

Improving Data Quality via Pre-Task Participant Screening in Crowdsourced GUI Experiments

Takaya Miyama
Meiji University
Tokyo, Japan
cs252040@meiji.ac.jp

Satoshi Nakamura
Meiji University
Tokyo, Japan
nkmr@meiji.ac.jp

Shota Yamanaka
LY Corporation
Tokyo, Japan
syamanak@lycorp.co.jp

Abstract

In crowdsourced user experiments that collect performance data from graphical user interface (GUI) interactions, some participants ignore instructions or act carelessly, threatening the validity of performance models. We investigate a pre-task screening method that requires simple GUI operations analogous to the main task and uses the resulting error as a continuous quality signal. Our pre-task is a brief image-resizing task in which workers match an on-screen card to a physical card; workers whose resizing error exceeds a threshold are excluded from the main experiment. The main task is a standardized pointing experiment with well-established models of movement time and error rate. Across mouse- and smartphone-based crowdsourced experiments, we show that reducing the proportion of workers exhibiting unexpected behavior and tightening the pre-task threshold systematically improve the goodness of fit and predictive accuracy of GUI performance models, demonstrating that brief pre-task screening can enhance data quality.

CCS Concepts

• **Human-centered computing** → **HCI theory, concepts and models; Pointing; Empirical studies in HCI.**

Keywords

Crowdsourcing, Pointing, Fitts' law, Graphical user interfaces

ACM Reference Format:

Takaya Miyama, Satoshi Nakamura, and Shota Yamanaka. 2026. Improving Data Quality via Pre-Task Participant Screening in Crowdsourced GUI Experiments. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 22 pages. <https://doi.org/10.1145/3772318.3791332>

1 Introduction

Crowdsourced experiments are conducted over the Internet, enabling participation regardless of time or place. Accordingly, researchers can execute studies quickly while recruiting numerous and diverse participants, and more readily detect statistical significance from the resulting data [35]. However, because experimenters cannot observe participants' circumstances in crowdsourced settings, some participants may disregard instructions or act carelessly,

which degrades data quality [34]. Thus, data-quality management is a critical challenge in crowdsourced experiments [38].

In HCI, graphical user interface (GUI)-based user experiments have long been conducted in controlled laboratory settings for purposes such as (i) validating and comparing predictive models of interaction performance (e.g., Fitts' law [31, 42]), (ii) evaluating the effectiveness of new interaction techniques such as the Bubble Cursor [19, 51], and (iii) comparing performance across different user groups (e.g., younger vs. older adults or people with/without motor impairments [13, 39]). In such studies, the primary aim is to obtain precise, reliable measurements of how users typically perform under standardized task conditions for model fitting and controlled comparison, rather than to capture the full diversity of uncontrolled real-world behaviors.

However, laboratory-based GUI experiments are typically limited to relatively small numbers of participants [14, 27]. Crowdsourcing has therefore become attractive for this class of GUI experiments because it allows researchers to collect much larger and more demographically diverse samples for estimating central tendencies of user performance and model parameters. At the same time, however, crowdsourced settings may increase the risk that some participants will ignore instructions or act inattentively, which threatens the validity of performance measurements and model predictions. For example, there are reports that crowdsourced participants tend to perform pointing tasks less accurately than laboratory participants, with more than twice the error rate observed [14]. Thus, the characteristics of recruited participants may influence outcomes and raise concerns about degraded data quality. This motivates the need for mechanisms that can detect and filter out such non-compliant participants when the research goal is (e.g.) precise model evaluation under controlled task conditions.

This work thus aims to improve data quality in crowdsourced GUI experiments by ensuring that only participants conforming to requirements in a GUI-based pre-task proceed to the main experiment. While our experiments focus on target pointing, we expect the same screening principle to transfer to other downstream GUI-interaction tasks that rely on quantitative performance measures (e.g., goal crossing, path steering, dragging, typing, and object selection; see Section 7.5). Although our approach shares the general goal of identifying low-quality workers with existing techniques such as gold-standard tasks [26] and attention checks [34], it differs in that our pre-task is designed to yield a task-relevant continuous error value that directly reflects operational carefulness in GUI interaction (even though gold tasks and attention checks can also be summarized using continuous measures such as accuracy or response time). These continuous scores can then be thresholded with varying strictness, analogous to choosing different statistical



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

CHI '26, Barcelona, Spain

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2278-3/26/04

<https://doi.org/10.1145/3772318.3791332>

outlier criteria, providing a tunable *minimum-quality* requirement for participation in the main experiment. We instantiate this idea using an image-resizing task as the pre-task and a standard pointing experiment as the main task. Resizing a card image to a prescribed physical size is a simple GUI manipulation whose error sensitively reflects how carefully participants follow instructions, whereas pointing tasks are a canonical testbed in HCI with well-established procedures and performance models (e.g., Fitts' law). This combination allows us to quantify how pre-task-based screening influences the reliability and predictive validity of interaction-performance models in a representative GUI-experiment setting.

We hypothesize that participants who follow instructions in the pre-task are more likely to act conscientiously in the main task, and we examine whether users who will perform the main task as intended (judged by the lawful regularity of their performance) can be identified using only pre-task data. To this end, we simulate experimental outcomes under varying levels of inclusion of participants screened as nonconforming to pre-task requirements, and we assess the utility of this method. Concretely, we vary (i) the proportion of nonconforming participants in the overall pool and (ii) the pre-task threshold T of resizing error used to judge nonconformity. Then, we compute the goodness of fit of models that estimate GUI performance (movement time and error rate) to evaluate whether screening improves model fit.

Because researchers must choose a screening strictness T without knowing how much nonconforming behavior will occur in a crowdsourced pool, we use a simulation framework that systematically varies N , T , and X (the nonconforming proportion) and quantifies how these choices affect model fit (R^2) and, in smartphone settings, predictive accuracy for unseen task conditions.

Importantly, the goal of this screening is *not* to select “top-performing” participants. Instead, it serves as a minimum-quality check that removes workers who do not follow instructions or who behave inattentively, in the same spirit as common exclusion rules based on $\pm 3SD$ or IQR. Thus, the screening threshold controls how lenient or strict this filtering is, rather than privileging high-skill participants. We discuss potential biases of this design choice, for example, against participants with motor or visual impairments, in Section 7.4.

Across three experiments involving mouse and smartphone interaction, we show that reducing the proportion of nonconforming participants and tightening the nonconformity threshold both improve the goodness of fit of GUI performance models and the prediction accuracy for unseen task conditions. Because the information used to judge nonconformity does not rely on outcomes from the main GUI experiment, these results indicate that a brief pre-task can improve the data quality of the main experiment. This enables researchers to screen workers before they proceed to the main experiment, thereby mitigating the risk of drawing incorrect conclusions (e.g., that a given model estimates GUI performance poorly).

Our contributions are threefold:

- We propose a new screening method in which a brief pre-task requiring simple GUI interactions (finished in nine seconds on average) is used to determine which workers proceed to the main experiment.
- We show that changing the nonconformity threshold in the pre-task (namely, error of image resizing from the prescribed size) and the proportion of such workers in the sample systematically affects the goodness of fit and prediction accuracy of models that estimate GUI performance (movement time and error rate) in the main experiment.
- We demonstrate that these effects are robust across mouse interaction (Experiment 1) and two smartphone-interaction settings with different experimental configurations (Experiments 2 and 3).

2 Related Work

This work lies at the intersection of two research areas: (1) crowdsourcing-based experimentation, which enables large-scale data collection but raises concerns about data quality, and (2) GUI interaction research, particularly studies on pointing performance. Because our method introduces a GUI-based pre-task into a crowdsourced experiment in order to screen participants, both strands of prior work are needed to motivate and situate our approach. Accordingly, Section 2.1 reviews prior work on crowdsourcing experiments and data-quality management, whereas Section 2.2 summarizes key concepts from pointing-task research (including movement time MT , error rate ER , and performance models such as Fitts' law) that underpin our experimental design and analyses.

2.1 Data Quality in Crowdsourced Experiments

To address concerns about data quality in crowdsourced experiments, prior work has compared outcomes between crowdsourced and laboratory settings. In psychology and sociology topics, results are often similar across both environments, although some studies report differences [9, 21]. There are also efforts to make online environments more laboratory-like; for example, Li et al. [29] proposed a task that matches the size of a physical card (e.g., a credit card) to an on-screen card image to compute a display's pixel density and thereby control the physical size of visual stimuli.

Research has also examined inattentive participants, with one report finding that 45.9% of crowdsourced participants exhibited inattentive behaviors [7]. Curran [10] proposed effective methods for detecting inattentive outcomes in response times and open-ended text responses. Oppenheimer et al. [34] proposed the Instructional Manipulation Check (IMC) to assess whether participants read instructions. They showed that results differ between participants who pass the IMC and those who do not. Despite these concerns, one review reports that 55% of psychology studies conducting online experiments did not evaluate data quality [18].

For GUI experiments as well, crowdsourced and laboratory results have been compared. Komarov et al. [27] validated the Bubble Cursor [20], reporting that its movement-time advantage over a standard cursor observed in the laboratory also replicated in crowdsourcing. Schwab et al. [37] conducted panning and zooming tasks and reported that Fitts' law holds in both PC and mobile environments. Findlater et al. [14] found that, in mouse and touch pointing tasks, crowdsourced participants exhibited shorter movement times and higher error rates than laboratory participants, suggesting a tendency to prioritize speed over accuracy when instructed to act “as fast and accurately as possible.”

Taken together, diverse evaluations and countermeasures have been explored for data quality in crowdsourced experiments, but the applicability of existing quality-improvement methods to task results of GUI-based experiments (e.g., error rates in pointing tasks) remains unexplored. Our work investigates whether a pre-task tailored to GUI-specific operations can improve data quality upon these efforts.

2.2 Pointing and Other GUI-Related Tasks

When users interact with GUI elements like icons and hyperlinks, their performance, such as how long it takes to reach a target and whether the selection succeeds or fails, has been studied extensively in HCI. A central line of work models these behaviors using movement time (MT) and error rate (ER), which capture how task difficulty and user strategy jointly shape interaction outcomes. A representative model for pointing is Fitts' law, which predicts MT for the first click in a trial based on the target distance A and target width W [16, 31]. There is also research on ER , defined as the percentage of trials in which clicks fall outside the target region [33, 45], showing that ER increases as actions become faster and decreases as actions become more cautious. There are prediction models for ER as well for rectangular and circular targets [4, 53].

Although participants are typically instructed to perform the task as fast and accurately as possible, individuals might unintentionally shift their biases toward either speed or accuracy [39, 49]. Moreover, by running tasks under bias-emphasized conditions, it is possible to evaluate pointing behavior as a function of emphasis. For example, movement time MT decreases while error rate ER increases under speed emphasis [51]. We adopt this approach (emphasizing either speed or accuracy in pointing tasks) to examine whether participants follow the given instructions.

Beyond target pointing, many other GUI interaction tasks have been studied in HCI, including goal crossing [52], path steering [1], dragging [32], typing [3], and object selection [57]. These tasks share the property that user performance can be quantified (e.g., via time or accuracy), and many rely on regularities similar to those captured by Fitts' law or its extensions. Because our screening approach evaluates how carefully participants perform simple GUI manipulations before the main experiment, we expect the same idea to be applicable to a broader class of GUI experiments beyond pointing, especially when the research goal is to obtain reliable performance measurements under controlled task conditions.

3 Approach

3.1 Concept

Our approach is to improve data quality in crowdsourced GUI experiments by screening participants with a pre-task before the main experiment and assigning the intended main task only to conforming workers. Concretely, a pre-task requires operations related to those in the main task and identifies suitable workers based on their performance in this pre-task. In this way, the main experiment can be administered only to the screened workers. This increases the likelihood that most participants will perform interactions as intended, thereby improving data quality. An overview of the proposed approach is shown in Figure 1.

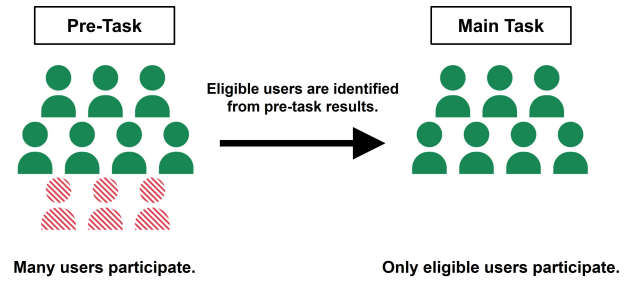


Figure 1: Overview of the proposed approach

3.2 Pre-Task for Participant Screening

In this study, we assume that researchers have a primary GUI experiment that they ultimately wish to analyze, and that this experiment is administered only after screening participants using a brief pre-task. We refer to this primary GUI experiment as the “main task.” Because our main task in this study is target pointing, a pre-task including operations related to pointing can be used to screen conforming participants. We focus on the size-adjustment task proposed by Li et al. [29], in which participants adjust an on-screen card image to match the size of a physical card (e.g., a credit card) placed on the screen, see Figure 2. The accuracy of the operation can be evaluated by measuring the error between these sizes. We assume that this is suitable as a pre-task when the main task requires accuracy as in pointing.

Although the size-adjustment task is typically administered prior to an experiment to control the physical size of visual stimuli [29], in our previous experiments we found some participants exhibiting unexpected behavior (e.g., submitting without resizing at all, or completing the resizing in an extremely short time). Therefore, by adopting the size-adjustment task as a pre-task, we expect to identify such participants in advance and restrict the main task to participants conforming to pre-task requirements.




