2018. ACM.
- [22] F. H. Alqahtani and F. A. Alsulaiman. Is image-based CAPTCHA secure against attacks based on machine learning? An experimental study. *Computers & Security*, 88:101635, 2020.
- [23] M. Belk, P. Germanakos, C. Fidas, A. Holzinger, and G. Samaras. Towards the Personalization of CAPTCHA Mechanisms Based on Individual Differences in Cognitive Processing. In A. Holzinger, M. Ziefle, M. Hitz, and M. Debevc, editors, *Human Factors in Computing and Informatics*, pages 409–426, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [24] J. P. Bigham and A. Cavender. Evaluating Existing Audio CAPTCHAs and an Interface Optimized for Non-Visual Use. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '09, page 1829–1838, New York, NY, USA, 2009. ACM.
- [25] K. Bock, D. Patel, G. Hughey, and D. Levin. unCaptcha: A Low-Resource Defeat of reCaptcha's Audio Challenge. In *11th USENIX Workshop on Offensive Technologies (WOOT 17)*, Vancouver, BC, Aug. 2017. USENIX Association.
- [26] E. Bursztein, R. Beauxis, H. Paskov, D. Perito, C. Fabry, and J. Mitchell. The Failure of Noise-Based Non-continuous Audio Captchas. In *2011 IEEE Symposium on Security and Privacy*, pages 19–31, 2011.
- [27] E. Bursztein, S. Bethard, C. Fabry, J. C. Mitchell, and D. Jurafsky. How Good Are Humans at Solving CAPTCHAs? A Large Scale Evaluation. In *IEEE Symposium on Security and Privacy*, 2010.
- [28] V. G. Cerf. Guidelines for Internet Measurement Activities. RFC 1262, Oct. 1991.
- [29] J. Chen, X. Luo, Y. Guo, Y. Zhang, and D. Gong. A Survey on Breaking Technique of Text-Based CAPTCHA. *Security and Communication Networks*, 12 2017.
- [30] M. Chmielewski and S. C. Kucker. An MTurk crisis? Shifts in data quality and the impact on study results. *Social Psychological and Personality Science*, 11(4):464–473, 2020.
- [31] F. Consulting. State of online fraud and bot management. https://services.google.com/fh/files/misc/google_forrester_bot_management_tlp_post_production_final.pdf, 2021.
- [32] M. Darnstädt, H. Meutzner, and D. Kolossa. Reducing the Cost of Breaking Audio CAPTCHAs by Active and Semi-supervised Learning. In *2014 13th International Conference on Machine Learning and Applications*, pages 67–73, 2014.
- [33] Y. Feng, Q. Cao, H. Qi, and S. Ruoti. Sencaptcha: A mobile-first captcha using orientation sensors. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 4(2), jun 2020.
- [34] C. A. Fidas, A. G. Voyiatzis, and N. M. Avouris. On the Necessity of User-Friendly CAPTCHA. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, page 2623–2626, New York, NY, USA, 2011. ACM.
- [35] H. Gao, W. Wang, and Y. Fan. Divide and conquer: an efficient attack on Yahoo! CAPTCHA. In *2012 IEEE 11th International Conference on Trust, Security and Privacy in Computing and Communications*, pages 9–16. IEEE, 2012.
- [36] H. Gao, J. Yan, F. Cao, Z. Zhang, L. Lei, M. Tang, P. Zhang, X. Zhou, X. Wang, and J. Li. A Simple Generic Attack on Text Captchas. In *Network and Distributed System Security Symposium (NDSS)*, San Diego, California, United States, 2016.
- [37] H. Gao, D. Yao, H. Liu, X. Liu, and L. Wang. A Novel Image Based CAPTCHA Using Jigsaw Puzzle. In *2010 13th IEEE International Conference on Computational Science and Engineering*, pages 351–356, 2010.
- [38] S. Gao, M. Mohamed, N. Saxena, and C. Zhang. Emerging-Image Motion CAPTCHAs: Vulnerabilities of Existing Designs, and Countermeasures. *IEEE Transactions on Dependable and Secure Computing*, 16(6):1040–1053, 2019.
- [39] I. J. Goodfellow, Y. Bulatov, J. Ibarz, S. Arnaud, and V. Shet. Multi-digit number recognition from street view imagery using deep convolutional neural networks. *arXiv preprint arXiv:1312.6082*, 2014.
- [40] M. Guerar, L. Verderame, M. Migliardi, F. Palmieri, and A. Merlo. Gotta CAPTCHA 'Em All: A Survey of Twenty years of the Human-or-Computer Dilemma. *CoRR*, abs/2103.01748, 2021.
- [41] C. J. Hernandez-Castro and A. Ribagorda. Pitfalls in CAPTCHA design and implementation: The Math CAPTCHA, a case study. *Computers & Security*, 29(1):141–157, 2010.
- [42] C. J. Hernandez-Castro and A. Ribagorda. Pitfalls in captcha design and implementation: The math captcha, a case study. *Computers & Security*, 29(1):141–157, 2010.
- [43] C.-J. Ho, C.-C. Wu, K.-T. Chen, and C.-L. Lei. DevilTyper: A Game for CAPTCHA Usability Evaluation. *Comput. Entertain.*, 9(1), apr 2011.
- [44] M. I. Hossen and X. Hei. A Low-Cost Attack against the hCaptcha System. *CoRR*, abs/2104.04683, 2021.
- [45] M. I. Hossen, Y. Tu, M. F. Rabby, M. N. Islam, H. Cao, and X. Hei. An Object Detection based Solver for Google's Image reCAPTCHA v2. In *23rd International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2020)*, pages 269–284, San Sebastian, Oct. 2020. USENIX Association.
- [46] Imperva. Imperva bad bot report. <https://www.imperva.com/resources/resource-library/reports/bad-bot-report/>, 2022.

