Reflow: Automatically Improving Touch Interactions in Mobile Applications through Pixel-based Refinements

JASON WU*, HCI Institute, Carnegie Mellon University, USA

TITUS BARIK, XIAOYI ZHANG, COLIN LEA, JEFFREY NICHOLS, JEFFREY P. BIGHAM, Apple, USA
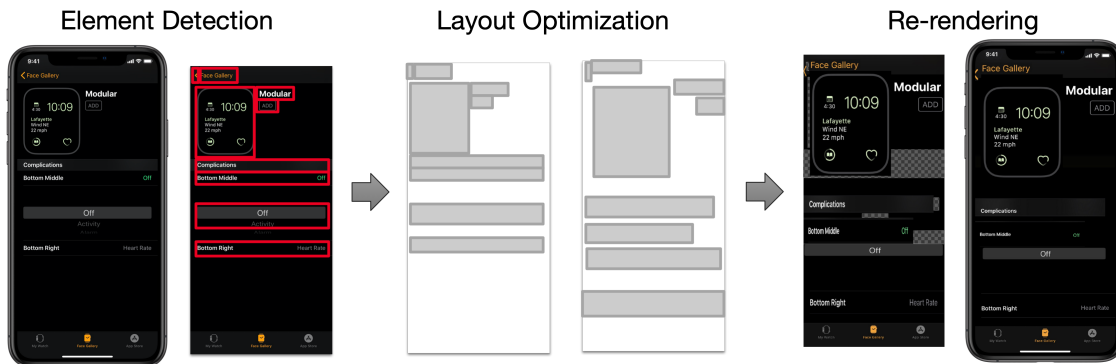
Fig. 1. Reflow *(i)* detects UI elements from pixels, *(ii)* optimizes the UI layout for a personalized difficulty model, and then *(iii)* re-renders the visual UI with the new layout. In this example, the UI elements are correctly detected, and the new layout includes larger buttons that are more spread apart. Simply moving and stretching the UI elements causes gaps and distortions, which Reflow fixes using additional post-processing methods.

Touch is the primary way that users interact with smartphones. However, building mobile user interfaces where touch interactions work well for all users is a difficult problem, because users have different abilities and preferences. We propose a system, Reflow, which automatically applies small, personalized UI adaptations, called *refinements*—to mobile app screens to improve touch efficiency. Reflow uses a pixel-based strategy to work with existing applications, and improves touch efficiency while minimally disrupting the design intent of the original application. Our system optimizes a UI by *(i)* extracting its layout from its screenshot, *(ii)* refining its layout, and *(iii)* re-rendering the UI to reflect these modifications. We conducted a user study with 10 participants and a heuristic evaluation with 6 experts and found that applications optimized by Reflow led to, on average, 9% faster selection time with minimal layout disruption. The results demonstrate that Reflow's refinements useful UI adaptations to improve touch interactions.

---

*This work was done while Jason Wu was an intern at Apple.

Authors' addresses: Jason Wu, HCI Institute, Carnegie Mellon University, Pittsburgh, PA, USA, jsonwu@cmu.edu; Titus Barik, Xiaoyi Zhang, Colin Lea, Jeffrey Nichols, Jeffrey P. Bigham, Apple, Cupertino, CA, USA, {tbarik,xiaoyiz,colin_lea,jwnichols,jbigham}@apple.com.

---

arXiv:2207.07712v1 [cs.HC] 15 Jul 2022

1   INTRODUCTION

Touch is an ubiquitous way of interacting with smartphones. However, building mobile user interfaces (UI) where touch interactions work well for all users is a difficult problem, because users have different motor abilities, skills, and even preferences [51]. For example, consider a right-handed user who wants to access a menu item on the left side of the screen. For larger screens, this menu item is more difficult for the user to touch with their right hand than for a left-handed user.

UI adaptivity is a promising approach towards improving touch interactions, because it allows systems to dynamically personalize the UI and tailor the UI to the users' needs. But for UI adaptivity to be practically useful for real-world apps, it must support two goals. First, the technique should be generally useful across a broad range of existing mobile applications. Second, the technique should apply adaptations in a way that respects the design intentions of the original applications. In other words, we expect that drastic UI adaptations are likely to make the user interface less familiar to the user and disruptive to the overall user experience [7].

To operationalize these goals, we built a system, *Reflow*, which automatically applies small UI adaptations—called *refinements*—to mobile existing app screens to improve touch efficiency. Towards the first goal of supporting broad applicability, Reflow is entirely pixel-based: the system does not need knowledge of the applications' dependencies or view hierarchy to make its UI adaptations. Towards the second goal of respecting design intent and minimally disrupting the user experience, Reflow incorporates the theory of microstrategies in its model [12, 13, 17]. Microstrategies suggest that even small, principled adaptations to the user interface can significantly improve task efficiency—particularly over cumulative usage—and we postulate that the same principle applies when personalizing touch-based mobile applications.

Reflow supports personalized optimization by constructing a spatial map from usage data, which identifies difficult-to-access areas of the screen (*e.g.*, elements on edges of the screen requires users to reach and reposition their hand to select). Reflow then *(i)* automatically detects the UI elements contained on the screen, *(ii)* uses a machine learning model to optimize the UI layout to better support the difficulty map, and then *(iii)* re-renders the existing UI pixels to match the new layout (Figure 1). Reflow improves on existing approaches because it works with a range of existing mobile applications and enables an end-to-end pipeline from layout optimization to re-rendering the application screens.

To evaluate Reflow, we first conducted a study with 10 participants, where we found it improved interaction speed by 9% on average, and improved interaction speeds by up to 17% for some UIs. From lessons learned, we made further improvements to this model by detecting and applying an additional set of UI constraints (relative positioning, alignment). We then conducted a heuristic evaluation based heuristics for evaluating UI layouts [43, 49] with 3 accessibility and 3 design experts to validate if these improvements make acceptable trade-offs between touch efficiency and layout preservation. Feedback from our expert evaluators indicated that refinements were likely to improve selection time while avoiding significant disruption to the UI. The results of this work demonstrate that refinements are a useful UI adaptation technique to improve task efficiency for touch interactions.

The contributions of our paper are as follows:

Fig. 2. Reflow optimizes existing third party apps using only UI information from pixels and a spatial difficulty map corresponding to a user's abilities. An app screen with short menu items near the top of the screen is difficult to use (Left). Reflow automatically optimizes the layout of on-screen elements, making each menu item taller and shifting items down (Right).

- We propose an approach based on the theory of microstrategies [17] for improving touch interactions through *refinements*, which are small modifications to the original UI. Based on this approach, we present Reflow, an end-to-end system for personalizing any existing mobile app using only its pixels. We describe our implementation in several modular steps: *(i)* element detection, *(ii)* layout refinement, *(iii)* UI re-rendering, which may be applied to other UI adaptation systems.
- We conduct two evaluations that provide evidence for the effectiveness of Reflow's automatic UI refinements. From our user study ($n = 10$), we find that the refinements automatically applied by Reflow result in more efficient touch interaction (average speedup of 9%, up to 32%). Furthermore, we conduct a heuristic evaluation with 3 accessibility and 3 design experts, validating that the changes induced by Reflow are likely to improve selection time while being minimally disruptive. Qualitative feedback from our expert evaluators provide additional rationale for the acceptability of these trade-offs.

2   EXAMPLE USAGE SCENARIO FOR REFLOW

To motivate how Reflow can be used to improve touch interactions on mobile apps, we provide an expected usage scenario where Reflow would be enabled globally through the mobile operating system.

Scenario: Alice has recently purchased a new smartphone. Alice's new smartphone is able to install and run all of her favorite apps, but the device is slightly larger than her old smartphone. Because of this, she can no longer comfortably reach UI controls on the far edges of the screen while using the device with one hand. While she is still able to use her favorite apps, Alice finds her routine interactions with the apps inconvenient.

The Reachability accessibility feature built into iOS[1] addresses the challenge of touching items in the upper half of the screen by enabling users to swipe down on the lower portion of the screen to move the upper half of the screen

---

[1]https://support.apple.com/guide/iphone/touch-iph77bcdd132/14.0/ios/14.0#iph145eba8e9

down to the lower half. While this is useful to many people, it requires an additional touch interaction (swipe down) which slows Alice down and feels like more support than she requires.

Setup: Alice opens the settings application on her smartphone and turns on the "Reflow" toggle button. If this is the first time the setting has been enabled, Alice is brought to a setup screen that that initializes the Reflow system. This screen is similar to the first-time setup process of some biometric authentication features (*e.g.,* Face ID). The setup screen asks Alice to perform a calibration task, and measures how quickly she can select targets located at different parts of the screen. With this usage information, Reflow creates a personalized profile for her that identifies difficult-to-reach areas on the screen.