Figure 2: The size-adjustment task by Li et al. [29]

3.3 Evaluation and Research Questions

We evaluate the proposed approach by running crowdsourced experiments in which the size-adjustment task is administered as a pre-task followed by a main task. As mentioned, this pre-task is not for detecting a top-performing group; rather, it serves as a minimum-quality criterion that excludes workers who do not follow instructions or who behave inattentively in the pre-task. In this sense, the threshold T of the error between physical and on-screen

cards to detect nonconforming workers plays a role similar to common statistical exclusion rules (e.g., $\pm 3SD$ or IQR-based filtering), allowing researchers to choose more lenient or more conservative settings depending on their study design without biasing the sample toward only high-skill users. Participants whose error is below a threshold T pass, whereas those at or above the threshold fail. Building on this, we simulate experimental outcomes when participants who fail the size-adjustment task (the non-passing group) are mixed into the sample. Specifically, we vary the proportion of the non-passing group in the overall pool and the error threshold, compute the goodness of fit of models that estimate GUI operation performance, and evaluate whether screening improves model fit.

We address the following research questions (RQs):

- RQ1** Can a simple GUI pre-task (image resizing) effectively screen out nonconforming participants?
- RQ2** How do the proportion of the non-passing group and the error threshold affect the goodness of fit and predictive accuracy of performance models (movement time and error rate) for the main pointing task?
- RQ3** Is the proposed screening method consistently effective across device types (PC/smartphone) and across different error-handling policies (with or without re-aiming when missing a target)?

Our expectation is that model fit will improve as the proportion of the non-passing group decreases and as the threshold is tightened, and that nonconforming participants can be effectively screened regardless of device. We conducted the following three experiments to evaluate these expectations:

- Experiment 1** PC mouse-based experiment; in case of an error, participants re-aim until the target is successfully selected.
- Experiment 2** iPhone-based experiment; in case of an error, participants re-aim until the target is successfully selected.
- Experiment 3** iPhone-based experiment; even if an error occurs, the next trial begins immediately.

We first conducted a preliminary check in Experiment 1 to confirm that screening functions as intended, and then decided to perform a more rigorous evaluation in Experiments 2 and 3. We targeted the iPhone environment in Experiments 2 and 3 because, by restricting the device to iPhone, we can control the physical size of on-screen stimuli in millimeters (mm) even in crowdsourced experiments and thus evaluate models more rigorously [44, 55]. We also prepared multiple error-handling policies because under a “re-aim until success” policy, inattentive participants may try to perform as quickly and accurately as possible to finish a task in a short time, making them difficult to distinguish from conscientious participants. Previous studies have used both with [8, 17, 20, 56] and without re-aiming when errors occur [23, 45, 47].

All experiments were accessible through a web system we developed, and participants were recruited via Yahoo! Crowdsourcing [30] without demographic stratification (e.g., no age balancing). Workers voluntarily chose to participate in the tasks from our listings, and we did not pre-screen them by demographic or performance attributes. Thus, our sample reflects a natural, unstratified subset of the platform’s worker population. We chose Yahoo! Crowdsourcing because our previous GUI experiments on this platform revealed a non-trivial number of inattentive workers,

making it a suitable environment for evaluating the effectiveness of the proposed screening method. Demographic information (age and gender), as reported to the platform, was recorded and is summarized in the supplementary materials.

This study was approved by the research ethics committee at our institution including the legal department (serving as the equivalent of an Institutional Review Board (IRB)) as well as by the crowdsourcing platform. We obtained informed consent from all workers, who agreed that we may record, analyze, and publish data to the extent necessary for research.

4 Experiment 1: Mouse-Based Experiment on PC

4.1 Task and Design

4.1.1 Overview. In Experiment 1, we conducted a mouse-based study in a PC environment to preliminarily verify whether screening via the size-adjustment pre-task works as intended. Upon accessing the experimental system, participants first viewed an explanation, then completed the size-adjustment task twice, one practice block of the pointing task, and four main (data-collection) blocks of the pointing task. In the practice block, participants were instructed to act “as fast and accurately as possible” to familiarize themselves with the pointing interaction. Then, in the main blocks, we asked them to act either “as fast as possible” or “as accurately as possible” to test if they faithfully follow our instructions.

4.1.2 Size-Adjustment Task. As described in Section 3.2, the size-adjustment task required participants to place a physical card on the display and adjust an on-screen card image to match its size. Because pixel density varies across participants’ displays in a PC-based crowdsourcing setting, a single size-adjustment trial is insufficient to judge whether a participant adjusted accurately. Therefore, we administered the task twice and assessed each participant’s accuracy based on the discrepancy between the two adjustments.

To avoid cases in which the initial on-screen card size happened to be close to the physical card, thereby eliminating the need to adjust, we set the initial size to be either extremely small or large. Specifically, the initial width was set to 100 px or 900 px, and these two conditions were assigned once each across the two repetitions in random order. We allowed the use of physical cards conforming to the ISO/IEC 7810 ID-1 standard (53.98 mm \times 85.60 mm) [24]. Representative examples widely used in Japan include credit cards, transportation IC cards such as Suica and PASMO, health insurance cards, and driver’s licenses.

The size adjustment was performed on the screen shown in Figure 3 by dragging the lower-right corner of the card image. This behavior closely follows the virtual-chinrest plugin¹ of jsPsych [11], developed based on Li et al.’s work [29]. Although size adjustment is typically used to control the physical size of visual stimuli, across Experiments 1–3 we did not control target sizes in the pointing task based on the size-adjustment results. This is because some participants would operate inaccurately during the size adjustment; if we had computed pixel density (pixels per mm) from their results and used it to scale target sizes in the main task, we might have displayed excessively large or small stimuli.

¹<https://www.jspsych.org/v7/plugins/virtual-chinrest/>

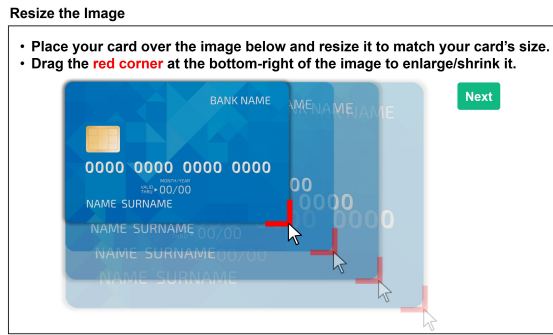


Figure 3: Size-adjustment task in Experiment 1

Thus, the size-adjustment task served solely as a pre-task to identify conforming users.

4.1.3 Pointing Task: Design and Procedure. Participants clicked circular targets on the screen according to the given instruction. After clicking a button labeled *Start*, targets appeared within a 1280×720 px region (Figure 4). When a target was clicked, the next target appeared, and trials proceeded in this manner. If a click missed the target, participants continued to *re-aim* until they succeeded. Upon a successful click, the next trial began.

The target distance A was fixed at 510 px, and the target diameter W took values of 8, 38, and 78 px. Target positions were random apart from satisfying the distance A constraint. In the main blocks, participants performed the task under two instructions, following prior work [50, 58]:

- **As fast as possible:** Click the target as quickly as you can after it appears, but do not act more carelessly than necessary such as clicking without aiming.
- **As accurately as possible:** Avoid errors (clicks outside the target), but do not spend more time than necessary.

Each practice and main block consisted of 15 trials. Among the 15 targets, each W condition appeared five times in random order. Participants completed four main blocks in total, with the two instructions assigned to two blocks each in random order. Thus, each participant completed $3W \times 5$ trials \times 2 instructions \times 2 blocks = 60 trials in the main blocks. Each ($W \times$ instruction) condition consisted of 5 trials \times 2 blocks = 10 trials.

4.2 Participants

A total of 500 participants completed the experiment. Each received a reward of 200 JPY (approximately 1.3 USD). The entire task completion time was 4 min 4 s on average, yielding an effective hourly payment of 2,951 JPY (approximately 19.1 USD). We excluded 29 participants who reported using a trackpad rather than a mouse in the questionnaire and 16 participants with missing data, leaving 455 participants (354 male, 101 female) for analysis.

4.3 Results

4.3.1 Size-Adjustment Task. Figure 5 shows the distribution of the long-side length (px) of the card image after the two adjustments. Participants' adjusted sizes had a mean of 359.45 px (SD=144.17 px) in the first trial and 359.41 px (SD=121.20 px) in the second; 390

Click the target as quickly as possible (15 left).




Figure 4: Pointing task in Experiment 1

participants (86%) had both adjustments between 200 px and 600 px. Because each participant used a different display with its own pixel density, the absolute pixel length corresponding to the physical card naturally varied across participants. For any given participant, however, the two adjustment trials should yield similar pixel lengths if they follow the instructions and actually align the image with the physical card. Figure 5 shows that many participants indeed produced consistent sizes within a plausible range (200–600 px)², while a noticeable subset left the image at the initial sizes (100 or 900 px) or produced clearly inconsistent pairs of sizes. These patterns indicate that some participants did not perform the required resizing operation at all or did so inattentively, highlighting the need to use the pre-task as a basis for screening.

Figure 6 shows the distribution of absolute adjustment error across the two trials. Most participants had small errors, specifically, 354 (78%) were below 20 px, but some had large errors, with 73 (16%) at or above 50 px. The drag-operation times were, on average, 8.03 s (SD=6.13 s) in the first trial and 6.27 s (SD=5.49 s) in the second. However, 38 participants in the first and 25 in the second had 0 s, indicating no adjustment at all.

4.3.2 Pointing Task. We first removed outliers. For each participant and each instruction \times W condition (10 trials), we excluded trials with outlier click coordinates. We rotated coordinates so that the direction from the previous target center to the current target center aligned with the positive x -axis; we then recorded the x -coordinate as positive if the click overshot the current target center and negative if it fell short [42, 47]. Using the standard deviation σ of these x -coordinates, we excluded trials whose x -coordinates deviated from the mean by more than 3σ [42]. For the first target in a block, the *Start* button is treated as the *previous target*. Similarly, for the 10 MT values in each instruction \times W condition, we excluded outliers using the same 3σ rule. Finally, using the mean MT across all 60 trials per participant, we applied the 3σ rule and excluded that participant. This procedure identified five participants (1.1%) as outliers. No trial-level outliers were detected for click coordinates or MT .

²If the long side of a credit card really corresponded to 200 px on the screen, the display would be approximately 37.3 inches assuming full HD resolution; if it were 600 px, it would be 12.4 inches. We set thresholds under the assumption that contemporary displays typically fall within this range, though displays outside this range may also be used. Therefore, in Experiments 2 and 3 we limited the devices to iPhones, for which we can obtain screen size in both mm and px, to evaluate the effectiveness of our screening more rigorously.




Figure 5: Distribution of the adjusted long-side length of the card image across the two size-adjustment trials in Experiment 1.




Figure 6: Distribution of absolute discrepancies between the two size-adjustment trials in Experiment 1.

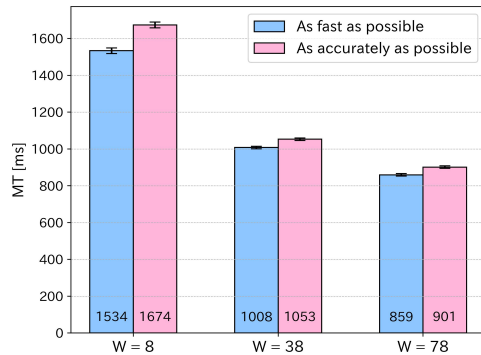


Figure 7: MT results for each W under the fast and accurate instruction conditions

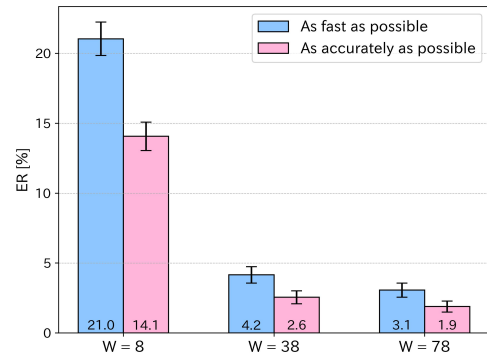


Figure 8: ER results for each W under the fast and accurate instruction conditions

Figure 7 shows MT and Figure 8 shows ER results. Throughout the paper, error bars indicate 95% confidence intervals. Because prior research has shown that ANOVA is robust against violations of the normality assumption [6, 36, 43], we consistently ran RM-ANOVAs (independent variables were instruction and W ; dependent variables were MT and ER). Bonferroni correction was used to adjust the p -values in pairwise tests. Analyses not directly related to our research questions (e.g., normality tests for MT or complete results of omnibus tests across conditions) are provided in the supplementary materials.

For MT , we found significant main effects of instruction ($F_{1,449} = 173.4$, $p < 0.001$, $\eta_g^2 = 0.026$) and W ($F_{2,898} = 4053$, $p < 0.001$, $\eta_g^2 = 0.65$). The interaction effect of $W \times$ instruction was significant ($F_{2,898} = 62.00$, $p < 0.001$, $\eta_g^2 = 0.010$). For ER , we found significant main effects of instruction ($F_{1,449} = 56.61$, $p < 0.001$, $\eta_g^2 = 0.022$) and W ($F_{2,898} = 612.5$, $p < 0.001$, $\eta_g^2 = 0.29$). The interaction effect of $W \times$ instruction was significant ($F_{2,898} = 34.41$, $p < 0.001$, $\eta_g^2 = 0.014$). Briefly, MT and ER decreased as the target size increased ($p < 0.001$ for all pairwise tests), and the participants changed their MT and ER according to the given instructions ($p < 0.001$ for both), consistent with prior work [50, 58]. Thus, on average, participants exhibited typical pointing behavior.

4.4 Simulation

4.4.1 Overview. To assess the utility of our screening method, we simulated cases in which participants whose behavior is identified by the pre-task screening as nonconforming are mixed into the sample. Assuming a researcher's goal is to evaluate the goodness of fit of existing performance models, we examine whether screening improves model fit. To this end, we compute the fit of pointing-performance models as the proportion of nonconforming participants varies.

