



Optimizing User Interface Layouts via Gradient Descent

Peitong Duan
Intel AI
Santa Clara, USA
peitong.duan@gmail.com

Casimir Wierzynski
Intel AI
San Diego, USA
casimir.wierzynski@intel.com

Lama Nachman
Intel Labs
Santa Clara, USA
lama.nachman@intel.com

ABSTRACT

Automating parts of the user interface (UI) design process has been a longstanding challenge. We present an automated technique for optimizing the layouts of mobile UIs. Our method uses gradient descent on a neural network model of task performance with respect to the model's inputs to make layout modifications that result in improved predicted error rates and task completion times. We start by extending prior work on neural network based performance prediction to 2-dimensional mobile UIs with an expanded interaction space. We then apply our method to two UIs, including one that the model had not been trained on, to discover layout alternatives with significantly improved predicted performance. Finally, we confirm these predictions experimentally, showing improvements up to 9.2 percent in the optimized layouts. This demonstrates the algorithm's efficacy in improving the task performance of a layout, and its ability to generalize and improve layouts of new interfaces.

Author Keywords

Optimization; data-driven design; gradient descent; deep learning; mobile interfaces; LSTM; performance modeling;

CCS Concepts

•**Human-centered computing** → *User models*; •**Computing methodologies** → *Continuous space search*;

INTRODUCTION

User interface (UI) design is a very difficult process. There are many factors to consider, such as ensuring the UI is efficient to navigate, and that the interface is intuitive so that users can quickly figure out how to use it. A vast number of tools and techniques have been developed to aid designers in this process. They range from evaluative metrics [2] and models of human performance [16, 4] that designers can use to assess their designs, to tools and techniques that can optimize aspects of the design [27, 24, 9, 25]. Tremendous progress has been made in modeling human performance on interaction tasks, and recently, deep learning approaches have been introduced, specifically for modeling menu item selection [16] and selecting items in a grid-based interface [22]. These neural network

models are able to find complex patterns in large datasets and can be configured to account for various factors affecting task performance time, such as saliency of the target element for the task, and learning effects from completing a similar task earlier in the sequence. These data-driven models have been shown to outperform analytical models. [16]

Neural networks are differentiable. They are trained to fit a dataset via gradient-based updates to their weights that aim to minimize the difference between the model's predicted value and the observed value in the data. Similar to how a neural network is trained, gradients can also be computed with respect to the network's inputs and be used to update the input to minimize the model's predicted output. This makes neural networks a viable tool for optimization. We decided to apply these neural network models of human performance to the well-studied problem of UI optimization. Specifically, we explore the use of a task performance model's gradients to make updates to user interfaces that aim to minimize an objective function consisting of the model's predicted task completion time and error rate. Since task completion time and error rate are both useful metrics for evaluating UIs, optimizing for them may lead to an interface with better usability [20].

To perform this optimization, we first need a predictive model of task performance. To date, task performance modeling has been done on menus and grid interfaces where tapping and scrolling are the only possible interactions. Thus, we extended the model by Li. et. al. ("Deep Menu") [16] to predict task performance times on 2D mobile user interfaces given a UI and a task sequence. Our model also accounts for user error, where the completion time is increased by a penalty if users made a mistake on the task. Hence, our model predicts a metric consisting of both task time and error rate. In addition, our model supports UIs with a variety of element types, such as sliders, icons, and button groups, as well as task sequences with many different interaction types including tapping, dragging and dropping, and sliding (slider bar). Furthermore, our model also handles tasks consisting of multiple interactions. For instance, a task may require the user to tap two different UI elements in sequence. To accommodate this more complex prediction task, we increased the complexity of Deep Menu's architecture and added many more input features, including the location and size of each UI element.

To scope our work, we focus on tuning the size and location of each element in the UI to minimize task completion time and error rate. We first crowdsourced the completion times and error rates of a task sequence with 284 tasks on 108 different layout variations of a single user interface, a photo editing UI

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI '20, April 25–30, 2020, Honolulu, HI, USA.

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-6708-0/20/04 ...\$15.00.