Usage: Once enabled, Reflow automatically intercepts and applies refinements to app screens before they are displayed to the user, a form of *manifest interface* [10]. The UI produced by Reflow is interactive, as it automatically redirects input events to the original app screen, using techniques similar to previous work [45, 57]. Alice notices that the location and size of UI elements on apps have only slightly changed, and that the overall appearance and structure of app UIs are still very similar (Figure 2). Because of this, she is once again able to comfortably use her favorite apps, as she was previously able to do with her old phone, but she does not need to re-familiarize herself with the refined UIs.

Customization: Alice may decide to customize the behavior of Reflow by opening its settings panel. In this panel, she may define lists of apps that she wishes the refinement feature to include or exclude. This could be useful for disabling the feature on apps that Alice already finds easy-to-use on her new smartphone. Similarly, Alice can define a shortcut (*e.g.,* gesture or detected activity) to quickly toggle Reflow's functionality based on when it is useful. Finally, the settings panel allows her to reset or re-calibrate the feature to reflect updated preferences or physical affordances (for example, a smartphone case that makes it easier to grip the device).

3 RELATED WORK

Several areas of related work have informed the design of Reflow: *(i)* difficulties with touch interaction, *(ii)* adaptive user interfaces, and *(iii)* improving existing applications.

3.1 Difficulties with Touch Interaction

We first review literature related to the characterization of touch input on smartphones to identify why users like Alice experience interaction difficulty. Touch interaction can be analyzed using existing models of cursor-based selection [30], which has been extended to the touch screen [5]. This type of analysis gives recommendations about the relative size and spacing of UI elements on touch-based apps. Other work has focused on properties specific to touch interaction, such as the effect of the finger choice on selection accuracy. As examples, touch input is imprecise due to factors such as finger deformations (i.e., the "fat finger" problem [50]), occlusions when the hand covers up UI elements [18, 19], and varying finger-to-screen ratios [3].

Many of these factors are spatially dependent (*e.g.,* the finger occludes more of the screen when it makes contact with the touchscreen at an extreme angle). Based on common hand postures used to grip a smartphone with one hand, Le et al. [27] characterized the region of the phone that could be comfortably reached. Mayer et al. [33] give further evidence through their analysis of 45 million touch events collected during touch-based gameplay, where they identified a region of the screen most likely to be comfortable to tap, known as the "sweet spot." Based on this evidence, Reflow takes a spatially-dependent approach to personalizing touch interactions by modifying the position and size of UI elements.

3.2 Adaptive User Interfaces

Some UIs are constructed so that they can automatically reconfigure themselves dynamically depending on usage context. For example, adaptive UIs have been constructed to create alternate layouts for additional contexts [55] and provide accessibility benefits to older adults with motor and cognitive impairments [41]. This functionality can be manually programmed by app developers who expect usage difficulty to occur in certain contexts [22, 32].

Another approach is to define an objective function that measures how well a layout is perceived based on desired qualities [38, 39, 47]. SUPPLE is one example which uses this approach to automatically personalize an interface to facilitate faster access and lower error [15]. This approach has also shown encouraging results for optimizing application mockups for low-vision and motor impaired users [16]. In contrast to this approaches, Reflow uses an objective function that is learned from personalized usage data. Specifically, we construct a neural network that predicts the time required to select on-screen elements based on performance on a calibration task, and our system aims to minimize this predicted value.

The approach by Duan et al. [11] employs gradient descent to optimize UIs with respect to estimated task completion time is most similar to our adaptation technique. In our work, we improve on this approach by: *(i)* incorporating personalization through user-specific calibration data, *(ii)* applying constrained optimization to minimize disruptions to the original UI, and *(iii)* supporting end-to-end optimization of existing mobile app screens. A limitation of many approaches is that they cannot be applied to real-world apps, as they often require the UI to be defined in a certain way (*e.g.,* a layout definition) or implemented using specialized toolkits [15, 16, 20, 54]. Reflow uses a pixel-based approach without having to rely on the applications' underlying implementation.

3.3 Improving Existing Applications

Adaptive UIs that need to be built from scratch, using specific UI toolkits, and with specific requirements on developers, is severely limiting in practice. It is difficult to get developers to adopt any new UI framework, and this approach does not address the substantial legacy of existing applications not created to support this use. An alternative is to repurpose and augment existing applications. For example, existing UIs can be retargeted to support new modalities [46] or support responsive resizing [21, 52]. *Interaction proxies* apply runtime modification to existing mobile apps by "repairing" inaccessible or difficult-to-use UI elements [57]. Interaction proxies and other approaches that apply input/output redirection [45] use specialized UI elements to re-render parts of the original UI while maintaining interactivity.

To support a broad range of existing applications by not having to rely on application dependencies or other application metadata, some approaches improve applications only from their pixels (visual appearance). Template-based pixel matching has been used to locate icons or UI elements of interest on a screen to support early screen readers [42] and end-user scripting [53]. Similarly, Prefab enables custom interaction techniques (*e.g.,* target-aware pointing) for desktop applications through pixel-based identification of user interface elements [10]. An important contribution of the Reflow system is to demonstrate how UI refinements supporting better touch interaction can be made to existing mobile UIs directly from their pixels.

4 REFLOW

Reflow is an end-to-end system that produces a refined UI from an original app's screenshot. Reflow's operations occur in several stages (Figure 3). First, the user is asked to perform a calibration task, which is used by Reflow to construct a difficulty map characterizing areas of the screen which may be hard-to-reach. During runtime, an element detector
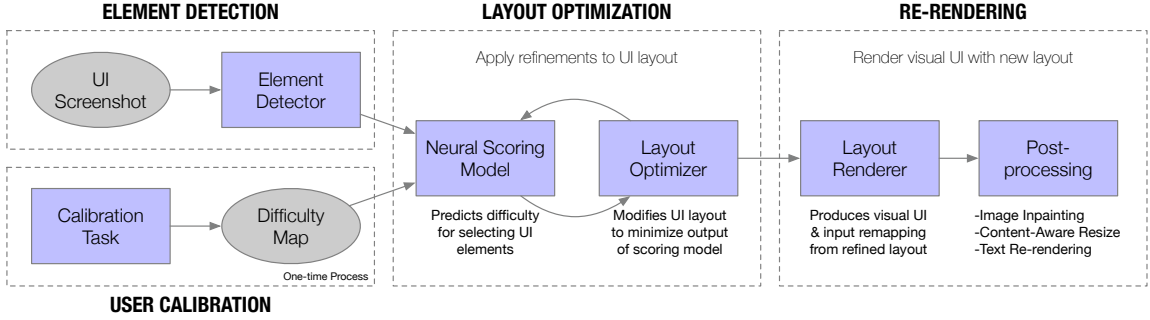
ELEMENT DETECTION   **LAYOUT OPTIMIZATION**   **RE-RENDERING**

Apply refinements to UI layout   Render visual UI with new layout

UI Screenshot → Element Detector

Calibration Task → Difficulty Map

One-time Process

Neural Scoring Model ⇄ Layout Optimizer

Predicts difficulty for selecting UI elements

Modifies UI layout to minimize output of scoring model

Layout Renderer → Post-processing

Produces visual UI & input remapping from refined layout

-Image Inpainting
-Content-Aware Resize
-Text Re-rendering

USER CALIBRATION

Fig. 3. A block diagram describing the architecture of Reflow. Reflow consists of 4 main stages: *(i)* user calibration, *(ii)* element detection, *(iii)* layout optimization, and *(iv)* re-rendering. Together, these stages allow an existing UI screenshot to be adapted to a personalized difficulty model.

is used to extract the layout of an existing app screen. This layout is then optimized by repeatedly *(i)* predicting the selection time for UI elements given the UI layout and user-specific difficulty map and *(ii)* modifying the layout to minimize this predicted value while respecting stopping conditions designed to detect large disruptions (*e.g.,* overlapping, other constraints). Finally, the refined layout is re-rendered into a refined graphical user interface, where it can be presented to the user and made interactive.

4.1 User Calibration

4.1.1 Calibration Task. To characterize the areas of the screen that are hard-to-reach for a user, we designed a one-handed calibration task that a user performs when first enabling the feature. We focused on one-handed touch interactions because these motor activities are more likely to cause selection challenges in user interfaces, particularly given the limited range and flexibility of the thumb [3, 9, 35]. Users were asked to select on-screen targets, and the system recorded the *(i)* selection time and *(ii)* selection error (*i.e.,* difference between the target position and recorded tap position). These targets were uniformly spaced in a 4x8 grid, with a single, randomly selected target highlighted at a time. After a target was selected, the next one was not immediately displayed; instead, it was delayed for the remainder of a timeout value. This was done to allow the finger to return to a "rest position" after selection and reduce the influence of the previous target's location on the finger's starting position. We set the timeout value to 3 seconds, based on our early observations and estimation of how long this process would take.