The GUI performance models we employ are given in Equations (1)–(3). Lowercase italics a – f are regression coefficients estimated from the data. Equation (1) is Fitts' law, which predicts movement time MT from the target distance A and diameter W [31]. Equation (2) estimates spreads of click coordinates based on target size. For each W , after rotating the task axis (x -axis) towards the current target as explained in Section 4.3.2, we regress the standard deviations σ_x and σ_y of click coordinates (with the target center as origin) along the x - and y -axes [48, 54]. Equation (3) predicts the success probability $P(D)$ that a click (x, y) falls within a circular target region D of diameter W . Using σ_x and σ_y predicted by Equation (2), we obtain $P(D)$ for each W [54].

$$MT = a + bID, \quad ID = \log_2 \left(\frac{A}{W} + 1 \right) \quad (1)$$

$$\sigma_x = c + dW, \quad \sigma_y = e + fW \quad (2)$$

$$P(D) = \iint_D \frac{1}{2\pi\sigma_x\sigma_y} \exp\left(-\frac{x^2}{2\sigma_x^2} - \frac{y^2}{2\sigma_y^2}\right) dx dy \quad (3)$$

These models have shown high goodness of fit (typically R^2) in prior work [42, 54]. Hence, if data quality is sufficient, we should obtain similar results. Conversely, if data quality is poor, for instance, if participants do not vary MT or ER with W , model fit will degrade.

4.4.2 Procedure. Our screening aims to split participants into a “passing group” and a “non-passing group” based on the pre-task, and to administer the main task only to the passing group. After applying the outlier handling in Section 4.3.2, we classified conforming participants into the passing group as follows: those whose first and second adjusted sizes both fell between 200 px and 600 px (the proper range, see Figure 5) and whose absolute discrepancy between the two adjustments was smaller than a threshold T px (see Figure 6). All others were assigned to the non-passing group. We incorporated the range criterion to exclude participants whose two adjustments were both near extreme bounds of the slider (suggesting nonconformity) and those who might not have used a physical card yet happened to exhibit a small adjustment discrepancy.

In the simulation, for a total sample size N and a non-passing proportion $X\%$, we randomly drew $N \times X\%$ participants from the non-passing group and $N \times (100 - X)\%$ from the passing group. We then computed model fit for Equations (1) and (3) on the pointing task data of the sampled set. We repeated this sampling-and-fitting procedure 1,000 times and averaged the fit.

We varied N , T , and X to analyze how the inclusion of nonconforming participants affects researchers’ model-fit results. Specifically, we set $N \in \{10, 20, 40, 80\}$, T from 5 px to 50 px in steps of 5 px, and X from 0% to 100% in steps of 10%. We capped N at 80 and T at 50 px because larger N requires more non-passing data, and overly lenient T reduces the number of non-passing participants, making it difficult to secure sufficient data for analysis. Because Experiment 1 serves as a preliminary PC-based check, in the main text we explain the results of $N = 40$ as an example. The other sample sizes ($N = 10, 20, 80$) showed similar quantitative trends and are thus summarized in the supplementary materials. Figure 9 shows the numbers of passing and non-passing participants as T varies.




Figure 9: Numbers of passing vs. non-passing participants as a function of the threshold

4.5 Simulation Results

4.5.1 Movement Time (MT). Figures 10 and 11 show R^2 of Fitts’ law for Equation (1) when varying T and X for movement time MT . Figure 10 corresponds to the “fast” instruction and Figure 11 to the “accurate” instruction. For $N = 40$, R^2 remained above 0.98 for all T and X in both the fast and accurate conditions, with very little variation. A likely reason is that Fitts’ law tends to achieve high fit even with relatively small samples and repetitions compared to ER models [48]. In addition, participants with larger discrepancies in the pre-task are expected to act more carelessly in the main task, but Equation (1) models MT , so accuracy-based screening plays a smaller role. Consequently, even a higher proportion of careless participants may have limited impact on MT -model fit.

Therefore, when validating an MT model such as Equation (1) via crowdsourcing, the utility of our screening may be limited. Note that we did not evaluate cross-validated prediction for unseen task conditions in Experiment 1, because the number of W conditions (three) was too small to support an appropriate cross-validation design.

4.5.2 Error Rate (ER). Figures 12 and 13 present the analogous analysis for error rate ER , showing R^2 for Equation (3). As with MT , the R^2 values for $N = 40$ did not remarkably change with T and X . Yet, in Figure 13, R^2 tends to decrease from the upper-left to the lower-right. This indicates that the model fit declines as the threshold T is more lenient (large) and as the non-passing proportion is larger. When T is lenient, participants with very large pre-task discrepancies populate the non-passing group; as their proportion increases, model fit worsens. Hence, when validating an ER model such as Equation (3) via crowdsourcing, screening participants and restricting the main experiment to conforming users can improve model fit.

The high prediction accuracy of Equation (3) has been confirmed in previous studies [48, 54], and thus this simulation result demonstrates that our screening method enables researchers to draw valid experimental conclusions on model evaluation. Because our screening leverages accuracy in the size-adjustment pre-task, it is particularly beneficial when evaluating ER models. The trend that lower non-passing proportions yield higher R^2 further supports the usefulness of screening.

We also examined the other total sample sizes $N \in \{10, 20, 80\}$; larger N tended to yield slightly higher R^2 , but the qualitative pattern that stricter T and smaller X improve model fit was unchanged. Detailed heatmaps for these sample sizes are provided in the supplementary materials.

By contrast, Figure 12 does not show the same trend. Under the “fast” instruction, participants are forced to act quickly, and thus the difference in accuracy between the passing (presumably accurate) and non-passing (presumably careless) groups may be attenuated, which could explain the lack of a clear effect of X .

4.6 Limitations

Although Figure 13 suggests that screening is effective, the top-left cells (threshold $T = 5$, non-passing proportion $X = 0$) are not always the best-fitting across the heatmap. Since $T = 5$ is the strictest threshold and $X = 0$ includes no non-passing participants, this cell represents the ideal sample for an experimenter. A plausible reason

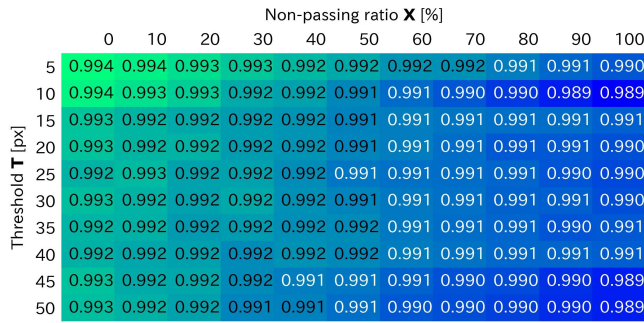


Figure 10: Experiment 1: Goodness of fit to Equation (1) for MT (fast)




Figure 11: Experiment 1: Goodness of fit to Equation (1) for MT (accurate)

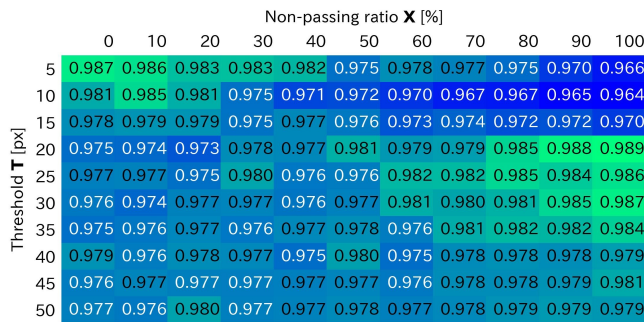


Figure 12: Experiment 1: Goodness of fit to Equation (3) for ER (fast)

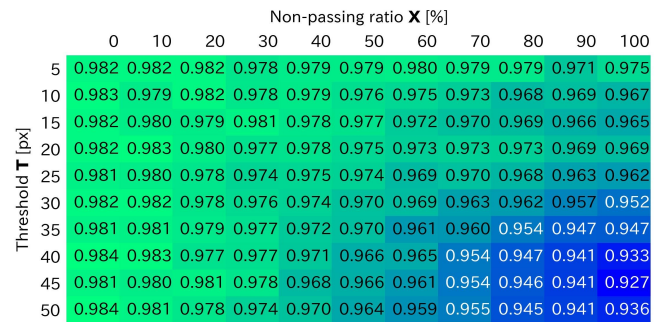


Figure 13: Experiment 1: Goodness of fit to Equation (3) for ER (accurate)

behind this non-optimal result is that we used somewhat limited data size for precise model evaluation: only three W conditions and 10 trials per instruction \times W condition.

To more rigorously assess the utility of our screening, we should increase the number of W levels and trials per condition in subsequent experiments. Furthermore, our evaluation of the size-adjustment task was based on the discrepancy between the two trials, but a participant who did not use the designated physical card might, by chance, exhibit a small discrepancy. Thus, we could not guarantee that the adjusted size actually matched the physical card. Addressing these limitations is necessary for a stricter validation of the size-adjustment task; we do so in Experiment 2.

5 Experiment 2: iPhone-Based Experiment with Re-Aiming

Regarding participant overlap across experiments, workers who took part in Experiment 1 were not prevented from later participating in Experiments 2 or 3, and we did not track whether such overlap occurred. By contrast, Experiments 2 and 3 were posted as separate tasks at different times, and we ensured that no worker participated in both. Because Experiment 1 used a PC-based mouse interaction setting whereas Experiments 2 and 3 used smartphone-based touch interaction with different task structures, we consider any learning effects across experiments to be negligible.

5.1 Task and Design

5.1.1 Overview. In Experiment 2, we examined more rigorously whether screening via the size-adjustment pre-task works as intended by restricting the devices to iPhones. We can identify each device’s display pixel density and accurately evaluate whether a participant’s adjusted size matches the physical card. Specifically, we used iPhone screen-resolution data³ to infer each participant’s device PPI from the resolution obtained by the experimental system. We then converted the adjustment results from px to mm using that PPI, enabling strict measurement of error relative to the physical card.

We also tightened control in the pointing task. First, targets were rendered in mm based on each participant’s device PPI. We further increased the number of W conditions to nine and the number of trials per instruction \times W condition to 20, allowing a more rigorous evaluation of model fit. This addressed the limitation in Experiment 1, where few W levels and repetitions hindered thorough model evaluation.

The procedure largely followed Experiment 1: after instructions, participants completed one size-adjustment task, one practice block of the pointing task, and four main blocks. Whereas Experiment 1 used two size adjustments, here a single adjustment sufficed because

³<https://www.ios-resolution.com/>

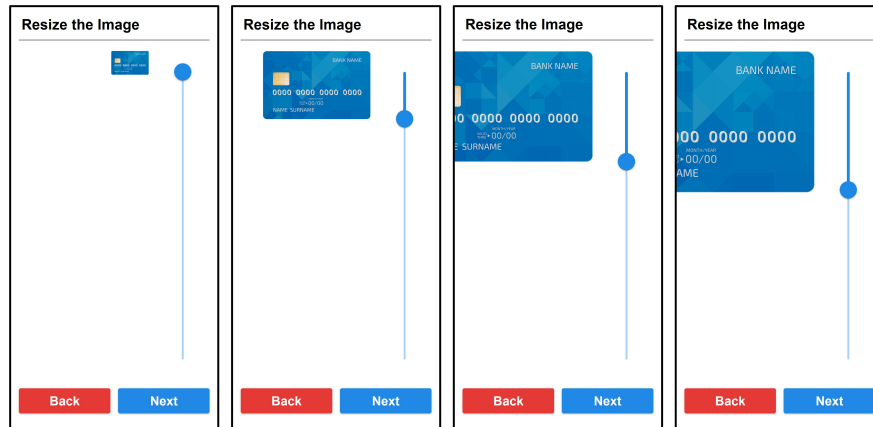


Figure 14: Size-adjustment screen in Experiment 2




Figure 15: How participants operated the size-adjustment task

restricting to iPhones provided a correct reference for each participant. The study targeted iPhone 7 and later models to properly render targets and texts.

5.1.2 Size-Adjustment Task. This pre-task was conducted on the screen shown in Figure 14; participants resized the image via a slider. As in Figure 15, participants placed a physical ISO/IEC 7810 ID-1 card on the iPhone and matched the shorter side of the on-screen card image to the card's shorter side. We set the initial size of the shorter side of the image to 50 px, because this ensures that, on any iPhone model, it is smaller than the shorter side of the physical card (53.98 mm), thereby requiring slider dragging.⁴ As in Experiment 1, we did not scale pointing task targets based on the pre-task's inferred pixel density. Instead, we controlled target size using the PPI identified from screen resolution.

5.1.3 Pointing Task. As shown in Figure 16, two horizontally long rectangular targets, which could have different heights, were displayed on the screen. Each target spanned the full-screen width and was W mm in height. Participants tapped them alternately according to the instructions. The top target initially served as the start

button. When participants tapped a target, its height changed to the next W condition. As in Experiment 1, if a tap missed, participants continued re-aiming until they succeeded. Upon a successful tap, they could proceed to tap the opposite target.

The center-to-center distance of targets A was 30 mm (one level). The target height W took nine levels: 2, 2.8, 3.6, 4.4, 5.2, 6, 6.8, 7.6, and 8.4 mm. This range covered very easy to difficult conditions, yielding error rates from approximately 0% up to 29% [53, 55]. Target centers appeared at two fixed vertical positions and alternated top and bottom. As in Experiment 1, the two instructions were "as fast as possible" and "as accurately as possible," and the practice instruction was "as fast and accurately as possible." Participants were instructed to hold the smartphone in the non-dominant hand and tap with the dominant index finger (Figure 17) [4, 53].