<http://dx.doi.org/10.1145/3313831.3376589>

shown in Figure 2. We then fitted our model to this dataset, achieving an R^2 of 0.79 for the target level metric described in [16]. Then, we developed an optimization algorithm that takes in a user interface and a task sequence and makes iterative adjustments to the x,y position, width, and height of each UI element using gradient descent. We applied this optimization algorithm on various layouts of the photo editing UI. To assess the generalizability of our technique, we applied this algorithm using the same trained model to layouts of a new interface that the model has not been trained on. This is important because designers should not have to collect training data for every UI they plan to optimize. Fortunately, relationships between aspects of the layout and task performance are mostly universal (e.g. the relationship between the UI element's size and the time it takes to point to the element is governed by Fitts' Law [8]), so our model should be able to infer these relationships from the dataset and apply it towards optimizing a new interface layout.

This optimization technique produced layouts that have better predicted task performance for both the photo editing UI and the new UI. To verify this experimentally, we crowdsourced task completion times and error rates for a few initial and optimized layouts of both UIs. The observed task performance metric also showed improvements in the optimized layouts, with improvements up to 9.2 percent. These results demonstrate our optimization algorithm's ability to make effective improvements to a layout, as well as its ability to generalize to new interface layouts. From a practical standpoint, our system can facilitate the design process. A designer could start with a set of hand-crafted candidate layouts and use the model to compare their task performance. The designer could also use the optimizer to improve their layouts, generating layout alternatives with better task performance.

To summarize, we made the following contributions:

- An extension of the model from Deep Menu to predict task performance of various task types for 2D mobile UI's. We crowdsourced a dataset of task performance times from 379 participants and evaluated our model on this data.
- A new technique for optimizing the layout of mobile UI's using the gradients of a trained task performance predictor network. This technique is generalizable and can improve layouts of a new UI that the model has not been trained on.

RELATED WORK

Substantial work has been done in both predicting human performance on interaction tasks and optimizing user interfaces.

Modeling Human Behavior

Modeling human behavior started with simple analytical models that focused on very specific aspects of human performance in isolation. For instance, the well-known Fitts' Law predicts how long it takes users to point to a visual target as a function of distance to the target and target width [8]. However, this does not include other factors that affect task performance, such as learning effects from past tasks and visual search time for the target, which have been accounted for by a model on menu item selection proposed by Bailly et. al. [4].

Recently, there has been a shift towards using neural networks to model human behavior, as deep learning models can discover complex patterns in data and do not require extensive feature engineering that is often necessary for analytical models. In particular, Li. et. al. [16] and Swearngin et. al. [26] collected large datasets via crowdsourcing on Amazon Mechanical Turk, and then fitted neural network models to their datasets. Li. et. al.'s model (Deep Menu) takes in a menu and a sequence of menu items to select, and predicts the time for each selection. Their model incorporates factors that affect visual search and utilizes recurrent layers to capture learning effects. Swearngin et. al.'s model takes in a mobile UI and an element in the UI and predicts whether users would view the element as tappable. Pfeuffer et. al. also collected a dataset via a 20-user study and fitted a neural network model to predict the time it takes to tap on items in a scrollable mobile grid layout [22]. To our knowledge, no one has used deep learning to predict task performance for a general mobile layout, as well as for interaction types beyond tapping and scrolling. Our model expanded the range of interactions to include drag and drop and sliding. We also extended task modeling to incorporate tasks with multiple interactions.

UI Optimization

Because UI design is complex and multifaceted, many techniques have emerged to optimize UI designs, and several different metrics have been used as the objective function. For online games, the goal is usually to maximize user engagement, or how long users spend playing the game. Bayesian optimization [13] and multi-arm bandits [17] have been used to tune features of the game, including font-size and how users enter in input, to maximize user engagement. For general user interfaces, Krzysztof et. al. created a system (SUPPLE) that automatically generates interfaces, optimizing for a complex function estimating the user effort for a given user trace [9]. For optimizing the UI layout specifically, Quiroz, et. al. used genetic algorithms to evolve the layout of a single UI [24]. The color and location of each element in the interface can be evolved, but the layout was restricted to a grid, and changes to the layout consisted of swapping locations of elements within the grid. Furthermore, human input was required to compute the fitness of the layouts in each generation.