We conducted a data collection study with 10 participants (7M/2F/1 prefer not to disclose, ages 22-40, recruited within our organization) to characterize the input error for users and to initialize our system. We performed data collection remotely using video conference software. Participants were asked to install an app on an iPhone, which was needed to run our software. 4/10 of our participants used an iPhone Xs, 2/10 used an iPhone 11, 3/10 used an iPhone Xs Max, and 1/10 used an iPhone 11 Pro. Users were asked to hold their device with one hand and tap on targets using the same hand that they were holding the device with.

During the study, participants were asked to select targets placed at different locations on the screen. In total, the study required less than 30 minutes and included both a practice and evaluation session. 9/10 participants held the phone in their right hand.
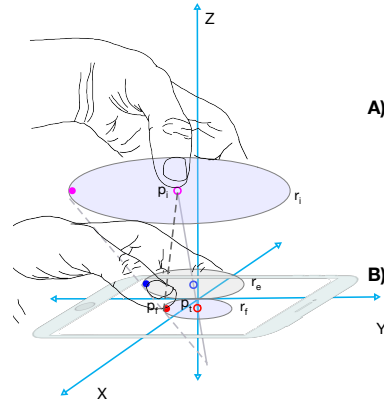
Fig. 4. A user taps a location on the screen by lowering a finger from a starting position p_i (A) to a target location on the screen p_t (B), whose path is shown as a dark gray dotted line. The circles represent region on the screen where the finger is expected to land given the current position. As the finger descends towards the screen, the circle shrinks, as the user "hones in" on the target location. The user's finger makes contact with the screen at p_f, which is a distance r_f away from p_t (B).
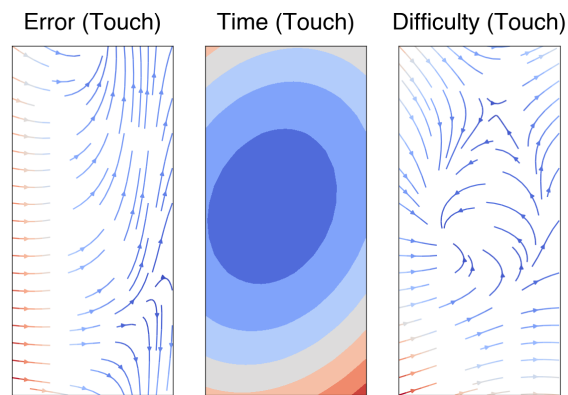

Fig. 5. An example of a difficulty map generated by our user calibration process. Data from the calibration task is first used to compute input error (Left) and selection time (Center) independently. We combine these two measurements to produce the "adjusted error", or difficulty map (Right).

4.1.2 Difficulty Map. Using calibration data from each user, we generated a personalized difficulty map [34] that estimates the relative difficulty of accessing any on-screen location (including locations between two of the original calibration points).

We calculated difficulty by combining two measurements from the calibration data: *(i)* input error and *(ii)* selection time. We introduce a procedure to "normalize" the raw error measurements by selection time. Our procedure is based on the intuition that if a user selects two targets with equal accuracy (*i.e.,* equal input error), the one that took longer to select (*i.e.,* higher selection time) is more difficult. We refer to this computed value as the "adjusted error", and it represents an estimation of input error given a constant selection time. In other words, if a user was only given one second to select anywhere on the screen, what is the error we would expect for a given location? Our approach is

informed by prior work which suggests that selection-based interactions experience a tradeoff between speed and accuracy [30].

Our computation of adjusted error is based on the standard Fitts's law equation. Figure 4 shows a user's finger, initially located at p_i (*i.e.,* resting position), selecting a target located at p_t, which is located a distance A away.

$$t = a + b \cdot \log_2(A)$$
$$t = a + b \cdot \log_2(||p_f - p_i||) \tag{1}$$

If the resting position p_i is assumed to be constant for a single user (which previous research has shown is plausible [26, 33]), we can estimate its location as a model parameter. To do this—in addition to the standard Fitts' model parameters a and b—we also learned p_i by fitting Equation 1 to a dataset of pairs $\langle t_i, p_f \rangle$ using non-linear ordinary least squares. The initial value of p_i was set based on the resting location area identified in previous work [26], and we constrained the position of p_i so that it lies at most 5 inches above the device screen (a reasonable upper bound for the thumb's distance to the screen when holding a smartphone) and within the screen's x-y bounds.

To estimate the error magnitude ϵ at a constant time t_n, we solve the following system of equations, where a and b are from the Fitts's model, and $\mathbf{d}$ is the distance the finger has traveled at t_n.

$$((A, r_f, 1) \times (0, r_i, 1))(\mathbf{d}, \epsilon, 1)^{\mathbf{T}} = 0 \tag{2}$$

$$\mathbf{d} = 2^{\frac{t_n - a}{b}}$$
$$\epsilon = \frac{(r_f - r_i) \cdot \mathbf{d}}{A} + r_i \tag{3}$$

Using this procedure, we compute the adjusted error at each of the original calibration points (4x8 grid), then fit a bivariate polynomial (similar to a 2-D spline) to interpolate values in between.

To summarize, the output of the user calibration step is a function that returns the adjusted error for any location on the screen. The adjusted error at location (x, y) is an estimate of the offset of where a user's touch would land if given t_n seconds to tap a target located at (x, y).

4.2 Element Detection

In this stage, Reflow extracts the locations of on-screen UI elements using a CNN-based object detector and grouping heuristics [56]. We performed additional post-processing on the output of the object detector to further improve performance. First, we removed any detection that overlaps (by keeping the one with the higher confidence) or contains another (if a detection contains a text element, then the container is removed; otherwise, the children are removed). The resulting layout contains no overlapping elements. Because the downstream re-rendering stage (Section 4.4) involves cropping and moving image patches, the quality of element bounding boxes identified in element detection have affect the quality of the final output. We employed a heuristic that refines the positions of bounding box edges by repeatedly *(i)* computing the mean color of the pixels of each edge, *(ii)* selecting the edge whose mean color is most dissimilar from the others, *(iii)* adjusting its value by a small increment in the direction of improvement.

The output of the element detector is the UI layout of the original screenshot: $\Theta = \{\theta_1, \theta_2, \cdots, \theta_n\}$ where $\theta_i = [x, y, w, h]$ describes the location and size of the i-th bounding box.

4.3 Layout Optimization

4.3.1 Neural Scoring Model. Using a difficulty map, we estimated the error a user would experience when selecting individual elements of the UI. For a given app layout we define a scoring function $S(\Theta)$ which quantifies how likely the
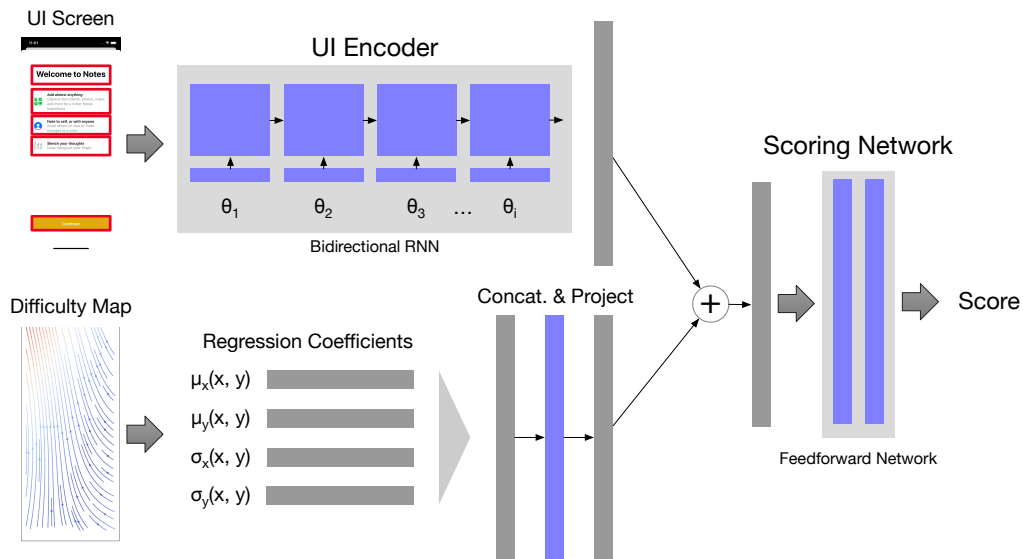
Fig. 6. The architecture for the neural scoring model used for scoring a UI screen given a spatial map of error. The network encodes the layout of UI elements using a bidirectional RNN and encodes the spatial difficulty map using the coefficients of a 2-D polynomial function fitted to the calibration points. These encoded representations are combined and fed into a feedforward network.

user is able to successfully select each element in the screen. We define the scoring function as

$$S(\Theta) = \sum_{\theta \in \Theta} \iint_{R_\theta} P_\theta \, dA \qquad (4)$$

where P_θ is the predicted distribution of where the user's finger will actually land when attempting to click the middle of an element θ. To estimate P_θ, we centered 2-D Gaussian on the center of UI element θ and set its covariance based on the adjusted error at that location. We compute $S(\Theta)$ using Monte Carlo integration. For each UI element, n samples are drawn from the 2-D Gaussian parameterized by the adjusted error at its center. We set $n = 30$ based on empirical observations of how many samples were needed for consistent results. We scored a screen by counting the number of true positives over total number of points.