Each block comprised 90 pointing trials, with each of the nine W levels appearing 10 times in random order. In the main phase, participants completed four blocks; the two instructions appeared twice each in random order. Thus, each participant completed $9W \times 10 \text{ trials} \times 2 \text{ instructions} \times 2 \text{ blocks} = 360 \text{ trials}$, yielding $10 \text{ trials} \times 2 \text{ blocks} = 20 \text{ trials per } (W \times \text{instruction}) \text{ condition}$.

⁴For example, iPhone 16 has 460 PPI and Scale Factor of 3, yielding 50 px as 8.28 mm on the screen.




Figure 16: Smartphone pointing-task screen in Experiment 2. Two horizontally oriented rectangular targets span the full screen width, and their height is set to the target width W in mm.



Figure 17: Instructed posture for the smartphone pointing task in Experiments 2. Participants were asked to hold the phone in the non-dominant hand and tap the targets with the index finger of the dominant hand.

5.2 Participants

A total of 584 participants completed the experiment. Each received 250 JPY (approximately 1.7 USD). The mean task completion time was 6 min 43 s, yielding an effective hourly payment of 2,233 JPY (approximately 15.2 USD). We excluded 8 participants with missing data and 42 for whom the iPhone PPI could not be identified, leaving 534 participants for analysis (297 male, 235 female, 2 other).

5.3 Results

5.3.1 Size-Adjustment Task. Figure 18 shows the distribution of absolute errors relative to the correct size in the size-adjustment task. Most participants had small errors, specifically, 317 (59%) were below 2 mm, but some showed large errors, with 144 (27%) at or above 10 mm. The operation time for the adjustment averaged 8.86 s (SD=7.06 s); five participants recorded 0 s, indicating no adjustment at all.




Figure 18: Distribution of size-adjustment errors

5.3.2 Pointing Task. To remove outliers on tap coordinates, we measured y as positive when the tap landed below the target center and negative when above, and excluded trials whose y -coordinates deviated from the mean by more than 3σ . Other criteria on trial- and participant-level MT followed Experiment 1 (see Section 4.3.2). This procedure excluded 1,042 trials for tap coordinates, 2,547 trials for MT , and four participants for MT . These counts are not mutually exclusive, as some trials and participants met multiple exclusion criteria, and we analyzed the remaining 187,347 trials (97.5%).

Figures 19 and 20 show the results of the pointing task in Experiment 2: MT by ($W \times$ instruction) and ER by ($W \times$ instruction), respectively. For MT , we found significant main effects of instruction ($F_{1,529} = 256.3, p < 0.001, \eta_g^2 = 0.042$) and W ($F_{8,4232} = 493.3, p < 0.001, \eta_g^2 = 0.035$). The interaction effect of $W \times$ instruction was significant ($F_{8,4232} = 38.75, p < 0.001, \eta_g^2 = 0.0014$). For ER , we found significant main effects of instruction ($F_{1,529} = 121.9, p < 0.001, \eta_g^2 = 0.014$) and W ($F_{8,4232} = 1751, p < 0.001, \eta_g^2 = 0.51$). The interaction effect of $W \times$ instruction was significant ($F_{8,4232} = 22.22, p < 0.001, \eta_g^2 = 0.0084$). When W increased from 2.0 to 8.4 mm under the fast condition, MT changed from 498 to 438 ms, i.e., change ratio was $(498 - 438)/498 = 12\%$. In contrast, ER changed from 32.1 to 0.8%, i.e., change ratio was $(32.1 - 0.8)/32.1 = 98\%$. A similar trend was found for the accurate condition. This result indicates that ER is more sensitive to W than MT , which is also supported by the larger effect size ($\eta_g^2 = 0.035$ for MT vs. 0.51 for ER). As intended, ER reached roughly 30% at the smallest W .

5.4 Simulation

5.4.1 Overview. As in Experiment 1 (Section 4.4), we simulated outcomes when participants whose behavior is identified by the




Figure 19: MT results for each W under the fast and accurate instruction conditions

pre-task screening as nonconforming are mixed into the sample. In addition to Equation (1) to model MT (Fitts' law), the GUI performance models used in Experiment 2 are Equations (4) and (5), where g and h are regression coefficients estimated from the data. Equation (4) shows that the variance of tap y -coordinates is linearly related to the square of target size [4, 53]. We can then predict the tap success probability $P(-W/2 \leq Y \leq W/2)$ for each W by Equation (5) using σ_y from Equation (4).

$$\sigma_y^2 = g + hW^2 \quad (4)$$

$$P\left(-\frac{W}{2} \leq Y \leq \frac{W}{2}\right) = \text{erf}\left(\frac{W}{2\sqrt{2}\sigma_y}\right) \quad (5)$$

5.4.2 Procedure. In Experiment 2, participants whose adjustment error was smaller than a threshold T (mm) were assigned to the passing group; all others were assigned to the non-passing group. We varied T from 1 mm to 10 mm in 1 mm steps. The values of N and X were the same as in Experiment 1. Figure 21 shows the numbers of passing and non-passing participants as T varies. After applying the outlier handling in Section 5.3.2, we followed the same simulation procedure as in Experiment 1.




Figure 21: Numbers of passing vs. non-passing participants as a function of the threshold

5.5 Simulation Results

5.5.1 Movement Time (MT). Figures 22 and 23 show R^2 for Equation (1) when varying T and X . In both figures, R^2 tends to decrease from the upper-left to the lower-right of each heatmap: when the




Figure 20: ER results for each W under the fast and accurate instruction conditions

threshold is lenient and the non-passing proportion is large, model fit declines. Thus, unlike Experiment 1, the non-passing proportion clearly affects the evaluation of Equation (1).

A likely reason for this difference is the device environment. In Experiment 1 (PC with mouse), even an inattentive participant attempting to finish quickly still needed to move the pointer according to target distance, which may have produced relatively normal MT values. In Experiment 2 (smartphone), target positions were fixed and alternating, and thus, an inattentive participant could finish quickly by using, for example, two fingers or both hands to tap without moving a single finger. Although we instructed participants to hold the phone in the non-dominant hand and tap with the dominant index finger (Figure 17), inattentive participants may have ignored these instructions, leading to abnormal MT . This may also explain the higher outlier rate compared to Experiment 1. Hence, in crowdsourced smartphone settings where controlling participants' tapping techniques is difficult, our screening is particularly useful for validating MT models like Equation (1).

As in Experiment 1 (see Section 4.5.2), increasing N improved R^2 in the upper-left regions of the heatmaps and mitigated the impact of larger X . Nevertheless, R^2 still varied with X in Figures 22 and 23. Thus, when N is not large enough to absorb the inclusion of nonconforming participants, our screening can suppress the impact of nonconforming participants and improve model fit.

We additionally conducted leave-one- W -out cross-validation (LOOCV), in which we iteratively hold out one width level W as unseen data, train the model on the remaining levels, and then evaluate prediction accuracy on the held-out level. An example for $N = 40$ under the "accurate" instruction is shown in Figure 26. Compared to using all W levels, R^2 is slightly lower, but the same trend holds: model accuracy deteriorates from the upper-left to the lower-right. Thus, beyond model fit, our screening is effective for predicting user performance under new (unseen) task conditions. Complete LOOCV results are included in the supplementary materials.

5.5.2 Error Rate (ER). Figures 24 and 25 show R^2 for Equation (5) when varying T and X . Again, R^2 declines from the upper-left to the lower-right of each heatmap: lenient thresholds and large non-passing proportions degrade fit. Thus, the degradation observed in Experiment 1 is replicated. Because Experiment 2 increased the

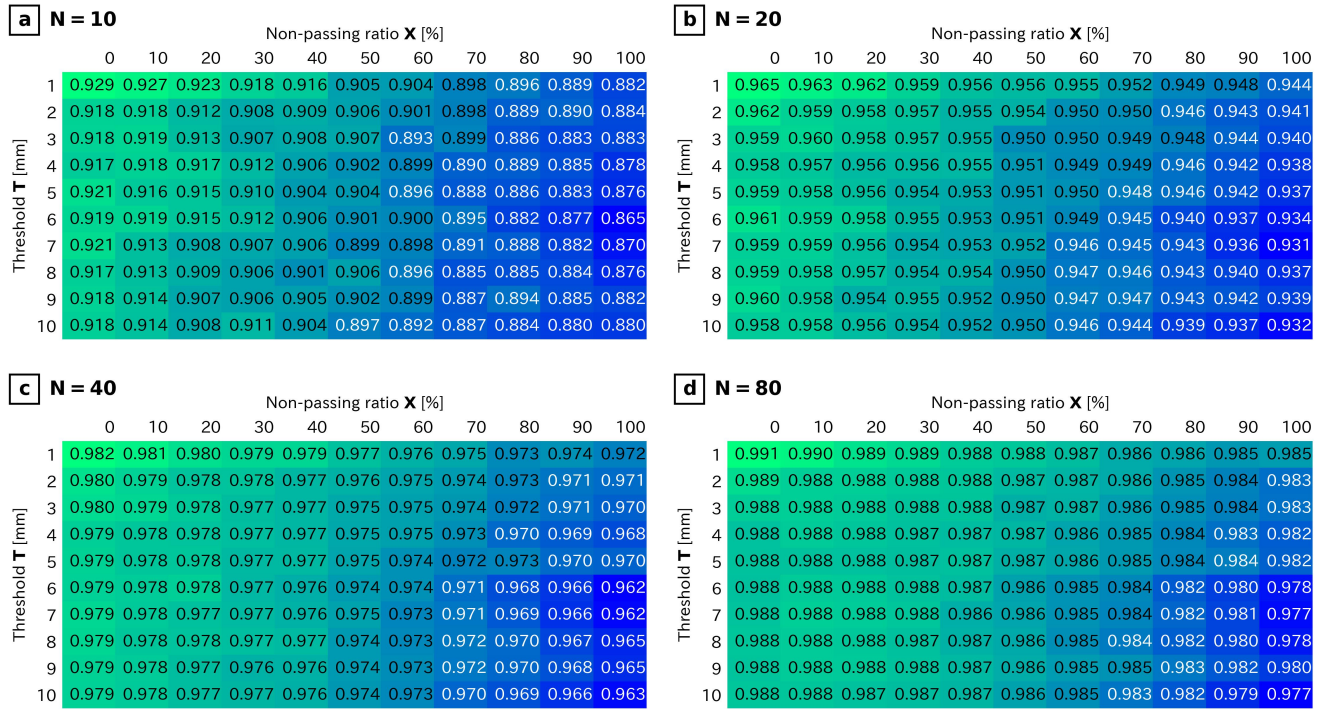


Figure 22: Experiment 2: Goodness of fit to Equation (1) for MT (fast)

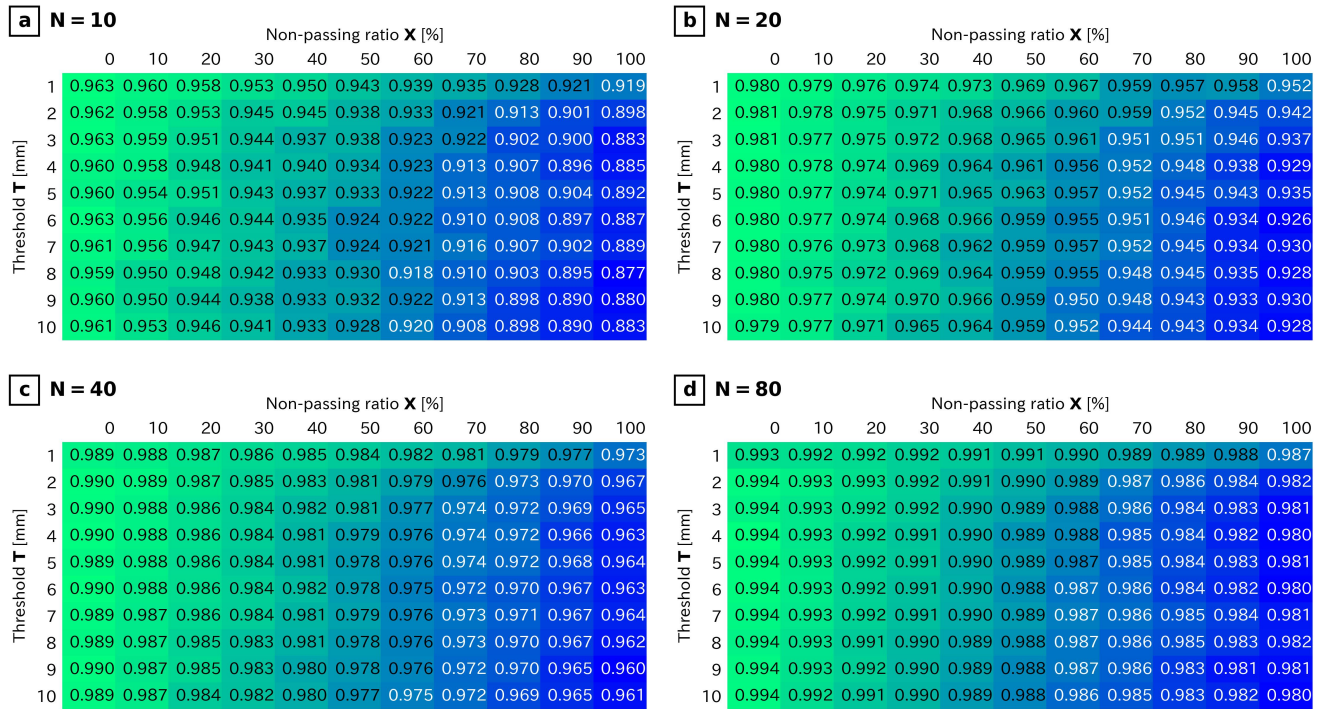


Figure 23: Experiment 2: Goodness of fit to Equation (1) for MT (accurate)