A common technique is to optimize for an objective function combining many layout metrics (e.g. visual clutter). This technique was used in the studies Sketchplore [27] and AIDE [25]. Sketchplore is perhaps most similar to our work, in that it optimizes an objective function accounting for both usability and aesthetics. Sketchplore's usability component consists of a weighted summation of an analytical model for visual search and Fitts' Law for target acquisition. Our data-driven model expanded upon this by including additional factors that affect task performance. For instance, like Deep Menu, our model also takes in the semantics and saliency of the text labels on each UI element, which should have a strong effect on visual search [18], whereas Sketchplore does not factor in text labels in their visual search component. In addition, our task performance predictor also accounts for errors users may make and differentiates amongst types of user interactions, such as drag and drop and sliding, whereas Sketchplore's usability

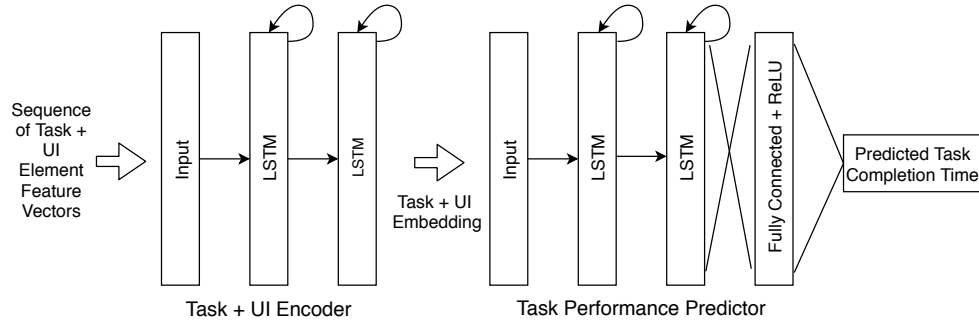


Figure 1. The Task Performance Predictor model's architecture.

model groups all user interactions as visual search and target acquisition.

Until now, no work has been done to apply gradients of a deep neural network model that predicts human task performance to tune the layout of a user interface for better task performance. Li et. al. did compute gradients of task completion time with respect to input features, but it was to study the model's memory effects. Our work is also the first to use deep learning to predict task completion time and error rate of a general 2D user interface with many different element and interaction types, and also handles multi-step tasks.

TASK PERFORMANCE PREDICTION MODEL

We first expanded the Deep Menu model [16] to predict task performance for mobile user interfaces. Like [16], the mobile task performance predictor takes in a UI and a sequence of tasks performed on the UI and predicts the completion time of each task. We also increase the task completion time with an error penalty, which is described later. Our model accounts for interfaces with a wide selection of individual and grouped element types. Grouped element types (e.g. a group of icons) consist of a set of elements of the same type that are arranged in a rectangular container (see Figure 2, Section E). Individual element types are single elements, such as a slider bar (see Section C). Our model also handles different interaction types, namely tapping, drag and drop, and sliding. We draw largely from Deep Menu's model architecture, and most of our modifications are in the input features.

Model Architecture

Like Deep Menu, we utilize the LSTM's capabilities of learning and remembering information from the input task sequence. We have an encoder network that generates an embedding for each task. The embedding is then input into a predictor network, which outputs a prediction for the task's performance time. The network's hierarchical architecture is depicted in Figure 1. Since the UI may change from the user's interactions, features of the UI are fed into the encoder (along with the task information) for every task embedding. In particular, each element in the UI is represented by a fixed length vector, which also contains information about the task. Details for each feature in the vector are provided in the next section. A sequence of these feature vectors, one for each UI element, is input into the encoder network to generate the task embedding. The feature vectors are ordered by the location

of the top left corner of each element in a top-down, left-right manner. The predictor model takes in a sequence of task embeddings and generates a prediction for the completion time of each task. We provide a detailed figure illustrating this feature encoding process in the supplementary materials, and [16] also presents an explanation for feature encoding.