To further improve upon the speed and efficiency of our UI scoring function, we used a neural network to learn the result Monte Carlo scoring function. Our model architecture (Figure 6) consists of an LSTM-based UI layout encoder [11, 28] and a difficulty map encoder. Because we parameterize the difficulty map as a model input, it allows for personalization at runtime—allowing re-calibration and removing the need to retrain the network for every new user.

We trained our neural scoring network on two datasets. The first was a large annotated dataset of 77,000 iOS app screens introduced in prior work [56]. The second dataset consisted of spatial difficulty maps collected from our user calibration dataset. Because the second dataset is relatively small, we augmented it by introducing perturbations and randomly generated difficulty maps. The network was trained by randomly selecting an application screen and a spatial difficulty map and computing the score using our Monte Carlo algorithm, which was used as the ground truth. We trained our network using stochastic gradient descent until the validation loss stopped improving.
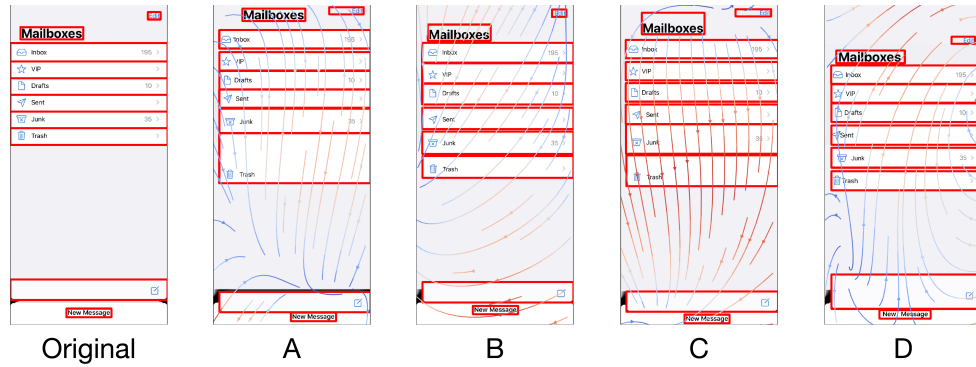
Fig. 7. An email app ("Original") is optimized using four different spatial difficulty maps ("A", "B", "C", "D"). The result of the optimization process moves UI elements away from regions of higher relative difficulty (red) to lower relative difficulty (blue). Depending on the direction of error where an element is located, it may be stretched horizontally or vertically.

4.3.2 Layout Optimization. To optimize a UI layout, we use predictions from our neural scoring model to make modifications that reduce expected interaction difficulty. One benefit of our model is that it is differentiable (since it is a neural network) and thus, given an initial screen layout and its corresponding score, we can use gradient-based optimizer [23] to find an improved layout with a better score. In service of our refinement approach to UI adaptation, we guide optimization and stop the process if disruptive changes are detected.

First, we add a regularization term that penalizes layouts that are proportionally dissimilar to the original. This regularization term is defined as the cosine distance (D_C) between the pairwise L_1 displacements of each UI element ($\phi(\Theta)$). This does not penalize the layout for increasing in size but attempts to maintains relationships between neighboring UI elements.

Next, we add "corrective procedures" after each optimization step to guide the optimization of certain properties (*e.g.,* size and element). We clamped element parameters to ensure they stay within a certain range and to prevent elements from becoming too small or large. We also use a routine that detects overlapping regions and resolves them by moving overlapped elements further apart. The overlap removal algorithm repeatedly tries to find intersecting region between pairs of UI elements, and if one exists, shifts them apart in the axis of least overlap. This process is repeated until no more overlaps are detected or a max number of iterations is reached. Despite our precaution, the final UI may still appear to contain overlaps due to inaccurate element detection (*e.g.,* detection bounding box does not fully enclose element or includes multiple elements) and artifacts introduced by our re-rendering process.

Finally, we implemented an early stopping condition that is triggered when the overlap removal algorithm detects overlaps that are unresolvable, that is, overlaps that cannot be resolved by moving elements further apart.

Figure 7 shows examples of our algorithm's output for an email app optimized using different difficulty maps.

4.4 Re-rendering

Using the refined UI layout produced by layout optimization, we produce artifacts that are needed for end-user interaction: *(i)* visual representation of the UI (UI screenshot) and *(ii)* mapping between the original and refined UIs, needed for input redirection [45, 57].
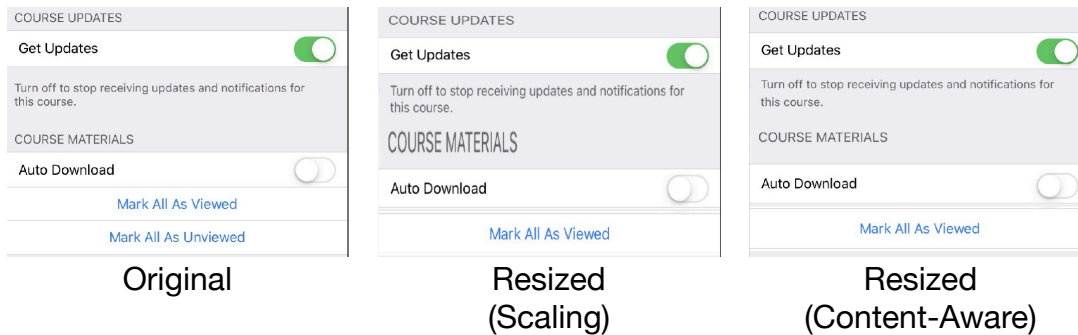
Fig. 8. An example of how content-aware resizing (Right) improves upon standard scaling (Center). Compared to the original screenshot (Left), scaling (Center) introduces distortions that make text less legible. Our approach (Right) preserves textual content.

4.4.1 Layout Renderer. We use the refined UI layout and re-render the UI back into a visual representation. First, Reflow outputs a mapping between regions of the original UI and the refined UI. This mapping can be used to update the interactive regions of the UI using input/output redirection methods [45, 57] (*e.g.,* clickable bounding box a button is updated to reflect its new optimized position). To align the screen's visual appearance with the updated regions, image patches from the original screen are translated and resized to their new locations.

4.4.2 Post-processing. We use several post-processing strategies to preserve relevant aesthetic qualities (*i.e.,* legibility of text). We briefly describe three such techniques we used, which improve the legibility of: *(i)* background areas, *(ii)* text, and *(iii)* image content.

Image Inpainting. Moving and resizing elements can result in "holes" so we employ an inpainting technique to generate visually plausible replacements. We experimented with many different inpainting algorithms [2, 4, 48], and found that most methods produced results of similar visual quality, since inpainted regions are unlikely to contain complex textures or structural features (most visual content is contained inside of the UI elements themselves). Our implementation uses flow-based inpainting algorithm [4] included in the OpenCV library [8].

Content-Aware Resizing. Because standard rescaling methods can distort elements (Figure 8), our system uses an optimized method for resizing these regions. We first create an image canvas with the target region's size. The source image patch resized (retaining the original aspect ratio) to maximum size fits within the target dimensions. Leftover space is filled using image inpainting.

Text Re-rendering. We explored a method specifically for resizing and replacing text (*e.g.,* translation example application). Our approach to resizing text involves detecting and extracting text using optical character recognition (OCR), re-rendering the text with the correct background and foreground colors, then inserting the result at the target location. We used an off-the-shelf OCR system that recognizes text from images [44]. We estimated the original text's font size by rendering it using a known font and comparing its dimensions to the size of the original image patch. Background and foreground color are extracted from the original image patch by performing k-means clustering on the pixels and extracting *(i)* the largest cluster (background color), and *(ii)* the cluster which is furthest away from the background color in pixel space (foreground color). We created an image

patch corresponding to the target dimensions and rendered the text with the correct foreground and background colors.