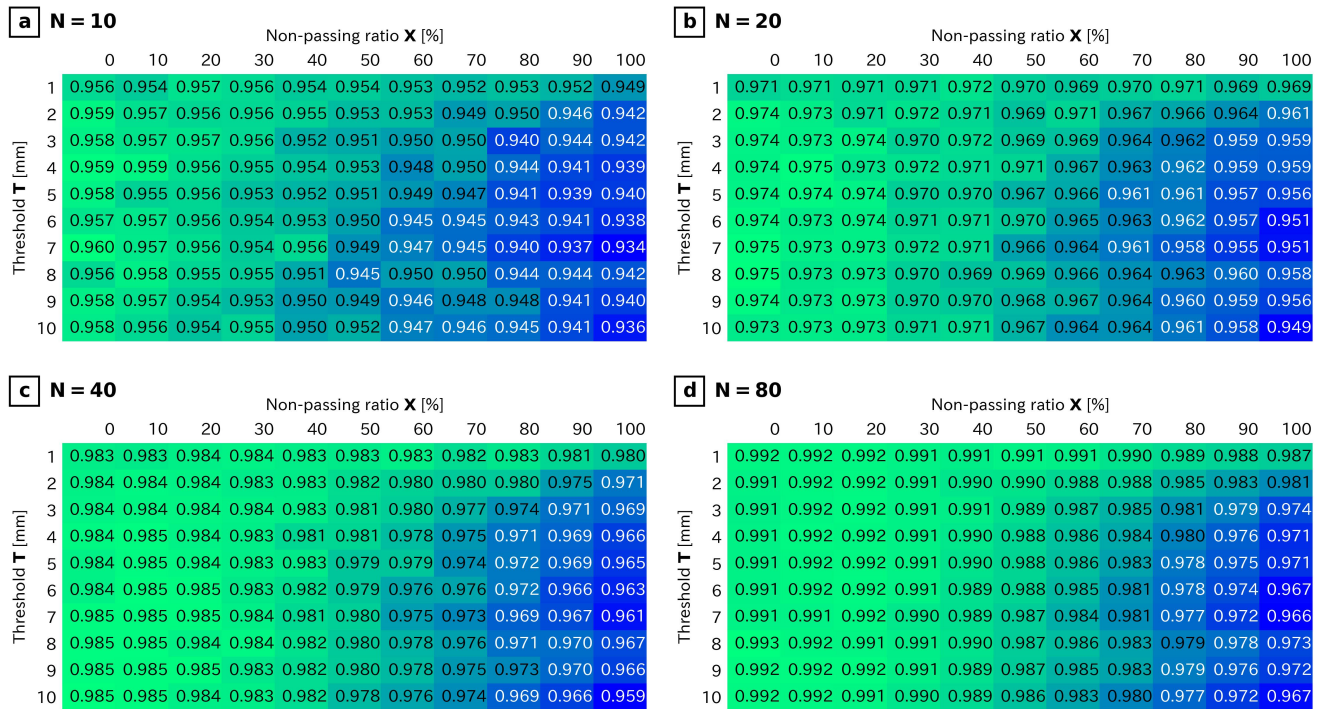


Figure 24: Experiment 2: Goodness of fit to Equation (5) for ER (fast)

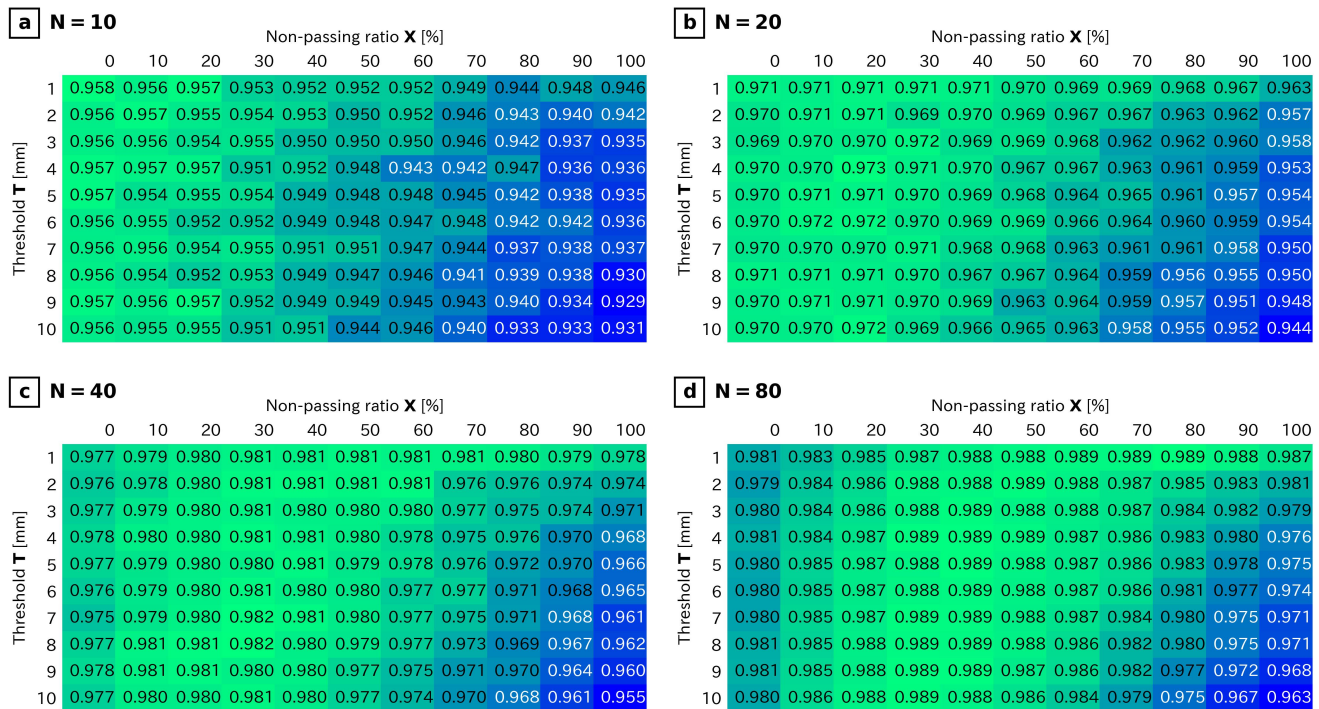


Figure 25: Experiment 2: Goodness of fit to Equation (5) for ER (accurate)




Figure 26: Cross-validation results of R^2 for MT ($N = 40$, accurate)

number of W levels and repetitions per task condition and controlled target sizes in mm, these results more strongly support that screening via the size-adjustment pre-task is effective. The effect of N on ER mirrored the MT results, as we observed that increasing N raised R^2 overall, most prominently in the upper-left region of the heatmaps (strict T with small X), and partially buffered the loss of fit at larger X .

Unlike Experiment 1, even under the “fast” instruction (Figure 24), increasing X reduced model fit. As described, because smartphone settings make it harder to control tapping operations using only the dominant hand’s index finger than PC settings, it may induce more noise to the data and thus amplify the negative effect of X .

As a cross-validation example, Figure 27 shows the “accurate” instruction with $N = 40$. As with MT , R^2 is slightly lower than when using all W levels, yet the same upper-left to lower-right degradation trend holds. Thus, when researchers aim to predict ER under unobserved conditions, our screening remains effective.

5.6 Limitations

By restricting devices to iPhones, Experiment 2 enabled comparison of the size-adjustment outcome to correct references. In both the pre-task and the main task, we controlled the visual stimuli in the mm scale, which enabled more rigorous evaluations. In the main task, increasing W levels and repetitions further allowed for precise validation of the models. However, as seen in Figures 24 and 25, particularly for the ER model, the top-left cells (threshold $T = 1$ mm, non-passing proportion $X = 0\%$) were not always the best-fitting, leaving some limitations from Experiment 1 unresolved. A plausible cause is that, in both Experiments 1 and 2, errors triggered re-aiming until success. An inattentive participant seeking to finish quickly might adopt a strategy of tapping both quickly and accurately, making it difficult to distinguish them from conscientious participants.

6 Experiment 3: iPhone-Based Experiment without Re-Aiming

6.1 Task and Design

As in Experiment 2, Experiment 3 rigorously evaluates both the size-adjustment pre-task and the pointing task on iPhones. The key




Figure 27: Cross-validation results of R^2 for ER ($N = 40$, accurate)

difference is that, in Experiment 3, the next trial begins immediately even when an error occurs (i.e., no re-aiming). This design aims to address the limitation observed in Experiments 1 and 2, where a re-aiming policy made inattentive participants difficult to distinguish from conscientious ones. Except for the handling of errors in the pointing task, the experimental design is identical to Experiment 2, and we conduct the same simulations.

6.2 Participants

A total of 575 participants completed the experiment. Each received 250 JPY (approximately 1.7 USD). The mean task completion time was 5 min 27 s, yielding an effective hourly payment of 2,752 JPY (approximately 18.8 USD). We excluded 1 participant with missing data and 55 for whom the iPhone PPI could not be identified, leaving 519 participants for analysis (250 male, 266 female, 3 other).

6.3 Results

6.3.1 Size-Adjustment Task. Figure 28 shows the distribution of absolute errors relative to the correct size in the size-adjustment task. Most participants exhibited small errors, specifically, 310 (60%) were below 2 mm, but some showed large errors, with 143 (28%) at or above 10 mm. The operation time for the adjustment averaged 8.44 s (SD=7.66 s). Five participants recorded 0 s, indicating no adjustment at all.




Figure 28: Distribution of size-adjustment errors




Figure 29: MT results for each W under the fast and accurate instruction conditions

6.3.2 Pointing Task. Applying the same outlier handling as in Experiment 2 (Section 5.3.2), we excluded 1,011 trials for tap coordinates, 2,717 trials for MT, and four participants for MT. In total, 181,796 trials (97.3%) were retained for analysis. Figures 29 and 30 present the pointing task results for Experiment 3: MT by ($W \times$ instruction) and ER by ($W \times$ instruction), respectively. For MT, we found significant main effects of instruction ($F_{1,514} = 205.8, p < 0.001, \eta_g^2 = 0.037$) and W ($F_{8,4112} = 252.2, p < 0.001, \eta_g^2 = 0.011$). The interaction effect of $W \times$ instruction was significant ($F_{8,4112} = 32.79, p < 0.001, \eta_g^2 = 0.00063$). For ER, we found significant main effects of instruction ($F_{1,514} = 67.46, p < 0.001, \eta_g^2 = 0.012$) and W ($F_{8,4112} = 1396, p < 0.001, \eta_g^2 = 0.27$). The interaction effect of $W \times$ instruction was significant ($F_{8,4112} = 18.61, p < 0.001, \eta_g^2 = 0.0026$). As in Experiment 2, ER was more strongly affected by W than MT; for example, under the fast condition, when W increased from 2.0 to 8.4 mm, the change ratio of ER was 86% while that of MT was 8.8%. In addition, relative to Figure 20, MT was slightly shorter and ER slightly higher.

6.4 Simulation

The simulation conditions and procedures followed Experiment 2 (Section 5.4). Figure 31 shows the numbers of passing vs. non-passing participants as a function of the threshold T in Experiment 3.




Figure 31: Numbers of passing vs. non-passing participants as a function of the threshold




Figure 30: ER results for each W under the fast and accurate instruction conditions

6.5 Simulation Results

6.5.1 Movement Time (MT). Figures 32 and 33 present R^2 for Equation (1) in Experiment 3 as T and X vary. In both figures, R^2 decreases from the upper-left to the lower-right of each heatmap: lenient thresholds combined with larger non-passing proportions lower model fit. Thus, the degradation observed in Experiments 1 and 2 is replicated in Experiment 3.

Compared with Experiment 2, the magnitude of degradation with increasing non-passing proportion is somewhat larger in Experiment 3. Specifically, in Figures 22 and 23, even when fit decreased, most cells remained above 0.88 for $N = 10$ or $N = 20$. By contrast, in Figures 32 and 33, values sometimes dropped to around 0.65 for $N = 10$ and around 0.78 for $N = 20$. This difference should stem from the change introduced in Experiment 3, i.e., skipping re-aiming and proceeding immediately after errors. This design mitigates the prior limitation where inattentive participants behaved similarly to conscientious (rapid and accurate) ones. Consequently, changes in model fit are more pronounced.

As before, increasing N improves R^2 in the upper-left regions and attenuates the effect of larger X . Nevertheless, R^2 still varies with X in Figures 32 and 33, and the declines are larger in Experiment 3. Thus, in settings where participants can complete tasks without following instructions (as here), our screening method becomes even more valuable.

As a cross-validation example, Figure 36 shows the case $N = 40$ under the “accurate” instruction. Similar to the results in Experiment 2, compared to using all W levels, overall predictive accuracy is slightly lower, yet the same upper-left to lower-right degradation trend appears.

6.5.2 Error Rate (ER). Figures 34 and 35 show R^2 for Equation (5) in Experiment 3 as T and X vary. Again, R^2 decreases from the upper-left to the lower-right in each heatmap: lenient thresholds and larger non-passing proportions reduce fit. Thus, the degradation observed in Experiments 1 and 2 is also observed here. The influence of N follows the same trend described above.