As shown in Figure 1, the encoder and predictor both have recurrent layers that account for previous tasks in the sequence while predicting the completion time of the current task. This captures the learning effect of users taking less time on tasks as they become more familiar with the UI. Since the tasks and UI for our model are more complex than menu item selection, we have two LSTM layers in both the predictor and encoder models compared to the single recurrent layers in Deep Menu's. The recurrent layers in our predictor model are followed by a feed-forward hidden layer with a ReLU activation function, and the final time prediction is a linear combination of this feed-forward hidden layer; this follows Deep Menu's architecture.

Model Features

The feature vector for UI element j of task s is as follows:

$$e_s^j = [\text{target}, \text{len}(\text{name}), \text{word2vec}(\text{name}), x, y, \text{width}, \text{height}, \text{orientation}, \text{container_x}, \text{container_y}, \text{container_width}, \text{container_height}, \text{element_type}] \quad (1)$$

These features are selected because they may impact task performance. The first three features (target , $\text{len}(\text{name})$, and $\text{word2vec}(\text{name})$) are taken from [16], which discussed their effect on task performance. Our target deviates from the definition in [16], and is instead a one-hot vector of length 3 indicating if the UI element is the target for an interaction (e.g. the specified button to be tapped), the drop or sliding target, or not a target. For drag and drop and sliding interactions, one UI element is the target being dragged or slid, and another is the drop target or sliding destination. Hence, these two elements must be differentiated. Likewise, all other categorical features are represented as one-hot vectors. $\text{len}(\text{name})$ and $\text{word2vec}(\text{name})$ are as defined in [16]. $\text{len}(\text{name})$ is the length of the text label on the UI element and represents its visual salience. $\text{word2vec}(\text{name})$ is the word2vec embedding (reduced to length 4) that captures the semantics of the element's text label. However, in our case, the element may have a symbol or image instead of a text label. In this situation, we

would take the *word2vec* embedding of the word most closely represented by the graphic (e.g. the embedding for "undo" is used for the undo icon). $len(name)$ for icons and images is set to a value suitable for its visual saliency. These features are normalized to a value between -1 and 1.

An element's location and size affect pointing time [8] and visual search [4]. Furthermore, the spatial relationships between two related elements affect the performance of tasks requiring interactions with both elements. Hence, we provide the x and y location of the element's (or its bounding box's) center and the element's *width* and *height*. We also include location and size of the grouped element's container (via *container_x*, etc.) In addition, for grouped elements, each element in the group has an input vector because individual members are usually the task target, as opposed to the group itself. For instance, users would tap one of the buttons in a button group. These spatial features are all in pixels that are then normalized to be between 0 and 1 by the screen width or height. *orientation* specifies the orientation of the element, which can be vertical, horizontal, or not applicable, since some interactions may be affected by the target element's orientation. For instance, sliding a horizontal slider requires a different motion from sliding a vertical slider and may take a different amount of time to perform. Finally, the element's type (slider, button, icon, etc) is given (*element_type*) since it determines the element's visual salience, as well as how users would interact with it.

We then extend our task encoding with the following task-specific features: $[interaction_type, step, total_steps]$. *interaction_type* indicates the type of interaction required by the task (tapping, drag and drop, etc.), since different interaction types require different gestures (e.g. tapping vs sliding) which should affect completion time. Our model also support tasks consisting of many interactions (e.g. a task may require tapping two different UI elements). These multi-step tasks are presented as a single task to the users, who would have to figure out the individual interactions. Since these multi-step tasks require more cognitive effort from users, they should not be modelled as a series of individual single-interaction tasks, where each step broken down and presented to the user. Our approach to modeling multi-step tasks are as follows: generate a task embedding for each interaction in the task, and use *step* and *total_steps* to identify the embedding as part of a multi-step task. For instance, if a task requires users to first tap element *A* and then tap element *B*, a task embedding will first be generated for tapping element *A* with the features $[tap, 1, 2]$ appended, where $[1, 2]$ specifies that tapping element *A* is step 1 of a 2-step task. Similarly the features $[tap, 2, 2]$ are appended to the embedding for tapping element *B*. For single-interaction tasks, *step* and *total_steps* are both set to 1.