- [47] M. Jain, R. Tripathi, I. Bhansali, and P. Kumar. Automatic Generation and Evaluation of Usable and Secure Audio ReCAPTCHA. In *The 21st International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '19, page 355–366, New York, NY, USA, 2019. Association for Computing Machinery.
- [48] K. Krol, S. Parkin, and M. A. Sasse. Better the Devil You Know: A User Study of Two CAPTCHAs and a Possible Replacement Technology. In *2016 NDSS Workshop on Usable Security*, pages 1–10, 2016.
- [49] C. Li, X. Chen, H. Wang, P. Wang, Y. Zhang, and W. Wang. End-to-end attack on text-based CAPTCHAs based on cycle-consistent generative adversarial network. *Neurocomputing*, 433:223–236, 2021.
- [50] D. Lorenzi, J. Vaidya, E. Uzun, S. Sural, and V. Atluri. Attacking Image Based CAPTCHAs Using Image Recognition Techniques. In V. Venkatakrishnan and D. Goswami, editors, *Information Systems Security*, pages 327–342, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- [51] N. Mitra, H. Chu, T. Lee, L. Wolf, H. Yeshurun, and D. Cohen-Or. Emerging images. In *Proceedings of ACM SIGGRAPH Asia 2009, SIGGRAPH Asia '09*, volume 28, pages 163:1–163:8, 2009. ACM SIGGRAPH Asia 2009, SIGGRAPH Asia '09 ; Conference date: 16-12-2009 Through 19-12-2009.
- [52] M. Mohamed, S. Gao, N. Saxena, and C. Zhang. Dynamic Cognitive Game CAPTCHA Usability and Detection of Streaming-Based Farming. In *2014 NDSS Workshop on Usable Security*, pages 1–10, 2014.
- [53] A. Moss. After the bot scare: Understanding what's been happening with data collection on MTurk and how to stop it. <https://www.clou dresearch.com/resources/blog/after-the-bot-scare-under-standing-whats-been-happening-with-data-collection-on-mturk-and-how-to-stop-it/>, Aug 2020.
- [54] M. Motoyama, K. Levchenko, C. Kanich, D. McCoy, G. M. Voelker, and S. Savage. Re: CAPTCHAs—Understanding CAPTCHA-Solving Services in an Economic Context. In *19th USENIX Security Symposium (USENIX Security 10)*, Washington, DC, aug 2010. USENIX Association.
- [55] D. Phillips. Secureimage: PHP CAPTCHA script. <https://www.php pcaptcha.org/>, 2023.
- [56] V. L. Pochat, T. van Goethem, S. Tajalizadehkhoob, M. Korczynski, and W. Joosen. Tranco: A research-oriented top sites ranking hardened against manipulation. In *26th Annual Network and Distributed System Security Symposium, NDSS 2019, San Diego, California, USA, February 24-27, 2019*. The Internet Society, 2019.
- [57] M. Prince and S. Isasi. Moving from reCAPTCHA to hCaptcha. <https://blog.cloudflare.com/moving-from-recaptcha-to-hcaptcha/>.
- [58] S. A. Ross, J. A. Halderman, and A. Finkelstein. Sketcha: A Captcha Based on Line Drawings of 3D Models. In *Proceedings of the 19th International Conference on World Wide Web*, page 821–830, New York, NY, USA, 2010. ACM.
- [59] S. Sano, T. Otsuka, and H. G. Okuno. Solving Google's Continuous Audio CAPTCHA with HMM-Based Automatic Speech Recognition. In K. Sakiyama and M. Terada, editors, *Advances in Information and Computer Security*, pages 36–52, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [60] Q. Scheitle, O. Hohlfeld, J. Gamba, J. Jelten, T. Zimmermann, S. D. Strowes, and N. Vallina-Rodriguez. A long way to the top: Significance, structure, and stability of internet top lists. In *Proceedings of the Internet Measurement Conference 2018, IMC '18*, page 478–493, New York, NY, USA, 2018. ACM.
- [61] H. Shekhar. Breaking Audio Captcha using Machine Learning/Deep Learning and Related Defense Mechanism. *San Jose State University Master's Projects*, 2019.