Compared to Experiment 2, the declines in R^2 with increasing non-passing proportion are somewhat larger in Experiment 3. In Figures 24 and 25, most cells remained above 0.93 even when fit decreased; in Figures 34 and 35, values sometimes dropped to around 0.85 for $N = 10$ or $N = 20$. Moreover, unlike Figures 24 and 25, Experiment 3 shows a tendency for fit to decline even when the

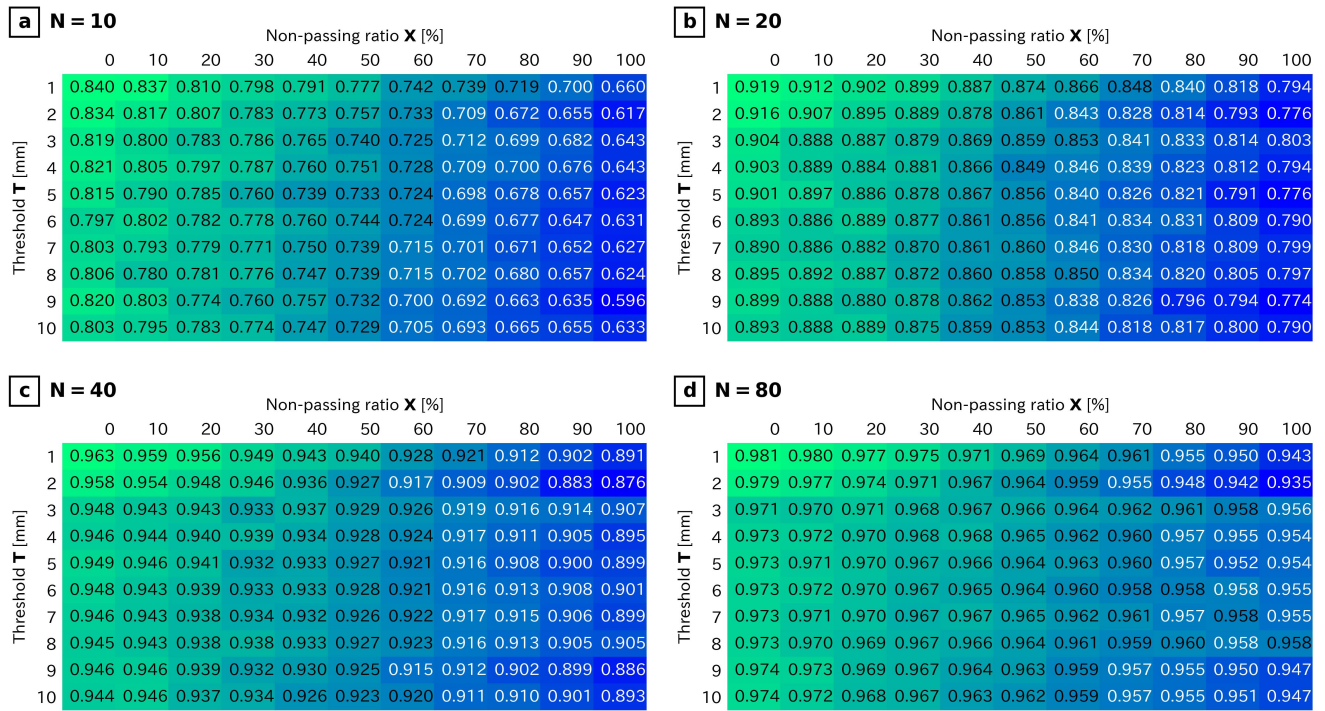


Figure 32: Experiment 3: Goodness of fit to Equation (1) for MT (fast)

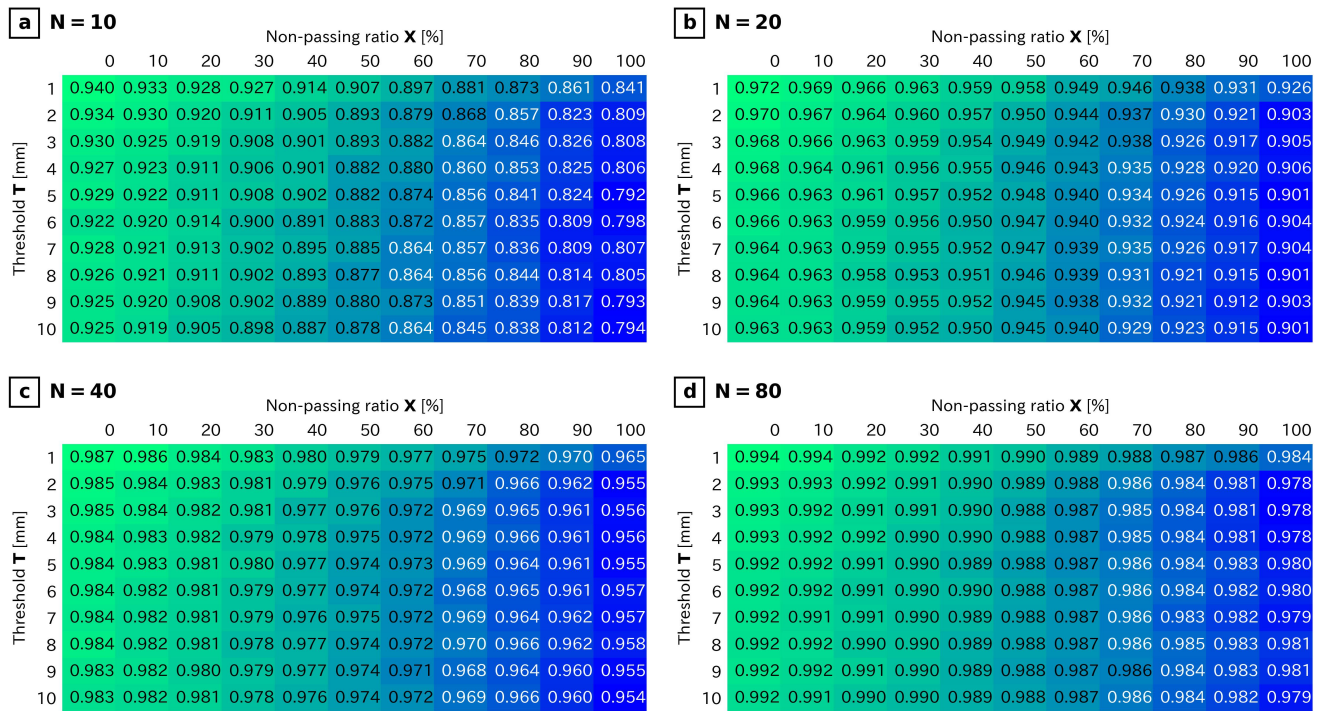


Figure 33: Experiment 3: Goodness of fit to Equation (1) for MT (accurate)

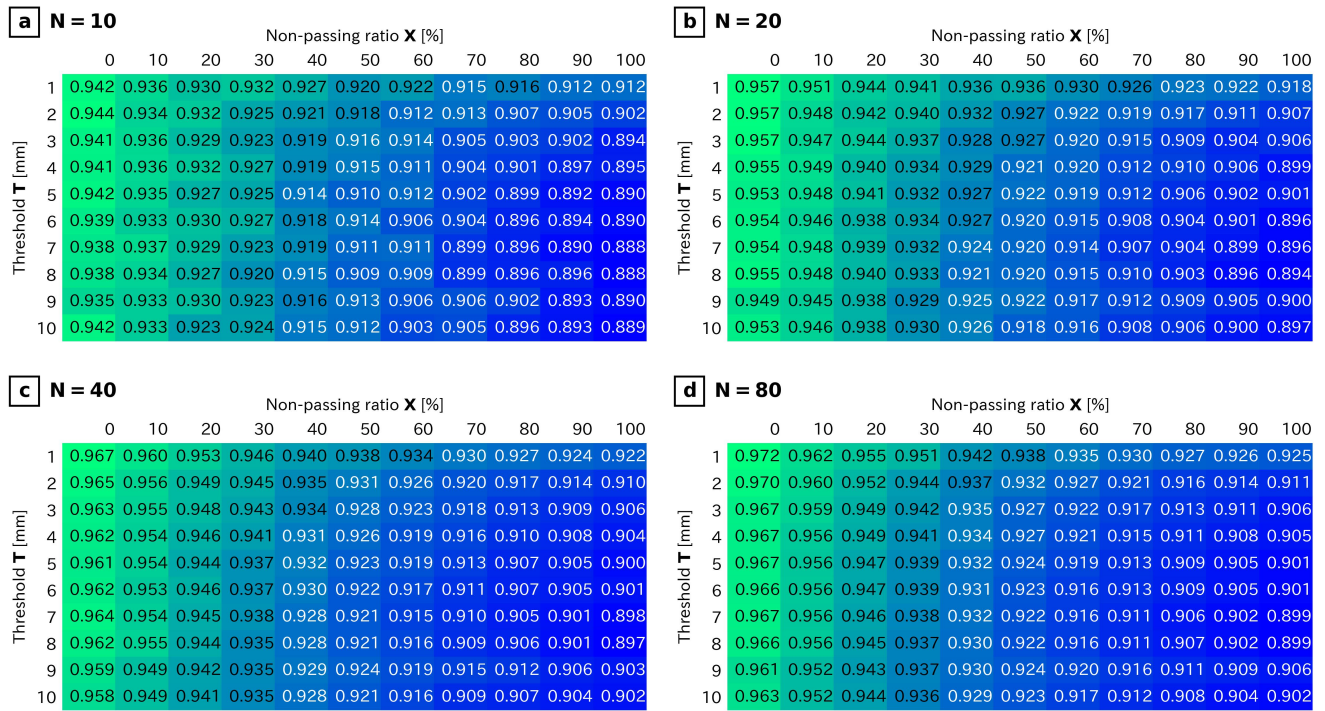


Figure 34: Experiment 3: Goodness of fit to Equation (5) for ER (fast)

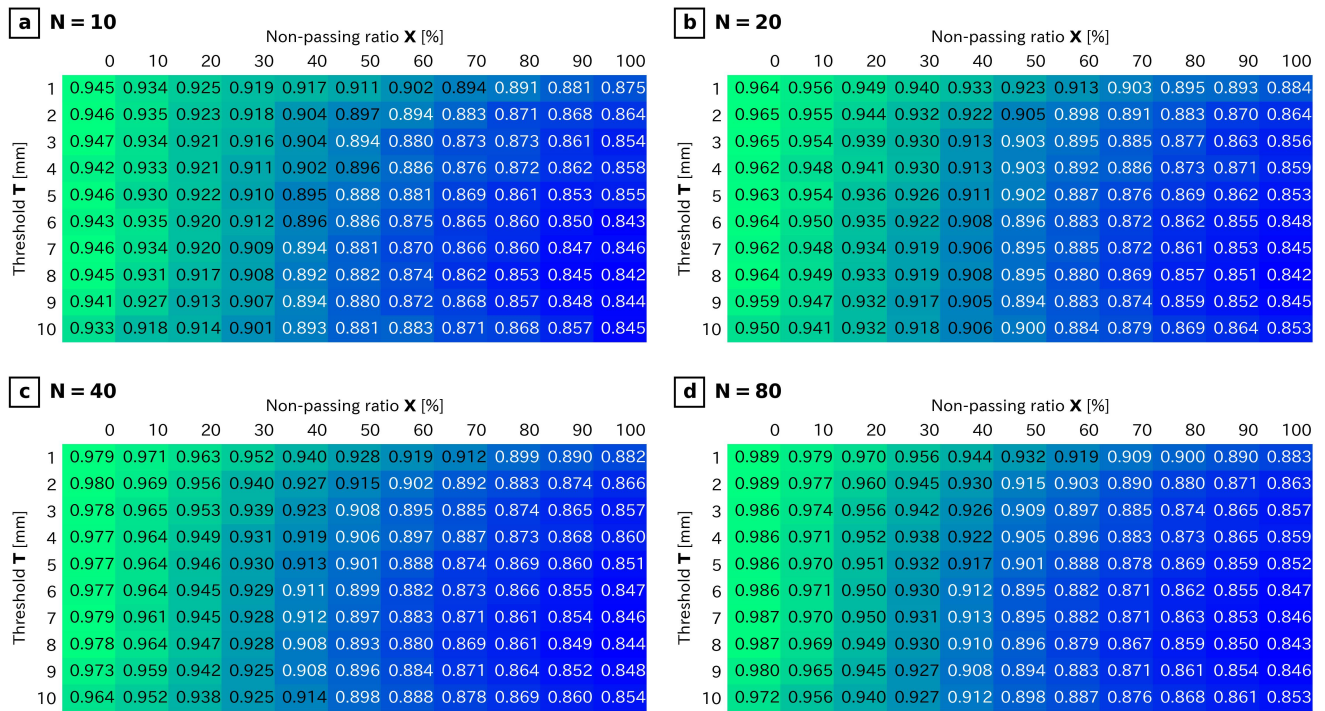


Figure 35: Experiment 3: Goodness of fit to Equation (5) for ER (accurate)




Figure 36: Cross-validation results of R^2 for MT ($N = 40$, accurate)

non-passing proportion is small. Another difference is that, in most heatmaps, the top-left cell (threshold $T = 1$, non-passing proportion $X = 0$) attains the best fit. Whereas Experiment 2 exhibited a limitation in which this cell was not always optimal (e.g., Figure 25), Experiment 3 generally achieves the best fit there (Figures 34–35), and this is also true for the MT results (Figures 32–33).

A likely reason is, again, the immediate progression after errors in Experiment 3. Inattentive participants may accept higher error rates to act quickly, deviating from the pointing behavior assumed by the ER model. This amplifies the effects of N and X on model fit and yields the best performance at the top-left cell (strict threshold, no non-passing participants).

As a cross-validation example, Figure 37 shows the case $N = 40$ under the “accurate” instruction. As with Figure 36, overall predictive accuracy is slightly lower than when using all W levels, but the same upper-left to lower-right degradation trend holds.