Drag and drop interactions, as well as sliding (the slider handle), involve first acquiring the target and then performing the interaction. Hence, these interactions consist of two steps and are modelled as a 2-step task where the first step is target acquisition and the second is the actual drag and drop or slide.

DATA COLLECTION AND EXPERIMENTS

In this section, we discuss the dataset we used to evaluate our model, the set-up for crowdsourcing this data on Amazon Me-

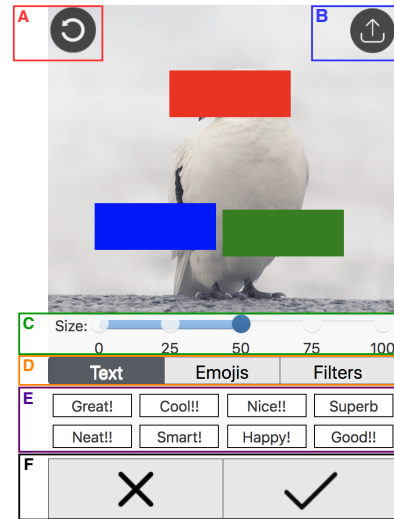


Figure 2. The photo editing UI with all the UI elements labelled. A is the undo icon, B is the upload icon, C is the slider, D is the button group that controls which set of stickers are displayed, E is the set of stickers (icon group type), and F is the button group with the save (checkmark) and cancel ("X") buttons. The colored rectangles in the photo are not part of the UI; they are the drop targets for Task Type 4 (drag and drop).

chanical Turk (mTurk), and the results of our data collection. To scope our data collection in a way that is useful for layout optimization, we collected data for various layouts of a single user interface. Our data collection protocol was subject to institutional review as detailed in the supplementary materials.

Dataset

The dataset consists of 108 different layouts for a user interface, with examples shown in Figure 3. This UI allows users to add stickers and filters to a photo, and consists of the following types of UI elements: button groups (save/cancel buttons and the "Text"/"Emoji"/"Filter" buttons, Sections F and D of Figure 2), icons (undo and upload icons, Sections A and B), icon groups (the stickers, Section E), slider bars (Section C), and static divs (the photo). An icon group is defined as a group of icons confined in a rectangular container with as many icons placed on a row as possible, with uniform and maximal horizontal and vertical spacing between adjacent icons, as shown in Figure 2, Section E. Button groups are defined similarly. Icons and buttons are both tappable, but an icon must have a fixed width to height aspect ratio, which is preserved across all 108 layouts. The different layouts in this dataset vary in the size, location, and orientation of each UI element as shown in Figure 3.

Task Sequence

Possible user interactions with this photo editing UI include tapping a button or icon, adding a sticker or filter to the photo by tapping on it, adjusting the size of added sticker with the slider, and dragging and dropping the added sticker to a section of the photo. A comprehensive task sequence is created with tasks consisting of individual and combinations of these interactions. Specifically, the different types of tasks include:

1. Selecting a sticker or filter to add to the photo

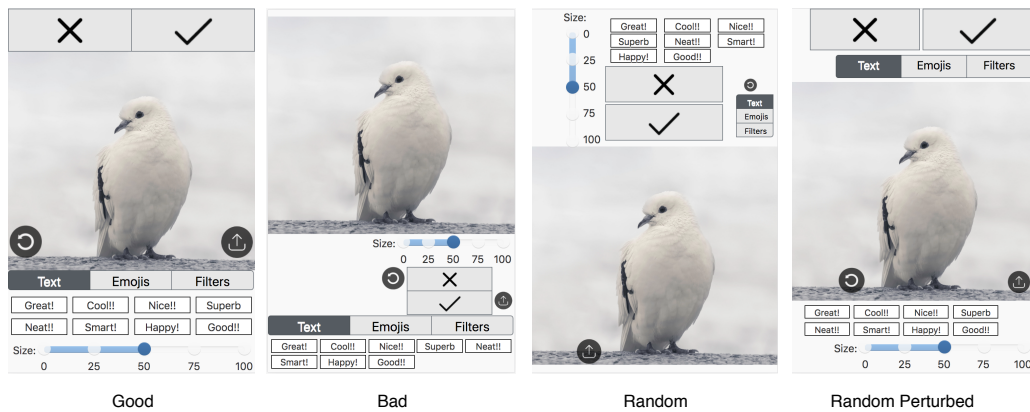


Figure 3. Examples of each layout type in the dataset. The "Random Perturbed" layout is a random perturbation of the "Good" layout.