4.5 Prototype Implementation

For the purposes of our prototype, the Reflow system was implemented on a remote server. Screenshot images are sent from the iOS device to the remote server, which returns a re-rendered screenshot image that contained the UI refinements. For purposes of the user study (described next), the re-rendered screenshot is displayed to the user and touch events are recorded, which allowed us to conduct our user study. Prior work has demonstrated how such a re-rendered graphical UI could be used in a more advanced proxy setup to control the original underlying mobile app, which is how we imagine it would be used with existing applications in practice. We leave that implementation, and some difficult details (*e.g.*, handling scrolling GUIs) to future work.

5 USER STUDY

5.1 Procedure

We empirically evaluated the performance of our system through a 45-minute user study. We recruited 10 participants (5M/4F/1 Prefer not disclose, ages 24-36) within our organization. Similar to how we collected data in the calibration task (Section 4.1), we conducted our usability study remotely using a specialized data collection app and video conferencing software. The breakdown of devices used by our participants were: 1/10 iPhone X, 5/10 iPhone Xs, 1/10 iPhone XR, 1/10 iPhone Xs Max, 1/10 iPhone 11, 1/10 iPhone 11 Pro. 8/10 participants held the phone using their right hand.

The usability study consisted of two phases, a calibration phase and a navigation phase, which included a practice session. During calibration participants performed the user calibration task, which was used to generate personalized difficulty maps. Participants were then given a short 5 minute break. During the navigation phase, participants were presented with a series of app screens with a target UI element that was highlighted. Participants were instructed to select the target element, which brought them to the next screen.

We chose 3 apps where users had to navigate through a total of 5 screens each (Figure 9). 3 apps were chosen to constrain the study design to fit under our time constraints, and we selected default applications on iOS that users are likely to be familiar with, since they come pre-installed. All of these apps and screens were outside of the training set used to train our scoring model. For each of the screens, we used our system (running on a remote server) to optimize its layout using the user's difficulty map generated during the calibration phase.

5.2 Results

Table 1 shows the results of our user study. We report both screen (time taken to navigate a single screen) and app (time taken to navigate all screens in an app) interaction times. For both, Reflow's refinements resulted in an average of 9% faster navigation. This speed-up was significant for screen-level navigation ($T = 2.76$, $p < 0.01$, Cohen's $d = 0.23$) and approached significance for app-level navigation ($T = 1.99$, $p = 0.06$, Cohen's $d = 0.38$), due to small sample size for apps. The small–medium effect size suggests that refinements lead to small, but significant improvements to interaction speed [25].

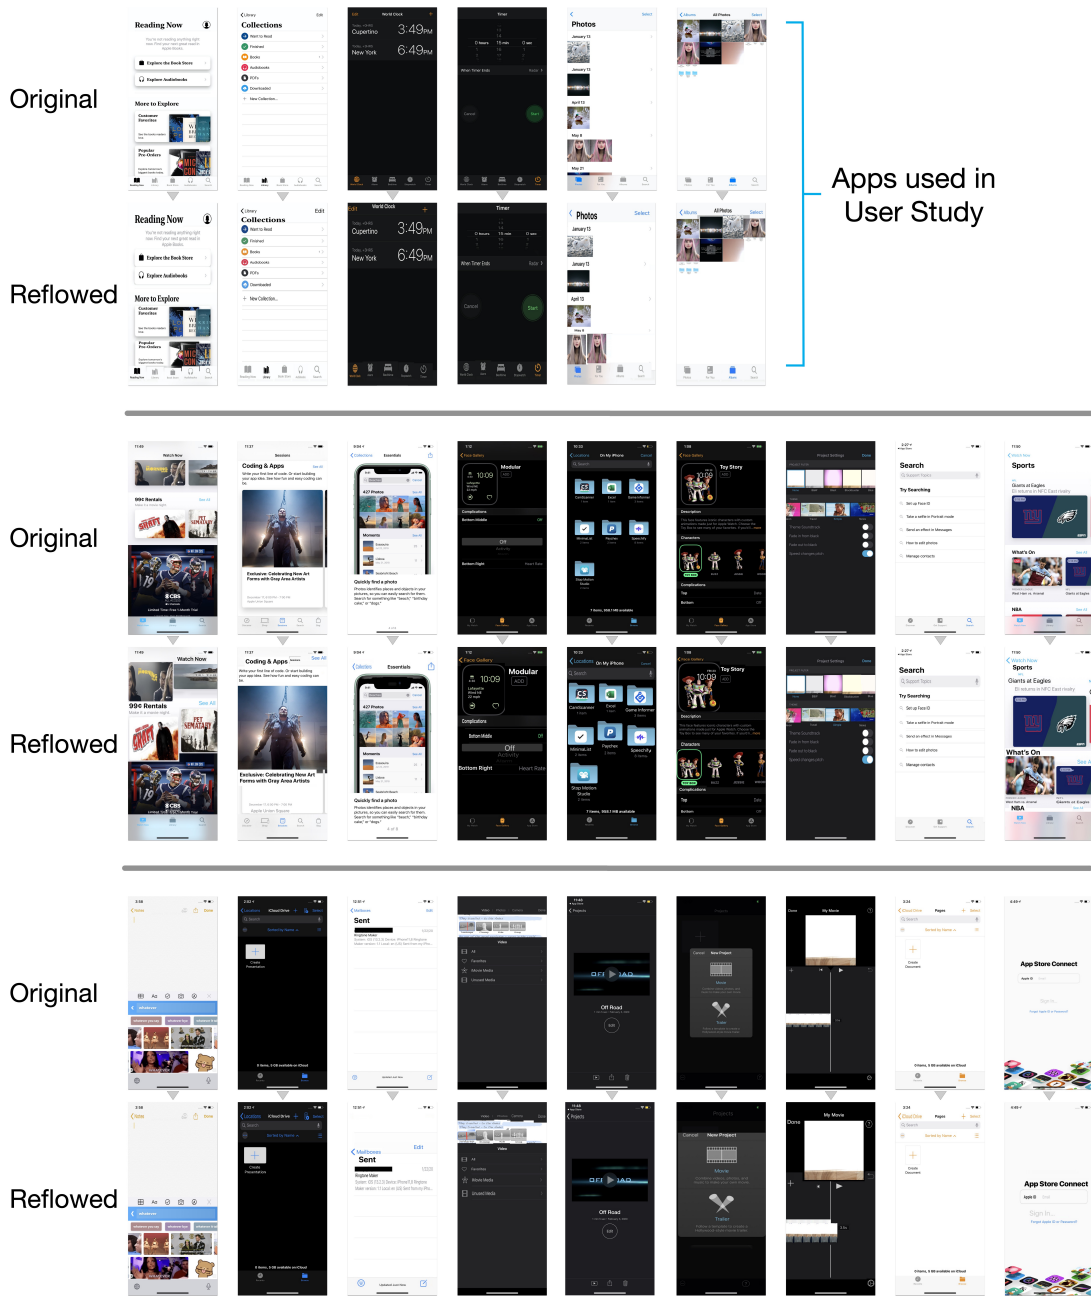Fig. 9. A gallery of app screens optimized by Reflow. For the apps used in the User Study, we only show the first and last screens of each app (each app has 5 screens). As was our intention, most changes by our system were conservative (resulting in little to no changes) and aimed at minimizing any negative effects introduced by drastic changes. Some text (*e.g.,* email address) is redacted in black.

Table 1. Navigation Time Results from our User Study

	Screen		App	
	Original	Reflow	Original	Reflow
Books	1.1 ± 0.4	1.0 ± 0.3	4.2 ± 1.0	3.9 ± 0.9
Clock	1.1 ± 0.6	1.0 ± 0.5	4.4 ± 1.5	3.8 ± 1.3
Photos	1.2 ± 0.5	1.0 ± 0.4	4.7 ± 1.4	4.1 ± 1.1
Overall	1.1 ± 0.5	1.0 ± 0.4	4.4 ± 1.3	4.0 ± 1.1

Navigation times from our user study. Values shown are the times ($M \pm SD$) needed to navigate a single screen. We report results for both per-screen and per-app navigation. Screen navigation refers to the time taken to advance one screen, while app navigation refers to the time taken to complete all screens for the app. For both measurements, Reflow's refinements resulted in 9% faster navigation, on average.

5.3 Post-study Improvements

The results from our user study demonstrate the feasibility of our scoring model to optimize the layout of existing UIs. Based on impromptu feedback from three participants during the study and our observations on how users interacted with Reflow, we made some improvements to further minimize the amount of disruption to the original UI. Specifically, participants commented that the adaptation process affected the ordering of elements (*e.g.,* an element was moved past an adjacent neighbor) and caused controls in UI structures navigation bars became mis-aligned. We implemented additional heuristics to extract these constraints from the UI and used them further guide the optimization process.