- [62] V. Shet. Street View and reCAPTCHA technology just got smarter. <https://security.googleblog.com/2014/04/street-view-and-recaptcha-technology.html>, 2014.
- [63] S. Sivakorn, I. Polakis, and A. D. Keromytis. I am Robot: (Deep) Learning to Break Semantic Image CAPTCHAs. In *2016 IEEE European Symposium on Security and Privacy (EuroSP)*, pages 388–403, 2016.
- [64] S. Solanki, G. Krishnan, V. Sampath, and J. Polakis. In (Cyber)Space Bots Can Hear You Speak: Breaking Audio CAPTCHAs Using OTS Speech Recognition, page 69–80. ACM, New York, NY, USA, 2017.
- [65] M. Tang, H. Gao, Y. Zhang, Y. Liu, P. Zhang, and P. Wang. Research on Deep Learning Techniques in Breaking Text-Based Captchas and Designing Image-Based Captcha. *IEEE Transactions on Information Forensics and Security*, 13(10):2522–2537, 2018.
- [66] N. Tanthavech and A. Nimkoompai. Captcha: Impact of website security on user experience. *ICIIT '19: Proceedings of the 2019 4th International Conference on Intelligent Information Technology*, pages 37–41, 02 2019.
- [67] E. Uzun, S. Chung, I. Essa, and W. Lee. rtCaptcha: A Real-Time Captcha Based Liveness Detection System. In *Network and Distributed System Security Symposium (NDSS)*, San Diego, California, United States, 02 2018.
- [68] L. von Ahn, M. Blum, N. J. Hopper, and J. Langford. CAPTCHA: Using Hard AI Problems for Security. In E. Biham, editor, *Advances in Cryptology — EUROCRYPT 2003*, pages 294–311, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.
- [69] M. A. Webb and J. P. Tangney. Too good to be true: Bots and bad data from mechanical turk. *Perspectives on Psychological Science*, 2022.
- [70] H. Weng, B. Zhao, S. Ji, J. Chen, T. Wang, Q. He, and R. Beyah. Towards understanding the security of modern image captchas and underground captcha-solving services. *Big Data Mining and Analytics*, 2(2):118–144, 2019.
- [71] Q. Xie, S. Tang, X. Zheng, Q. Lin, B. Liu, H. Duan, and F. Li. Building an open, robust, and stable voting-based domain top list. In K. R. B. Butler and K. Thomas, editors, *31st USENIX Security Symposium, USENIX Security 2022, Boston, MA, USA, August 10-12, 2022*, pages 625–642. USENIX Association, 2022.
- [72] Y. Xu, G. Reynaga, S. Chiasson, J.-M. Frahm, F. Monrose, and P. Van Oorschot. Security and usability challenges of moving-object captchas: Decoding codewords in motion. In *Proceedings of the 21st USENIX Conference on Security Symposium, Security'12*, page 4, USA, 2012. USENIX Association.
- [73] Y. Xu, G. Reynaga, S. Chiasson, J.-M. Frahm, F. Monrose, and P. C. van Oorschot. Security analysis and related usability of motion-based captchas: Decoding codewords in motion. *IEEE Transactions on Dependable and Secure Computing*, 11(5):480–493, 2014.
- [74] J. Yan and A. S. El Ahmad. A Low-cost Attack on a Microsoft CAPTCHA. In *Proceedings of the 15th ACM conference on Computer and communications security*, pages 543–554, 2008.
- [75] J. Yan and A. S. El Ahmad. Usability of captchas or usability issues in captcha design. In *Proceedings of the 4th Symposium on Usable Privacy and Security, SOUPS '08*, page 44–52, New York, NY, USA, 2008. ACM.
- [76] H. Yu and M. O. Riedl. Automatic generation of game-based captchas. In *Proceedings of the FDG workshop on Procedural Content Generation*, 2015.
- [77] Y. Zi, H. Gao, Z. Cheng, and Y. Liu. An End-to-End Attack on Text CAPTCHAs. *IEEE Transactions on Information Forensics and Security*, 15:753–766, 2020.