2. Tapping the appropriate "Text"/"Emoji"/"Filter" button to open the appropriate set of stickers and then selecting the target sticker
3. Adjusting the size of the added sticker with the slider
4. Dragging and dropping the added sticker to a specified target in the photo. The drop targets are colored rectangles as shown in Figure 2.
5. Tapping one of the icons (the undo or upload icon) or one of the buttons (the save or cancel button)

These task types are relatively simple to ensure all mTurk workers perform the same set of interactions in the exact same order as they work through the tasks. This is to guarantee consistent and accurate modeling of each task. We are able to incorporate more complex tasks that require two consecutive interactions (e.g. Task Type 2), which are more difficult to figure out and can be used to assess how intuitive the layout is. For Task Type 4 (drag and drop), the locations of its drop targets are changed (randomly) for every photo to simulate realistic usage of the UI. For Task Type 3, the sliding destination for the slider handle is a specified range on the slider bar (e.g. a value between 50 - 75). Each element in the interface is interacted with at least five times in this task sequence to capture both the learnability of the interface, which is reflected by the first time each task is performed, and the efficiency of the interface, which is shown in later repetitions of the task as the user becomes more familiar with the UI. To simulate realistic usage of this photo editing app, the mTurk participant works through 20 photos in the task sequence and saves or cancels their edits to the current photo before moving onto the next one. This results in a task sequence of 284 tasks.

Generating Random Designs

Out of the 108 different layout of the photo editing UI, 5 are manually designed to meet many of the design guidelines specified by Apple [12] and Nielsen [19] and 3 are manually designed to be bad and violate many of these design guidelines. Figure 3 shows an example of a good and bad layout, and all 8 layouts were designed by the authors. To ensure good coverage of the layout space, the remaining 100 layouts are generated with the layout parameters randomized or by randomly perturbing one of the good designs.

Fifty designs were generated by randomizing the layout parameters directly. The location, size, and orientation parameters of each UI element are first randomized and the elements are added to the interface one by one, checking for overlap. If the current element being added overlaps with any of the elements already in the UI, its parameters are rerandomized. Otherwise, the element is added to the interface. For grouped elements (e.g. icon groups), the size and location parameters of the container are also randomized. The remaining 50 designs were generated by randomly perturbing each of the 5 good layouts. Since the other set of randomized layouts likely violate design guidelines and are considered bad, only good layouts are perturbed to incorporate more somewhat good layouts to the dataset. For these perturbed layouts, the size of each element is adjusted by a random factor selected uniformly from the interval $[0.7, 1.3]$. Since there is not much white space in these good layouts, the locations are perturbed by randomly swapping adjacent elements with a probability of 0.15. Figure 3 shows examples of these random and perturbed layouts.

Data Collection App

We built a web application using the psiTurk API to crowd-source task completion times and error rates of the task sequence for each layout. psiTurk provides a backend API for recording data and a command line interface to recruit workers from Amazon Mechanical Turk [10]. Following the data collection procedure from [16], the task sequence is presented to the mTurk worker in the following manner: the worker first sees the instructions for the task, and when they are ready to complete the task, they tap the start button and are taken to the UI to complete the task. Our data collection app is described in detail in the supplementary materials, which includes examples of task instructions shown to workers and details on how we handled and recorded the workers' errors.

Data Collection Results

All workers were assigned the same task sequence to work through on one of the 108 different layouts, with at least 3 workers assigned to each. In total, there were 379 participants from Amazon Turk. There were 151 males and 228 females, and around 12 percent of the users were left-handed. In total, we collected completion times of $379 \times 284 = 107,636$ tasks.