Constraint Extraction: Our approach maintains a list of constraints represented as equations and tests them against pairs of elements in the layout to determine if they hold. We check for two types of constraints: *(i)* relative positioning *(ii)* alignment.

The enforcement of relative positioning ensures, for example, that if an element A was to the left of another element B, it remains to its left. Relative positioning is represented using a multi constraint described by the following equation.

$$x_i + w_i \leq x_j \tag{5}$$

To reduce the number of relative constraints generated by our approach (making optimization more efficient), we remove redundant relationships through transitive reduction. For example, if A is left of B and B is left of C, the relative positioning of A and C is implied and can be omitted as an explicit constraint.

Alignment constraints require that some part of two elements are arranged on the same line. We consider 3 types of alignment between elements: *(i)* beginning alignment, *(ii)* center alignment, and *(iii)* ending alignment. These three types are computed for both the x and y dimension, resulting in a total of six possible element alignments. Note that most interface builders may also allow mixing of these types *e.g.,* the left edge of element A is aligned with the center of element B. For simplicity, we omit these.

Because our UI element detector introduces a small amount of error in the bounding boxes, exact computation of alignment (*e.g.,* computing $x_i = x_j$) would be inaccurate, leading to many undetected constraints. Thus, we allow alignments to be met inexactly using slack variables. We consider three types of alignment: beginning alignment, center

alignment, and end alignment.

$$
\begin{array}{ll}
x_i + \epsilon_i = x_j + \epsilon_j & \epsilon_i \geq 0, \epsilon_j \geq 0 \\
x_i + w_i + \epsilon_i = x_j + w_j + \epsilon_j & \epsilon_i \geq 0, \epsilon_j \geq 0 \\
x_i + \frac{w_i}{2} + \epsilon_i = x_j + \frac{w_j}{2} + \epsilon_j & \epsilon_i \geq 0, \epsilon_j \geq 0
\end{array} \tag{6}
$$

Constrained Optimization: We used the non-convex optimization technique introduced by Platt and Barr [40] to augment our original layout optimizer. Specifically, we added the heuristically extracted constraints as "secondary functions" of the main objective function (*i.e.,* minimizing expected difficulty) that the optimizer aims to minimize. Note that the while this technique allows the consideration of constraints, it does not guarantee that they will be met in the final solution.

6   HEURISTIC EVALUATION

6.1   Procedure

To determine the acceptability of our improved system and to obtain qualitative feedback, we conducted a heuristic evaluation as described in Nielsen [37]. Heuristic evaluation has been previously been used as a design-centered analytic evaluation for user interfaces [1, 31]. In a heuristic evaluation, expert evaluators examine the interface or aspects of the interaction to identify usability problems. Nielsen [37] recommends the use of around three to five evaluators, so we contacted 10 people in our organization, in anticipation that some would be unavailable. 6 experts (3M, 3F) from different backgrounds (3 designers, 3 accessibility experts) agreed to participate and provided diverse feedback. Recruitment was done via convenience sampling—we reached out to potential evaluators using our organization's messaging software. All of our evaluators had multiple years of experience in their field and five had doctoral degrees in their respective areas of specialization.

The heuristic evaluation was conducted online, and evaluators were sent a link to the evaluation materials. The link first gave a brief description of Reflow and the evaluation task. Evaluators were instructed to watch a video clip that showed four usage scenarios (*e.g.,* while walking or holding a shopping bag with one hand) where a user interacted with app UIs. Original and adapted versions of the UI were overlayed on the video and allowed evaluators to assess different qualities of the UI in context. Evaluators could pause and replay the video as many times as they needed during the evaluation to more closely inspect aspects of the UI and usage scenario.

We provided a questionnaire with a set of UI heuristics for them to evaluate, informed by prior work on layout usability [43, 49]. We removed heuristics (*e.g.,* color harmony) that were not applicable to Reflow and settled on the following: alignment, selection time, visual clutter, saliency, and element grouping. For each heuristic, we provided a brief description and asked evaluators to rate adherence using a 5-point Likert scale for comparison ("Significantly Worse (1)" to "Significantly Better (5)"). We also asked evaluators to provide rationale for their ratings and encouraged elaboration on the positive and negative aspects on the layout changes introduced by Reflow.

6.2   Results

Table 2 shows each evaluator's 5-point Likert-type item responses. Overall, the results of the heuristic evaluation confirm that Reflow improves touch efficiency while minimally disrupting the user experience. We summarize the feedback from our evaluators along each of the heuristics.

Alignment: *Heuristic.* Alignment refers to the internal alignment of elements with each other. *Evaluation.* Most evaluators agreed that Reflow's refinements slightly impacted alignment. While our improved layout optimizer does

Table 2. Relative Likert Scores for Heuristic Evaluation

Heuristic	AX			DS		
	E1	E2	E3	E4	E5	E6
ALIGNMENT	−1		−1	−1	+1	−1
SELECTION TIME	+1	+1	+1		+1	+1
CLUTTER		−1	−2	−1	+2	−1
SALIENCY	−1	+1		−1	+1	−1
GROUPING	−1		−1		+1	

Expert evaluations from heuristic evaluation. Ratings are normalized to show deviation from the neutral option ("About the same (3)"). Positive scores indicate better change and negative scores indicate worse change. AX denotes accessibility expert while DS denotes design expert.

detect and account for UI constraints such as alignment, it does not guarantee that these constraints are met. We found that some evaluators' intepretations of "alignment" included aspects such as spacing (not explicitly handled by our system) as well as edge alignment (handled by our system) ($E3_{AX}$). Nevertheless, $E2_{AX}$ observed that some layouts *"probably could change more without hurting understanding,"* indicating that Reflow's refinements kept semantic relations are kept intact.

Selection Time: *Heuristic.* Selection time refers to the time needed to select the target element. Since the evaluators did not directly operate each UI, we asked evaluators to estimate the selection time from the relative positioning of elements (*i.e.,*), as in previous work [43, 49]. *Evaluation.* Feedback from our evaluators mostly indicated that our adaptations allow for more rapid selection times during use, but evaluators expressed some uncertainty as they were estimating purely from visual appearance ($E1_{AX}$, $E3_{AX}$). We did indeed empirically measure selection times from our user study (Table 1), and our expert's estimates are consistent with our findings. $E5_{DS}$ pointed out that our system allowed faster interaction *"without having to use reachability"* and would *"certainly save time,"* since performing the Reachability gesture (Section 2) would require more time than using Reflow.

Taken together with other feedback (comment by $E2_{AX}$ on alignment), this suggests that supporting manual adjustment of optimization levels may help Reflow better serve users with differing goals—either prioritizing selection time or preserving aspects of the original UI (*e.g.,* alignment).

Visual Clutter: *Heuristic.* Visual clutter refers to how confusing a display is. The more cluttered a display, the more difficult it is for an element to catch a user's attention *Evaluation.* $E2_{AX}$ pointed out that movement of semantically-important elements contributed to confusion. There was some level of disagreement between evaluators ($E3_{AX}$, $E5_{DS}$), which was in part influenced by different interpretations of the heuristic. For example, $E3_{AX}$ attributed the higher levels of visual clutter to *"worse alignment, unexpected/unintuitive negative space, and overlapping elements,"* indicating the consideration of multiple factors such as alignment, spacing, and overlap.

Saliency: *Heuristic.* Saliency refers to the degree to which important elements are more likely to catch the attention of the user. *Evaluation.* Evaluators suggested that our system's modifications did not induce significant changes in saliency, which is desirable since a large saliency change would imply a violation of design intent. $E3_{AX}$ indicated that the refinement process only resulted in *"subtle changes"* and did not make a single element any more or less salient

than the others. $E2_{AX}$ noted that size refinements, which allows elements in "difficult areas" to be selected more easily, positively impacted the saliency of small items by enlarging their tap targets; however the same behavior was viewed negatively by $E1_{AX}$: *"the adaptations didn't have much impact on saliency, so I almost selected about the same, but I noticed that the tabs have increased size so they may be over-emphasized."*

Element Grouping: *Heuristic.* Elements that are clustered near each other are perceived as a unified object. *Evaluation.* Element grouping was also minimally affected by Reflow's refinement process. While evaluators generally agreed that element grouping remained "about the same" ($E2_{AX}$, $E4_{DS}$, $E6_{DS}$), some pointed out instances where the adjustment of element spacing led to ambiguity: *"on the 4th screenshot pair, the inbox and arrows have scrunched down too closely to the App Store and To field, leaving not enough separation"* ($E1_{AX}$).