A Abandonment measurement

Tables 7, 8, 9, and 10 show the results from four groups of participants from the secondary study which aimed to measure abandonment. Columns represent the order of CAPTCHAs shown, while rows represent the CAPTCHA type. Cell values represent the number of MTurkers who abandoned.

Table 7: Abandonment in contextualized setting (\$0.75 payment)

	1	2	3	4	5	6	7	8	9	10	Total
reCAPTCHA (easy)	5	0	0	0	2	0	0	0	0	0	7
Geetest (slide)	3	1	2	1	3	0	0	1	1	1	13
Arkose (selection)	8	2	0	1	1	0	0	0	0	0	12
Arkose (rotation)	2	1	1	0	1	1	0	0	0	0	6
Distorted text (simple)	2	1	0	0	0	2	1	0	0	0	6
Distorted text (moving)	0	1	2	1	1	0	1	0	1	0	7
reCAPTCHA (difficult)	5	0	1	1	0	0	0	0	0	0	7
Distorted text (masked)	4	2	1	0	0	0	0	0	0	0	7
hCAPTCHA (easy)	2	2	2	0	1	0	0	0	0	0	7
hCAPTCHA (difficult)	4	1	2	1	0	0	1	0	0	0	9
Total	35	11	11	5	9	3	3	1	2	1	81

Table 8: Abandonment in contextualized setting (\$1.50 payment)

	1	2	3	4	5	6	7	8	9	10	Total
reCAPTCHA (easy)	2	1	0	0	0	0	0	0	0	0	3
Geetest (slide)	4	0	0	0	0	1	0	1	2	0	8
Arkose (selection)	1	2	0	0	0	1	0	0	0	0	4
Arkose (rotation)	4	0	1	0	0	0	0	1	0	0	6
Distorted text (simple)	2	0	0	1	0	0	0	0	0	0	3
Distorted text (moving)	1	1	1	0	0	1	1	0	0	0	5
reCAPTCHA (difficult)	2	1	0	0	0	0	0	0	0	0	3
Distorted text (masked)	1	2	0	0	0	0	0	0	0	0	3
hCAPTCHA (easy)	1	1	0	0	0	0	0	0	0	0	2
hCAPTCHA (difficult)	0	0	1	0	0	0	0	0	1	0	2
Total	18	8	3	1	0	3	1	2	3	0	39

Table 9: Abandonment in direct setting (\$0.30 payment)

	1	2	3	4	5	6	7	8	9	10	Total
reCAPTCHA (easy)	0	0	0	1	0	0	0	0	0	0	1
Geetest (slide)	1	1	0	0	1	0	1	2	0	0	6
Arkose (selection)	2	1	1	0	0	1	0	0	0	0	5
Arkose (rotation)	0	0	0	0	0	1	0	0	0	0	1
Distorted text (simple)	0	0	0	0	0	0	0	0	0	0	0
Distorted text (moving)	0	0	0	0	0	1	1	0	0	0	2
reCAPTCHA (difficult)	0	0	0	1	0	0	0	0	0	0	1
Distorted text (masked)	0	0	0	0	0	0	1	0	0	0	1
hCAPTCHA (easy)	1	1	0	1	0	0	0	0	0	0	3
hCAPTCHA (difficult)	1	0	0	1	0	0	0	0	0	0	2
Total	5	3	1	4	1	3	3	2	0	0	22