The task completion times for each task were first averaged across all workers for that layout; only completion times where the task was completed correctly were considered. For each task and layout, we remove outliers whose distance is greater than 1.5 median absolute deviations (MAD) from the median. We use MAD because it is more robust in detecting outliers, compared to standard deviation [15]. To incorporate errors, the performance metric assigned to each task and layout is equal to the averaged time for each task per layout increased by an error penalty as shown in the following equation:

$$task_perf_metric = avg_task_time * (1 + 0.5 * frac_err) \quad (2)$$

where *frac_err* refers to the fraction of workers who made an error on that task out of those who were assigned to that layout. For more severe errors where users incorrectly tapped the save or cancel button, which means they had to redo all the tasks for the current photo, the error penalty was increased from 0.5 to 0.8 and *frac_err* becomes the fraction of workers who erroneously tapped the save or cancel button. Since these error penalties increase fluctuation in task performance, they decrease the model's prediction accuracy. We carefully tuned these penalty constants to maximize emphasis on errors without sacrificing significant prediction accuracy.

As a sanity check, we computed the average task performance metric for each category of layouts: 564.0 (good, std.err.=10.7), 638.8 (bad, std.err.=46.7), 588.2 (random, std.err.=9.7), and 580.7 (random perturbations of good layouts, std.err.=9.1). As expected, the good designs had on average, a lower value for the task performance metric compared to the bad and random categories, with the bad category having the highest averaged value.

Model Performance Results

This section describes how we configured the model to fit the dataset we collected, and presents the results of an evaluation of the model's accuracy. We also computed the fraction of workers (assigned to each layout) who were left-handed and the average of their ages and appended these statistics to the task embedding. The majority of users would likely interact with this photo editing UI using their dominant hand, as shown in the results of a phone grip study on a focused task by Eardley et. al. [7]. Since certain elements are located at regions that are more difficult to reach for left-handed users and vice versa, a user's handedness may have an impact on task performance. Furthermore, a person's age is correlated with their UI design preferences, familiarity with mobile technology, and gesture mobility [23], all of which impact task performance.

Model Configuration and Loss Function

There are 8 different element types in the photo editing UI, which means the *element_type* feature from Equation 1 is a one-hot vector of size 8. This results in the per-element features vectors defined in Equation 1 to have size 27. The two recurrent layers in the encoder each have 23 LSTM cells. The predictor network has 30 LSTM cells in its two recurrent layers, followed by a feed-forward layer of size 28 to compute the task completion time. To regularize the model and prevent overfitting, we applied a dropout probability of 0.1 to the task

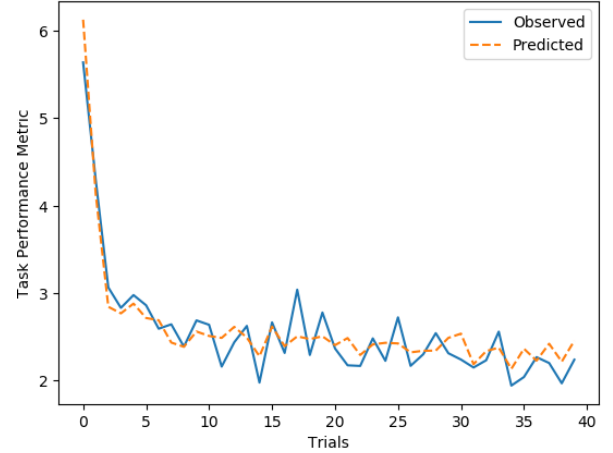


Figure 4. A plot showing the model's accuracy. This graph contains the predicted and observed task performance values for Task Type 3 (sliding), across trials.

embedding and a dropout probability of 0.4 to the feed-forward layer in the predictor (Figure 1).