7 DISCUSSION

7.1 Flexible Personalization through Difficulty Maps

An architectural decision choice we made in Reflow was to parameterize the difficultly map as an explicit input to the neural scoring model (Figure 3). Doing so has a number of advantages.

Difficulty maps are the key mechanism through which we enable personalization, because they allow users to recalibrate their touch interactions without having to retrain the neural scoring model for every user. Having to retrain a model for every user is not practical, especially if the retraining must be done on a mobile device.

An second advantage of parameterizing difficulty maps is that Reflow can be easily extended to other types of interactions. While our studies evaluated one-handed touch interactions, the Reflow is not restricted to this. Difficulty maps can also be constructed to support reachability, handedness, motor impairments, or other touch accommodations needed to support users. Difficulty maps can even go beyond touch interactions, for example, when using a pen input or other pointing device. Once a mode of interaction is translated to a difficulty map, the rest of the Reflow architecture can take advantage of this—without requiring any changes the remaining stages in the Reflow pipeline.

Finally, difficulty maps make it possible to offer users a set of pre-defined profiles for common touch interaction scenarios. Pre-defined maps could cover scenarios such as one-handed usage and allow many users to benefit from adaptation without having to perform manual calibration. Should these pre-defined maps not support the user, they can always fallback to performing a quick calibration to construct a personalized difficulty map.

7.2 Design Space between Touch Efficiency and Layout Preservation

Reflow applies *refinements*—small UI adaptations that minimally disrupt the UI layout—as a design choice for performing layout optimization. Our heuristic evaluation validates that this decision choice as an appropriate one. However, this decision choice is only one possible point across the full design space, which includes a spectrum of trade-offs between improving touch efficiency against preserving the existing layout. Users may have different preferences along this design space.

For example, all evaluators in our heuristic evaluation—except E4—indicated that selection time would be improved by Reflow. E4 desired to see more dramatic improvements to touch interactions, even if this would require Reflow to make more substantial disruptions to the layout. In contrast, E3 disliked the clustering that resulted from Reflow's layout optimization. Users such as E3 may find it acceptable to have fewer touch improvements, if this would result in smaller layout disruptions.

One possibility is to give users more control over this design space. Consider a slider control that allows the user to select points in the design space between maximizing touch interaction efficiency, minimizing layout disruption, or some balance in between these two extremes. Users may also want to selectively enable or disable specific types of optimizations, for example, they may prefer not to allow the size of UI elements to change. Users may even have different preferences between applications, with varying expectations about layout disruption for the different applications.

7.3 Opportunities for Reflow

Through our prototype implementation and experiments, we identified opportunities and future work for Reflow, including *(i)* extracting and incorporating screen semantics, *(ii)* improving interaction fidelity, and *(iii)* studying the effect of applying refinements over time.

7.3.1 Extracting Screen Semantics. Several of these opportunities involve addressing existing limitations of Reflow. First, the set of constraints that we automatically extract should be extended to support those available in conventional UI authoring tools, such as vertical and horizontal alignment guides, distributing vertical and horizontal spacing, and resizing elements across axes. Including these constraints and inference techniques [6, 21, 24, 29] in Reflow would further minimize layout disruptions. Currently, Reflow internally represents UIs as a list of bounding boxes, but the limitations of this approach are that it is unable to capture semantic relationships *between* elements during inference. Because of this restriction, Reflow currently treats a list of *n* UI elements as *n* unrelated UI elements, which is potentially a lost opportunity for layout optimization.

7.3.2 Improved Interaction Fidelity. Reflow's pixel-based pipeline consumes an image (*i.e.,* screenshot) of the current UI as input and also generates an image as output. The refined UI is made interactive by making certain parts of the output image respond to touch events, which can then be forwarded to the original app. While this allows Reflow to be applied to any screen, it may lead to poor performance on screens with dynamic properties such as animated content and scrolling because this behavior cannot be adequately captured in a static screenshot.

One way to extend the current approach to handle these cases is to perform this image-to-image process multiple times per second, thus updating the output as frequently as dynamic behavior occurs. Besides the optimization and performance challenges this entails, repeatedly optimizing single video frames as independent inputs may lead to artifacts such as jitter. An alternative approach is to regenerate the interface from extracted semantics and interfaces, which has previously been applied to web applications [36]. Recent work in pixel-based semantic extraction [14, 52] suggests this may also be possible for mobile UIs.

7.3.3 Extended Evaluation. Finally, a longer-term usage study may reveal more detailed effects of adaptive UIs and, more specifically, our refinements approach. Although our studies demonstrate the effectiveness of refinements for improving touch efficiency, it would be important to evaluate Reflow in longitudinal studies. The benefits of Reflow can only be fully realized through cumulative use: the longer users use Reflow, the more opportunities they have to take advantage of the adapted UIs. As Gray and Boehm-Davis [17] observe—"milliseconds matter"—and even seemingly small improvements add up over frequent and repeated touch interactions.

8 CONCLUSION

In this paper, we introduced Reflow, a system that automatically applies small, personalized UI adaptations—called *refinements*—to mobile app screens to improve touch efficiency. Reflow supports real-world UIs without any source code

or metadata dependencies through pixed-based element detection. Reflow optimizes a UI by *(i)* extracting its layout from its screenshot, *(ii)* refining its layout, and *(iii)* re-rendering the UI to reflect these modifications. We conducted a user study with 10 participants and found that UIs optimized by Reflow were on average 9% faster to use. We conducted a heuristic evaluation with 6 experts to elicit qualitative feedback about Reflow and validate that the system's UI refinements improve selection time while minimizing layout disruption. The results of our work demonstrate that refinements applied by Reflow are a useful UI adaptation technique to improve touch interactions.

REFERENCES

[1] Titus Barik, Yoonki Song, Brittany Johnson, and Emerson Murphy-Hill. 2016. From quick fixes to slow fixes: Reimagining static analysis resolutions to enable design space exploration. In *2016 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE, 211–221.

[2] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. 2009. PatchMatch: A randomized correspondence algorithm for structural image editing. In *ACM Transactions on Graphics (ToG)*, Vol. 28. ACM, 24.

[3] Joanna Bergstrom-Lehtovirta and Antti Oulasvirta. 2014. Modeling the functional area of the thumb on mobile touchscreen surfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1991–2000.

[4] Marcelo Bertalmio, Andrea L Bertozzi, and Guillermo Sapiro. 2001. Navier-stokes, fluid dynamics, and image and video inpainting. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, Vol. 1. IEEE, I–I.

[5] Xiaojun Bi, Yang Li, and Shumin Zhai. 2013. FFitts law: modeling finger touch with fitts' law. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1363–1372.

[6] Pavol Bielik, Marc Fischer, and Martin Vechev. 2018. Robust relational layout synthesis from examples for Android. *Proceedings of the ACM on Programming Languages* 2, OOPSLA (2018), 1–29.

[7] Jeffrey P. Bigham. 2014. Making the Web Easier to See with Opportunistic Accessibility Improvement. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (Honolulu, Hawaii, USA) *(UIST '14)*. Association for Computing Machinery, New York, NY, USA, 117–122. https://doi.org/10.1145/2642918.2647357

[8] G. Bradski. 2000. The OpenCV Library. *Dr. Dobb's Journal of Software Tools* (2000).

[9] Youli Chang, Sehi L'Yi, Kyle Koh, and Jinwook Seo. 2015. Understanding users' touch behavior on large mobile touch-screens and assisted targeting by tilting gesture. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 1499–1508.

[10] Morgan Dixon and James Fogarty. 2010. Prefab: Implementing Advanced Behaviors Using Pixel-Based Reverse Engineering of Interface Structure. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) *(CHI '10)*. Association for Computing Machinery, New York, NY, USA, 1525–1534. https://doi.org/10.1145/1753326.1753554

[11] Peitong Duan, Casimir Wierzynski, and Lama Nachman. 2020. Optimizing user interface layouts via gradient descent. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.

[12] Katherine Eng, Richard L Lewis, Irene Tollinger, Alina Chu, Andrew Howes, and Alonso Vera. 2006. Generating automated predictions of behavior strategically adapted to specific performance objectives. In *Proceedings of the sigchi conference on human factors in computing systems*. 621–630.

[13] Sarah P Everett and Michael D Byrne. 2004. Unintended effects: Varying icon spacing changes users' visual search strategy. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 695–702.

[14] Michael Fischer, Giovanni Campagna, Silei Xu, and Monica S Lam. 2018. Brassau: automatic generation of graphical user interfaces for virtual assistants. In *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services*. 1–12.

[15] Krzysztof Z Gajos, Daniel S Weld, and Jacob O Wobbrock. 2010. Automatically generating personalized user interfaces with Supple. *Artificial Intelligence* 174, 12-13 (2010), 910–950.