7 DISCUSSION

7.1 RQ1: Can a simple GUI pre-task effectively screen out nonconforming participants?

Across three experiments using the size-adjustment task as a pre-task and a pointing task as the main experiment, we found that the goodness of fit of models estimating GUI performance (movement time and error rate) changes with the proportion of non-passing participants mixed into the sample (e.g., Figure 35). Consistently, keeping the proportion X of the non-passing group small (suppressing the inclusion of nonconforming participants) and setting a stricter error threshold T in the size-adjustment task tended to improve model fit. In other words, screening based on size-adjustment error contributes to higher goodness of fit for the performance models in the main experiment.

A key advantage is that nonconformity is judged without using outcomes from the main experiment, as we selected participants solely based on the pre-task UI interaction. This elevates the quality of those who proceed to the main experiment and reduces the risk of incorrect conclusions in model evaluation (e.g., that a given model poorly estimates UI performance). Moreover, the fact that screening can be achieved with a simple UI operation within nine seconds on average supports ease of adoption and practical utility for researchers. Taken together, size-adjustment-based screening

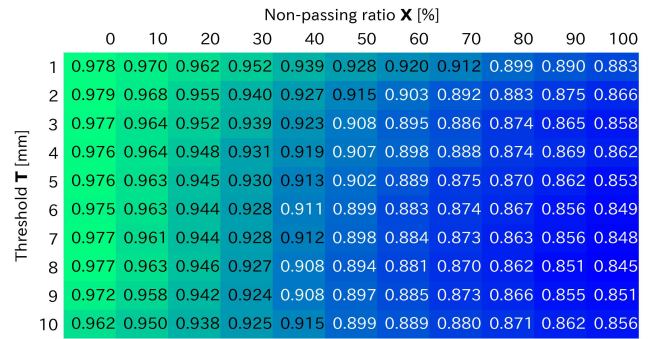


Figure 37: Cross-validation results of R^2 for ER ($N = 40$, accurate)

is an effective method to enhance data quality by improving the model fit of the main experiment.

Overall, the answer to RQ1 is: *a brief GUI pre-task—operationalized here as size adjustment—effectively screens out nonconforming participants and thereby improves the goodness of fit of performance models in the main experiment.*

7.2 RQ2: How do the non-passing proportion and the error threshold affect model fit and predictive accuracy in the main experiment?

We systematically manipulated the size-adjustment error threshold T and the non-passing proportion X and evaluated their effects on the fit and predictive accuracy of the MT and ER models. Overall, stricter T (smaller allowable error) and smaller X tended to yield higher R^2 .

In the PC mouse setting (Experiment 1), the MT model’s R^2 was stably above 0.98 and only weakly sensitive to T or X (Figures 10 and 11). By contrast, for the ER model under the “accurate” instruction, the heatmap showed a clear decline in R^2 as T was relaxed and X increased (Figure 13).

When restricting to smartphones (Experiment 2), declines in fit with relaxed T and larger X were observed not only for the ER model but also for the MT model (Figures 22, 23, 24, and 25). For unobserved target-width conditions, the same degradation trends were reproduced by leave-one- W -out cross-validation (e.g., Figures 26 and 27). Experiment 3 replicated these tendencies as tightening T and suppressing X improved fit (Figures 32, 33, 34, and 35).

We also observed cases where the seemingly ideal setting for experimenters (the smallest T with $X = 0$) did not always yield the highest R^2 . This suggests that design factors, such as the number of W levels, repetitions, and the error-handling policy (presence/absence of re-aiming), can affect evaluation resolution and warrant consideration in experimental design. In sum, stricter T and smaller X improve fit and predictive accuracy, especially for the ER model, and in smartphone settings the MT model benefits as well.

Overall, the answer to RQ2 is: *stricter thresholds T and smaller non-passing proportions X systematically increase model fit R^2 (and, where assessed, predictive accuracy), with especially strong effects for ER and clear benefits for MT in smartphone settings.*




Figure 38: Comparison of *ER* model fit heatmaps across Experiments 1–3 ($N = 40$, accurate)

7.3 RQ3: Is the proposed screening consistently effective across devices and error-handling procedures?

We evaluated the same pre-task-based screening method across three experiments spanning PC (mouse) and smartphone (touch). Comparing Experiment 1 (PC) with Experiments 2 and 3 (smartphone), we observed a common pattern that reducing the non-passing proportion X and tightening the error threshold T increased R^2 . As examples of device comparisons, see Figures 13, 25, and 35 for the *ER* results under the “accurate” instruction. Improvements were especially pronounced for the *ER* model and extended to the *MT* model in smartphone settings (e.g., Figures 22 and 32). To streamline figure references, Figure 38 summarizes the *ER* results across experiments.

A direct comparison of the two smartphone procedures, Experiment 2 (re-aim until success) and Experiment 3 (proceed immediately after errors), reveals similar improvement gradients with respect to T and X , indicating robustness across error-handling policies. This effect is also captured compactly in Figure 38. That is, across device differences and the presence/absence of re-aiming, stricter T and smaller X consistently yield higher fit. In addition, from Figure 38, we see that Experiments 1–2 (PC/iPhone with re-aiming) show only slight reductions in R^2 as T is relaxed and the non-passing proportion X grows, whereas Experiment 3 (no re-aiming) exhibits a marked, monotonic decline. Therefore, selecting a stricter T and keeping X small is especially crucial when errors are not corrected via re-aiming.

Even under smartphone conditions where target width W was controlled in mm via device PPI identification, we still observed fit improvements with stricter T and smaller X , supporting generalizability in practical deployments. At the same time, the best-looking setting (smallest T , $X = 0$) did not always produce the highest R^2 . Therefore, sufficient numbers of W levels and repetitions and the choice of error-handling can influence evaluation resolution and should be considered.

Overall, the answer to RQ3 is: *the method is consistently effective across device differences and across re-aiming procedures.*

7.4 Limitations and Future Work

Our screening cannot identify participants who act in a task-conforming manner in the pre-task yet unexpectedly in the main task. Thus, even after screening, the remaining sample may still include participants whose behavior in the main task is inattentive or only partially compliant with the instructions. Moreover, our approach requires setting a threshold T for the size-adjustment task. Because model fit and predictive accuracy change gradually with T and with the non-passing proportion X , uniquely determining the best T is difficult. A stricter T reduces the inclusion of nonconforming participants but induces a smaller sample size and thus leads to imprecise model evaluation. Further, design factors such as the error-handling policy (presence/absence of re-aiming), the number of W levels, and the number of repetitions may contribute to the phenomenon where the “best-looking” condition does not always yield the best fit.

Another limitation concerns potential bias introduced by the proposed screening method. Participants with motor or visual impairments may find it inherently difficult to achieve small resizing errors even when they follow the instructions conscientiously. Whether such participants should be excluded depends on the purpose of the experiment. In our study, the goal is to evaluate the estimation accuracy of GUI performance models under standardized task conditions. Hence, substantial motor or perceptual limitations would directly influence these models’ accuracy and complicate comparisons across participants, so excluding such cases is consistent with our objectives. However, for studies aiming to observe diverse real-world GUI behaviors or to investigate accessibility-related differences, removing these participants would suppress meaningful variability and could be undesirable. In addition, some apparent “nonconformity” may stem from unclear instructions or suboptimal interface design rather than worker intent, suggesting that improvements in task design and pre-task screening should be considered jointly.

Our findings and discussion are limited to the case where the pre-task is image resizing and the main task is pointing. Both tasks are relatively simple GUI operations with modest cognitive and motor demands and share similar low-level control characteristics (e.g., precise positioning on a 2D display), which may limit the extent to

which our findings directly generalize to more complex interaction scenarios. Whether the same screening idea works when one or both tasks differ requires further investigation. Future work will thus test the screening approach with other combinations of pre-tasks and main tasks. One conceptually interesting configuration would be to invert the roles of pre-task and main task, i.e., using a pointing task as the pre-task and a resizing task as the main task. If participants who yield low R^2 in pointing performance also tend to exhibit large resizing errors, this would support the idea that operational carefulness transfers across tasks. At the same time, prior work indicates that pointing-based performance models, especially error-rate models, can produce low R^2 even for conscientious participants when the number of trials per condition is small [48], and that people with motor-control difficulties may show high error rates despite carefully following instructions [15]. Therefore, interpreting a poor model fit in reversed configurations as straightforward indicators of “inattention” would require caution. We also examined using the number of re-aiming attempts in the main task as a direct indicator of careless behavior, but we did not observe a clear correlation with pre-task errors (Pearson’s $r < 0.2$; see the supplementary materials).

Across our three experiments, we consistently observed that stricter T and smaller X increased R^2 , and that appropriate screening was possible using only pre-task UI interaction data, without using outcomes from the main experiment to judge nonconformity. These results suggest that effective screening may be achievable even when the pre-task differs from the main task. Thus, researchers may not need to change the pre-task every time the main task changes. Rather, screening should transfer to other GUI-based main experiments, such as those to evaluate key-typing [41, 59] and path-steering performance [1, 25].

Alternative pre-tasks are also conceivable. Potential pre-tasks would be brief, require simple GUI-based interactions, and share operational characteristics with the main task. We will explore such pre-tasks for more effective screening. Another possible approach would be to employ a gold-standard task [26] or attention check [34] as a pre-task, thereby restricting the main task to workers who pass it. Still, we demonstrated that the results of pre-tasks, especially measures of operational carefulness (size-adjustment error), are in fact distributed continuously rather than dichotomously. Because gold tasks and attention checks are often used to make an explicit pass/fail decision of conformity (even though their outcomes can also be treated as continuous measures such as accuracy or response time), the advantage of our proposed method is that the pre-task is designed to yield a task-relevant continuous error signal and thus account for this gradient. Future work should empirically compare the screening capabilities of these candidate pre-task mechanisms to clarify their respective advantages and appropriate use cases in GUI experimentation.

7.5 Positioning within HCI and Suitable Downstream (Main) Tasks

From a broader HCI perspective, our work contributes to the methodological and infrastructural foundations that support empirical research on GUI interaction. As crowdsourcing becomes increasingly common for collecting large-scale behavioral data in tasks

involving pointing, selection, movement, or text entry, ensuring that the collected performance data are reliable is essential. Our screening method complements prior infrastructural contributions in HCI, such as refined statistical analysis procedures [12, 46] and standardized GUI performance metrics [25, 47], by providing a practical mechanism for improving data quality *before* conducting the downstream main experiment. Rather than proposing a new interaction technique, this work strengthens the basis on which such techniques, as well as comparative evaluations of existing models, can be reliably assessed in crowdsourced settings.

The proposed screening method is also broadly applicable to downstream GUI-interaction tasks beyond the pointing paradigm used in our experiments. Relevant examples include goal crossing [2, 52], path steering [1, 28], dragging [22, 32], typing [3, 40], and object selection techniques [5, 57]. These tasks all rely on quantitative performance measurements (e.g., speed, accuracy, error patterns) and are sensitive to noisy or inattentive behavior in crowdsourced environments. By screening out participants whose pre-task behavior indicates insufficient task compliance, our method can improve the reliability of collected data, thereby supporting the validity of experimental results in model fitting, performance comparisons, and replications.

8 CONCLUSION

This work introduced a pre-experiment screening based on a size-adjustment task and validated its effectiveness across three experiments spanning PCs (mouse) and smartphones (touch). For RQ1, we showed that selecting participants solely from pre-task UI outcomes (size-adjustment error) consistently improved the goodness of fit of performance models in the main experiment. In particular, the smaller the inclusion of the non-passing group and the stricter the threshold T , the higher the fit.

For RQ2, manipulating the threshold T and the ratio of inclusion of nonconforming participants X had systematic effects on model fit R^2 (and, in Experiments 2 and 3, on predictive accuracy via LOOCV), demonstrating that stricter T and smaller X are especially effective for the *ER* model, with improvements extending to the *MT* model in smartphone settings. For RQ3, we observed the same pattern of improvements across device differences (PC/smartphone) and error-handling procedures (re-aim until success vs. proceed immediately after errors), confirming the robustness of the proposed method.

These results indicate that a brief pre-task finished in nine seconds on average can enhance participant quality *before* the main experiment and reduce the risk of erroneous conclusions in model evaluation. In practice, a conservative (stricter) choice of T and keeping X small should be the default policy, while also ensuring sufficient sample size N , number of target-width levels W , and repetitions. On smartphones, the effect was confirmed even when stimulus size was controlled in millimeters via device PPI identification, supporting generalizability to real deployments. At the same time, uniquely determining T is nontrivial: there are trade-offs with X and participant recruitment, and residual influences from error-handling and trial design remain. Developing guidance tailored to study goals and constraints, and optimizing protocols accordingly, are important directions.