Table 10: Abandonment in direct setting (\$0.60 payment)

	1	2	3	4	5	6	7	8	9	10	Total
reCAPTCHA (easy)	0	0	0	0	0	0	0	0	0	0	0
Geetest (slide)	4	3	2	0	3	5	0	0	2	0	19
Arkose (selection)	0	0	1	0	0	0	0	0	0	0	1
Arkose (rotation)	1	0	0	2	1	0	0	0	0	0	4
Distorted text (simple)	0	0	0	0	0	0	0	0	0	0	0
Distorted text (moving)	1	0	0	0	0	0	0	0	1	0	2
reCAPTCHA (difficult)	0	0	0	0	0	0	0	0	1	0	1
Distorted text (masked)	2	0	0	0	0	0	0	0	0	0	2
hCAPTCHA (easy)	0	1	0	1	0	0	0	0	0	0	2
hCAPTCHA (difficult)	0	0	0	0	1	0	0	0	0	0	1
Total	8	4	3	3	5	5	0	0	4	0	32

B Questions asked in User Study

Table 11 shows the exact questions that were asked to the participants during the pre- and post-study questionnaire.

Table 11: Questions in user study

Question	Possible Answers
<i>Pre-study questions</i>	
Age	18 - 100
Gender	Male, Female, Non-binary
What is your country of residence?	[selected from list of countries]
What is your highest level of Education?	No formal education, High School, Associate, Bachelor's, Master's, Doctorate
Which of the following most closely describes the majority of your Internet use?	Work, Education, Browsing the Web, Gaming, Other
Which device type are you using for this survey?	Phone, Computer (Desktop / Laptop), Tablet
Which input method are you using for this survey?	Touchscreen, Keyboard, Other
[Only in the direct setting:] Are you familiar with the purpose of CAPTCHAs?	Yes, No
<i>Post-study question</i>	
On a scale of 1-5, how enjoyable was solving the following CAPTCHA types? (1 being the least, and 5 – the most, enjoyable). If the CAPTCHA type wasn't shown to you please put a 0 in that place. Note: You may not have seen the exact images shown, they are templates designed to represent different CAPTCHA types.	[single digit]

C Statistical Analysis of Solving Times

To confirm the validity of our conclusions, we conducted several standard tests on the measured solving times. We used the Holm-Bonferroni method to adjust for family-wise error in our statistical tests.

- First, we performed the *Shapiro-Wilk normality test* with a null hypothesis that solving times adhere to a normal distribution. For all CAPTCHA types, results showed that we can reject the null hypothesis ($p < 0.001$).
- Second, we ran a *skewness test* with a null hypothesis that the skewness of the sample population is the same as that of a corresponding normal distribution. For all CAPTCHA types, results allowed us to reject the null hypothesis in favor of the alternative: the distribution of solving times is skewed ($p < 0.001$).
- Third, we used the *tailedness test* with a null hypothesis that the kurtosis of the sample population is the same as that of a normal distribution. Results showed that, for all except distorted text (moving), the samples were drawn from a population that has a heavy-tailed distribution ($p < 0.001$).

Since solving times are: (1) not normally distributed, and (2) heavy tailed, we selected the *Brown Forsythe test* to compare the equality of variance between different types of CAPTCHAS. Results show that these distributions do not have equal variance, thus confirming our observations in Section 5.1. Given the result of the Brown Forsythe test, we selected the *Kruskal-Wallis test* to test the equality of mean. For two pairs: reCAPTCHA (easy image) - hCAPTCHA (easy) and reCAPTCHA (easy click) - (hard click), we didn't see any statistical evidence that the means differ. For the remainder, this test showed strong statistical evidence that the means differ ($p < 0.05$ between masked and moving distorted text and $p < 0.001$ for all other combinations).

D CAPTCHA Solving Times for Other Demographic Features

Figures 13 and 14 show participants' solving times analyzed across other demographic features.




Figure 13: Effects of Gender.




Figure 14: Effects of Education Level.