[16] Krzysztof Z Gajos, Jacob O Wobbrock, and Daniel S Weld. 2008. Improving the performance of motor-impaired users with automatically-generated, ability-based interfaces. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. 1257–1266.

[17] Wayne D Gray and Deborah A Boehm-Davis. 2000. Milliseconds matter: An introduction to microstrategies and to their use in describing and predicting interactive behavior. *Journal of experimental psychology: applied* 6, 4 (2000), 322.

[18] Christian Holz and Patrick Baudisch. 2010. The generalized perceived input point model and how to double touch accuracy by extracting fingerprints. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 581–590.

[19] Christian Holz and Patrick Baudisch. 2011. Understanding touch. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2501–2510.

[20] Yue Jiang, Ruofei Du, Christof Lutteroth, and Wolfgang Stuerzlinger. 2019. ORC layout: Adaptive GUI layout with OR-constraints. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.

[21] Yue Jiang, Wolfgang Stuerzlinger, and Christof Lutteroth. 2021. ReverseORC: Reverse Engineering of Resizable User Interface Layouts with OR-Constraints. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–18.

[22] Shaun K Kane, Jacob O Wobbrock, and Ian E Smith. 2008. Getting off the treadmill: evaluating walking user interfaces for mobile devices in public spaces. In *Proceedings of the 10th international conference on Human computer interaction with mobile devices and services*. 109–118.

[23] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

[24] Rebecca Krosnick, Sang Won Lee, Walter S Laseck, and Steve Onev. 2018. Expresso: Building responsive interfaces with keyframes. In *2018 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, 39–47.

[25] Daniël Lakens. 2013. Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for t-tests and ANOVAs. *Frontiers in psychology* 4 (2013), 863.

[26] Huy Viet Le, Sven Mayer, Patrick Bader, and Niels Henze. 2018. Fingers' Range and Comfortable Area for One-Handed Smartphone Interaction Beyond the Touchscreen. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–12.

[27] Huy Viet Le, Sven Mayer, Benedict Steuerlein, and Niels Henze. 2019. Investigating Unintended Inputs for One-Handed Touch Interaction Beyond the Touchscreen. In *Proceedings of the 21st International Conference on Human-Computer Interaction with Mobile Devices and Services*. 1–14.

[28] Yang Li, Samy Bengio, and Gilles Bailly. 2018. Predicting human performance in vertical menu selection using deep learning. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–7.

[29] Christof Lutteroth. 2008. Automated reverse engineering of hard-coded GUI layouts. In *Proceedings of the ninth conference on Australasian user interface-Volume 76*. 65–73.

[30] I Scott MacKenzie. 1992. Fitts' law as a research and design tool in human-computer interaction. *Human-computer interaction* 7, 1 (1992), 91–139.

[31] Jennifer Mankoff, Anind K Dey, Gary Hsieh, Julie Kientz, Scott Lederer, and Morgan Ames. 2003. Heuristic evaluation of ambient displays. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 169–176.

[32] Alex Mariakakis, Sayna Parsi, Shwetak N Patel, and Jacob O Wobbrock. 2018. Drunk user interfaces: Determining blood alcohol level through everyday smartphone tasks. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–13.

[33] Sven Mayer, Huy Viet Le, Markus Funk, and Niels Henze. 2019. Finding the Sweet Spot: Analyzing Unrestricted Touchscreen Interaction In-the-Wild. In *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces*. 171–179.

[34] Martez E Mott and Jacob O Wobbrock. 2019. Cluster Touch: Improving touch accuracy on smartphones for people with motor and situational impairments. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–14.

[35] Matei Negulescu and Joanna McGrenere. 2015. Grip change as an information side channel for mobile touch interaction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 1519–1522.

[36] Jeffrey Nichols, Zhigang Hua, and John Barton. 2008. Highlight: a system for creating and deploying mobile web applications. In *Proceedings of the 21st annual ACM symposium on User interface software and technology*. 249–258.

[37] Jakob Nielsen. 1994. *Heuristic Evaluation*. John Wiley & Sons, Inc., USA, 25–62.

[38] Peter O'Donovan, Aseem Agarwala, and Aaron Hertzmann. 2015. Designscape: Design with interactive layout suggestions. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 1221–1224.

[39] Peter O'Donovan, Aseem Agarwala, and Aaron Hertzmann. 2014. Learning layouts for single-pagegraphic designs. *IEEE transactions on visualization and computer graphics* 20, 8 (2014), 1200–1213.

[40] John C Platt and Alan H Barr. 1988. Constrained differential optimization for neural networks. (1988).

[41] Sayan Sarcar, Jussi PP Jokinen, Antti Oulasvirta, Zhenxin Wang, Chaklam Silpasuwanchai, and Xiangshi Ren. 2018. Ability-based optimization of touchscreen interactions. *IEEE Pervasive Computing* 17, 1 (2018), 15–26.

[42] Richard S. Schwerdtfeger. 1991. Making the GUI Talk. ftp://service.boulder.ibm.com/sns/sr-os2/sr2doc/guitalk.txt.

[43] Morteza Shiripour, Niraj Ramesh Dayama, and Antti Oulasvirta. 2021. Grid-based Genetic Operators for Graphical Layout Generation. *Proceedings of the ACM on Human-Computer Interaction* 5, EICS (2021), 1–30.

[44] Ray Smith. 2007. An overview of the Tesseract OCR engine. In *Ninth international conference on document analysis and recognition (ICDAR 2007)*, Vol. 2. IEEE, 629–633.

[45] Wolfgang Stuerzlinger, Olivier Chapuis, Dusty Phillips, and Nicolas Roussel. 2006. User interface façades: towards fully adaptable user interfaces. In *Proceedings of the 19th annual ACM symposium on User interface software and technology*. 309–318.

[46] Amanda Swearngin, Amy J Ko, and James Fogarty. 2017. Genie: Input Retargeting on the Web through Command Reverse Engineering. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 4703–4714.

[47] Amanda Swearngin, Chenglong Wang, Alannah Oleson, James Fogarty, and Amy J Ko. 2020. Scout: Rapid Exploration of Interface Layout Alternatives through High-Level Design Constraints. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.

[48] Alexandru Telea. 2004. An image inpainting technique based on the fast marching method. *Journal of graphics tools* 9, 1 (2004), 23–34.

[49] Kashyap Todi, Daryl Weir, and Antti Oulasvirta. 2016. Sketchplore: Sketch and explore with a layout optimiser. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*. 543–555.

[50] Daniel Vogel and Patrick Baudisch. 2007. Shift: a technique for operating pen-based interfaces using touch. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 657–666.

[51] Jacob O Wobbrock, Shaun K Kane, Krzysztof Z Gajos, Susumu Harada, and Jon Froehlich. 2011. Ability-based design: Concept, principles and examples. *ACM Transactions on Accessible Computing (TACCESS)* 3, 3 (2011), 1–27.

[52] Jason Wu, Xiaoyi Zhang, Jeffrey Nichols, and Jeffrey P. Bigham. 2021. Screen Parsing: Towards Reverse Engineering of UI Models from Screenshots. In *Proceedings of the 2021 ACM Symposium on User Interface Software and Technology*. 1–14.

[53] Tom Yeh, Tsung-Hsiang Chang, and Robert C. Miller. 2009. Sikuli: Using GUI Screenshots for Search and Automation. In *Proceedings of the 22nd Annual ACM Symposium on User Interface Software and Technology* (Victoria, BC, Canada) *(UIST '09)*. Association for Computing Machinery, New

York, NY, USA, 183–192. https://doi.org/10.1145/1622176.1622213

[54] Clemens Zeidler, Christof Lutteroth, Wolfgang Sturzlinger, and Gerald Weber. 2013. The auckland layout editor: an improved GUI layout specification process. In *Proceedings of the 26th annual ACM symposium on User interface software and technology*. 343–352.

[55] Clemens Zeidler, Gerald Weber, Wolfgang Stuerzlinger, and Christof Lutteroth. 2017. Automatic generation of user interface layouts for alternative screen orientations. In *IFIP Conference on Human-Computer Interaction*. Springer, 13–35.

[56] Xiaoyi Zhang, Lilian de Greef, Amanda Swearngin, Samuel White, Kyle Murray, Lisa Yu, Qi Shan, Jeffrey Nichols, Jason Wu, Chris Fleizach, et al. 2021. Screen Recognition: Creating Accessibility Metadata for Mobile Applications from Pixels. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15.

[57] Xiaoyi Zhang, Anne Spencer Ross, Anat Caspi, James Fogarty, and Jacob O Wobbrock. 2017. Interaction proxies for runtime repair and enhancement of mobile application accessibility. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 6024–6037.