References

- [1] Johnny Accot and Shumin Zhai. 1997. Beyond Fitts' law: models for trajectory-based HCI tasks. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) (CHI '97). Association for Computing Machinery, New York, NY, USA, 295–302. doi:10.1145/258549.258760
- [2] Johnny Accot and Shumin Zhai. 2002. More than dotting the i's – foundations for crossing-based interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Minneapolis, Minnesota, USA) (CHI '02). Association for Computing Machinery, New York, NY, USA, 73–80. doi:10.1145/503376.503390
- [3] Nikola Banovic, Ticha Sethapakdi, Yasasvi Hari, Anind K Dey, and Jennifer Mankoff. 2019. The limits of expert text entry speed on mobile keyboards with autocorrect. In *Proceedings of the 21st International Conference on Human-Computer Interaction with Mobile Devices and Services*. 1–12.
- [4] Xiaojun Bi and Shumin Zhai. 2016. Predicting finger-touch accuracy based on the dual Gaussian distribution model. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. 313–319.
- [5] Per Bjerre, Allan Christensen, Andreas Køllund Pedersen, Simon André Pedersen, Wolfgang Stuerzlinger, and Rasmus Stenholt. 2017. Predictive Model for Group Selection Performance on Touch Devices. In *Human-Computer Interaction. Interaction Contexts - 19th International Conference, HCI International 2017, Vancouver, BC, Canada, July 9-14, 2017, Proceedings, Part II*. Springer International Publishing, Cham, 142–161. doi:10.1007/978-3-319-58077-7_12
- [6] Maria Blanca, Rafael Alarcon, Jaume Arnau, Roser Bono, and Rebecca Bendayan. 2017. Non-normal data: Is ANOVA still a valid option? *Psicothema* 29 (10 2017), 552–557. doi:10.7334/psicothema2016.383
- [7] Florian Brühlmann, Serge Petralito, Lena F Aeschbach, and Klaus Opwis. 2020. The quality of data collected online: An investigation of careless responding in a crowdsourced sample. *Methods in Psychology* 2 (2020), 100022.
- [8] Géry Casiez, Nicolas Roussel, Romuald Vanbellegem, and Frédéric Giraud. 2011. Surfpad: Riding towards targets on a squeeze film effect. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems - CHI '11*. ACM, New York, 2491–2500. doi:10.1145/1978942.1979307
- [9] Matthew JC Crump, John V McDonnell, and Todd M Gureckis. 2013. Evaluating Amazon's Mechanical Turk as a tool for experimental behavioral research. *PLoS one* 8, 3 (2013), e57410.
- [10] Paul G Curran. 2016. Methods for the detection of carelessly invalid responses in survey data. *Journal of Experimental Social Psychology* 66 (2016), 4–19.
- [11] Joshua R De Leeuw. 2015. jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior research methods* 47 (2015), 1–12.
- [12] Lisa A. Elkin, Matthew Kay, James J. Higgins, and Jacob O. Wobbrock. 2021. An Aligned Rank Transform Procedure for Multifactor Contrast Tests. In *The 34th Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) (UIST '21). Association for Computing Machinery, New York, NY, USA, 754–768. doi:10.1145/3472749.3474784
- [13] Leah Findlater, Jon E. Froehlich, Kays Fattal, Jacob O. Wobbrock, and Tanya Dast-yar. 2013. Age-Related Differences in Performance with Touchscreens Compared to Traditional Mouse Input. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) (CHI '13). Association for Computing Machinery, New York, NY, USA, 343–346. doi:10.1145/2470654.2470703
- [14] Leah Findlater, Joan Zhang, Jon E Froehlich, and Karyn Moffatt. 2017. Differences in crowdsourced vs. lab-based mobile and desktop input performance data. In *Proceedings of the 2017 CHI conference on human factors in computing systems*. 6813–6824.
- [15] Leah Findlater and Lotus Zhang. 2020. Input Accessibility: A Large Dataset and Summary Analysis of Age, Motor Ability and Input Performance. In *Proceedings of the 22nd International ACM SIGACCESS Conference on Computers and Accessibility* (Virtual Event, Greece) (ASSETS '20). Association for Computing Machinery, New York, NY, USA, Article 17, 6 pages. doi:10.1145/3373625.3417031
- [16] Paul M Fitts. 1954. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of experimental psychology* 47, 6 (1954), 381.
- [17] Clifton Forlines and Ravin Balakrishnan. 2008. Evaluating Tactile Feedback and Direct vs. Indirect Stylus Input in Pointing and Crossing Selection Tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy) (CHI '08). ACM, New York, NY, USA, 1563–1572. doi:10.1145/1357054.1357299
- [18] Jaroslav Gottfried. 2024. Practices in Data-Quality Evaluation: A Large-Scale Review of Online Survey Studies Published in 2022. *Advances in Methods and Practices in Psychological Science* 7, 2 (2024).
- [19] Tovi Grossman and Ravin Balakrishnan. 2005. The bubble cursor: enhancing target acquisition by dynamic resizing of the cursor's activation area. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 281–290.
- [20] Tovi Grossman and Ravin Balakrishnan. 2005. The Bubble Cursor: Enhancing Target Acquisition by Dynamic Resizing of the Cursor's Activation Area. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Portland, Oregon, USA) (CHI '05). ACM, New York, NY, USA, 281–290. doi:10.1145/1054972.1055012
- [21] John J Horton, David G Rand, and Richard J Zeckhauser. 2011. The online laboratory: Conducting experiments in a real labor market. *Experimental economics* 14 (2011), 399–425.
- [22] Kori M. Inkpen. 2001. Drag-and-Drop versus Point-and-Click Mouse Interaction Styles for Children. *ACM Trans. Comput.-Hum. Interact.* 8, 1 (March 2001), 1–33. doi:10.1145/371127.371146
- [23] ISO. 2012. 9241-411:2012. Ergonomics of human-system interaction-Part 411: Evaluation methods for the design of physical input devices.
- [24] ISO. 2019. ISO/IEC 7810:2019. Identification cards-Physical characteristics.
- [25] Nobuhito Kasahara, Yosuke Oba, Shota Yamanaka, Anil Ufuk Batmaz, Wolfgang Stuerzlinger, and Homei Miyashita. 2024. Better definition and calculation of throughput and effective parameters for steering to account for subjective speed-accuracy tradeoffs. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 722, 18 pages. doi:10.1145/3613904.3642084
- [26] Gabriella Kazai, Jaap Kamps, Marijn Koolen, and Natasa Milic-Frayling. 2011. Crowdsourcing for book search evaluation: impact of hit design on comparative system ranking. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Beijing, China) (SIGIR '11). Association for Computing Machinery, New York, NY, USA, 205–214. doi:10.1145/2009916.2009947
- [27] Steven Komarov, Katharina Reinecke, and Krzysztof Z Gajos. 2013. Crowdsourcing performance evaluations of user interfaces. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 207–216.
- [28] Sergey Kulikov, I. Scott MacKenzie, and Wolfgang Stuerzlinger. 2005. Measuring the Effective Parameters of Steering Motions. In *CHI '05 Extended Abstracts on Human Factors in Computing Systems* (Portland, OR, USA) (CHI EA '05). ACM, New York, NY, USA, 1569–1572. doi:10.1145/1056808.1056968
- [29] Qisheng Li, Sung Jun Joo, Jason D Yeatman, and Katharina Reinecke. 2020. Controlling for participants' viewing distance in large-scale, psychophysical online experiments using a virtual chinrest. *Scientific reports* 10, 1 (2020), 904.
- [30] LY Corporation. 2025. Yahoo! crowdsourcing. <https://crowdsourcing.yahoo.co.jp>
- [31] I Scott MacKenzie. 1992. Fitts' law as a research and design tool in human-computer interaction. *Human-computer interaction* 7, 1 (1992), 91–139.
- [32] I Scott MacKenzie, Abigail Sellen, and William AS Buxton. 1991. A comparison of input devices in element pointing and dragging tasks. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 161–166.
- [33] David E Meyer, Richard A Abrams, Sylvan Kornblum, Charles E Wright, and JE Keith Smith. 1988. Optimality in human motor performance: ideal control of rapid aimed movements. *Psychological review* 95, 3 (1988), 340.
- [34] Daniel M. Oppenheimer, Tom Meyvis, and Nicolas Davidenko. 2009. Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology* 45, 4 (2009), 867–872. doi:10.1016/j.jesp.2009.03.009
- [35] Ulf Dietrich Reips. 2002. Standards for Internet-based experimenting. *Experimental psychology* 49, 4 (2002), 243.
- [36] Emanuel Schmider, Matthias Ziegler, Erik Danay, Luzi Beyer, and Markus Bühner. 2010. Is it really robust? Reinvestigating the robustness of ANOVA against the normal distribution assumption. *Methodology* 6, 4 (2010), 147–151. doi:10.1027/1614-2241/a000016
- [37] Michail Schwab, Sicheng Hao, Olga Vitek, James Tompkin, Jeff Huang, and Michelle A Borkin. 2019. Evaluating pan and zoom timelines and sliders. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–12.
- [38] Katie Seaborn and Satoshi Nakamura. 2025. Quality and representativeness of research online with Yahoo! Crowdsourcing. *Frontiers in Psychology* Volume 16 - 2025 (2025). doi:10.3389/fpsyg.2025.1588579
- [39] Ather Sharif, Victoria Pao, Katharina Reinecke, and Jacob O Wobbrock. 2020. The reliability of Fitts's law as a movement model for people with and without limited fine motor function. In *Proceedings of the 22nd international Acm Sigaccess conference on computers and accessibility*. 1–15.
- [40] R. William Soukoreff and I. Scott MacKenzie. 1995. Theoretical upper and lower bounds on typing speed using a stylus and a soft keyboard. *Behaviour & Information Technology* 14, 6 (1995), 370–379. doi:10.1080/01449299508914656
- [41] R. William Soukoreff and I. Scott MacKenzie. 2003. Metrics for text entry research: An evaluation of MSD and KSPC, and a new unified error metric. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems - CHI '03*. ACM, New York, 113–120. doi:10.1145/642611.642632
- [42] R William Soukoreff and I Scott MacKenzie. 2004. Towards a standard for pointing device evaluation, perspectives on 27 years of Fitts' law research in HCI. *International journal of human-computer studies* 61, 6 (2004), 751–789.
- [43] Theophanis Tsandilas and Géry Casiez. 2024. The illusory promise of the Aligned Rank Transform—A systematic study of rank transformations. <https://statransform.github.io/jovi>
- [44] Hiroki Usuba, Shota Yamanaka, Junichi Sato, and Homei Miyashita. 2022. Predicting touch accuracy for rectangular targets by using One-Dimensional task results. *Proc. ACM Hum.-Comput. Interact.* 6, ISS, Article 579 (nov 2022), 13 pages.

- doi:10.1145/3567732
- [45] Jacob O Wobbrock, Edward Cutrell, Susumu Harada, and I Scott MacKenzie. 2008. An error model for pointing based on Fitts' law. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1613–1622.
- [46] Jacob O Wobbrock, Leah Findlater, Darren Gergle, and James J Higgins. 2011. The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 143–146.
- [47] Jacob O. Wobbrock, Kristen Shinohara, and Alex Jansen. 2011. The Effects of Task Dimensionality, Endpoint Deviation, Throughput Calculation, and Experiment Design on Pointing Measures and Models. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Vancouver, BC, Canada) (CHI '11)*. ACM, New York, NY, USA, 1639–1648. doi:10.1145/1978942.1979181
- [48] Shota Yamanaka. 2021. Utility of crowdsourced user experiments for measuring the central tendency of user performance to evaluate error-rate models on gus. In *Proceedings of the AAI Conference on Human Computation and Crowdsourcing*, Vol. 9. 155–165.
- [49] Shota Yamanaka. 2022. Test-Retest Reliability on Movement Times and Error Rates in Target Pointing. In *Proceedings of the 2022 ACM Designing Interactive Systems Conference*. 178–188.
- [50] Shota Yamanaka. 2025. Relative merits of nominal and effective indexes of difficulty of Fitts' law: Effects of sample size and the number of repetitions on model fit. *International Journal of Human-Computer Interaction* 41, 1 (2025), 574–591. doi:10.1080/10447318.2024.2303201
- [51] Shota Yamanaka, Taiki Kinoshita, Yosuke Oba, Ryuto Tomihari, and Homei Miyashita. 2023. Varying subjective speed-accuracy biases to evaluate the generalizability of experimental conclusions on pointing-facilitation techniques. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [52] Shota Yamanaka and Wolfgang Stuerzlinger. 2020. Necessary and unnecessary distractor avoidance movements affect user behaviors in crossing operations. *ACM Transactions on Computer-Human Interaction (TOCHI)* 27, 6 (2020), 1–31.
- [53] Shota Yamanaka and Hiroki Usuba. 2020. Rethinking the dual gaussian distribution model for predicting touch accuracy in on-screen-start pointing tasks. *Proceedings of the ACM on Human-Computer Interaction* 4, ISS (2020), 1–20.
- [54] Shota Yamanaka and Hiroki Usuba. 2023. Tuning Endpoint-variability Parameters by Observed Error Rates to Obtain Better Prediction Accuracy of Pointing Misses. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 579, 18 pages. doi:10.1145/3544548.3580746
- [55] Shota Yamanaka and Hiroki Usuba. 2024. 0.2-mm-step verification of the dual gaussian distribution model with large sample size for predicting tap success rates. *Proc. ACM Hum.-Comput. Interact.* 8, ISS, Article 553 (October 2024), 20 pages. doi:10.1145/3698153
- [56] Shota Yamanaka, Hiroki Usuba, and Homei Miyashita. 2022. Bivariate effective width method to improve the normalization capability for subjective speed-accuracy biases in rectangular-target pointing. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems – CHI '22*. ACM, New York, Article 211, 13 pages. doi:10.1145/3491102.3517466
- [57] Shota Yamanaka, Hiroki Usuba, Wolfgang Stuerzlinger, and Homei Miyashita. 2022. The Effectiveness of Path-Segmentation for Modeling Lasso Times in Width-Varying Paths. *Proceedings of the ACM on Human-Computer Interaction* 6, ISS (2022), 640–659.
- [58] Shumin Zhai, Jing Kong, and Xiangshi Ren. 2004. Speed-accuracy tradeoff in Fitts' law tasks: On the equivalency of actual and nominal pointing precision. *International Journal of Human-Computer Studies* 61, 6 (2004), 823–856. doi:10.1016/j.ijhcs.2004.09.007
- [59] Mingrui Ray Zhang, Shumin Zhai, and Jacob O. Wobbrock. 2019. Text entry throughput: Towards unifying speed and accuracy in a single performance metric. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems – CHI '19*. ACM, New York, 1–13. doi:10.1145/3290605.3300866