We use the same loss function as [16], which is defined as

$$L_s = \frac{\sum_{i=1}^{|S|} (y_i - t_i)^2}{\sum_{i=1}^{|S|} (y_i - \bar{y})^2} \quad (3)$$

where $|S|$ refers to the length of the task sequence, y_i is the observed completion time of task i , and t_i is the predicted task time. \bar{y} is the average observed task completion time in sequence S , so the denominator $\sum_{i=1}^{|S|} (y_i - \bar{y})^2$ is the variance of the task performance times in the sequence. R^2 is a standard metric to assess a model's prediction quality and measures the correlation between observed and predicted sequences. This loss function is related to R^2 via the equation $R^2 = 1 - L_s$. Since we minimize the loss function L_s during training, we would be maximizing R^2 .

Our model was implemented in PyTorch, a deep learning framework for Python [21]. We trained the network using the Adam optimization algorithm [14] to minimize the loss function with a learning rate of $3e-4$ and a batch size of 8. We also clipped gradients so their norms do not exceed 1.0.

Results

We evaluated the model using 6-fold cross validation trained for 850 epochs at each fold, and computed the R^2 using the target-level R^2 defined by Bailly et. al. [4] and used by Li et. al. to evaluate Deep Menu. This target-level R^2 metric examines the relevant task performance for each UI element with varying amounts of practice (trials). Our model achieved a target-level accuracy of 0.79, averaged across all 6 folds. This is comparable to the target-level R^2 of 0.76 achieved by Deep Menu on their datasets. Figure 4 shows a plot of our model's predicted task performance for Task Type 3 (sliding) and the observed task performance, across trials, demonstrating the model's prediction accuracy.

UI LAYOUT OPTIMIZATION

Once we have a trained model to predict task performance, we proceed with layout optimization. We use the gradient descent algorithm, which makes iterative updates to the input based on computed gradients (of the objective function with respect to input) that minimize the objective function. In our case, the input is the location and size of each element in the layout, and our objective function is the model’s predicted task performance. This optimization algorithm is described in detail in the next sections and supplementary materials.

We then used this optimization algorithm to improve layouts of the photo editing UI and another UI that the model has not been trained on, since this technique should be able to generalize and improve layouts of new interfaces to be useful. Designers should not have to collect task performance data for every UI they hope to optimize; the trained model should be able to transfer patterns it learned from one interface to optimize the layout of another. We applied this optimization algorithm on several layouts of both interfaces, which led to predicted improvements in task performance. To verify actual improvements in human performance, we crowdsourced task completion times and error rates for the initial and optimized layouts. The optimized layouts also show improvements in observed task performance for both interfaces.

Optimization Algorithm

Optimization algorithms aim to minimize an objective function, and gradient descent is a particular optimization algorithm that is commonly used to train neural networks. Given objective function f and input x , a single update at step n is given by the following equation:

$$x_n = x_{n-1} - l_r \nabla f(x_{n-1}) \quad (4)$$

where l_r is the learning rate that controls the update step size. In our case, the objective function is the sum of the predicted completion times (with error penalty) of all tasks in a task sequence plus the penalty functions, and the inputs are the x , y , *width*, and *height* of each element, as well as its corresponding container dimensions. Each layout feature of every element in the UI is updated with their respective gradients individually, following Equation 4. For each UI element, its features (see Equation 1) are input into the model for every task in the sequence, which means there will be gradients for the element from each task. Since the size and location of every element in the UI must remain consistent throughout the task sequence, the average is taken over all the element’s gradients across tasks, and a single update is made using Equation 4 with this average gradient. The objective function F is formally defined for layout l by the following equation:

$$\begin{aligned} F(l) = & \text{task_seq_perf}(l) \\ & + (\text{penalty_constant})(\text{overlap_penalty}(l)) \\ & + (\text{penalty_constant}')(\text{boundary_penalty}(l)) \\ & + (\text{penalty_constant}'')(\text{additional_penalties}(l)) \end{aligned} \quad (5)$$

Where *task_seq_perf* is the sum of the predicted task performance of all tasks in the sequence, *overlap_penalty* is a differentiable function that is positive if there are overlapping UI elements in layout l and 0 otherwise. Likewise,

boundary_penalty is a differentiable function that is positive only if any element in l exceeds the boundary of the user interface. *additional_penalties* refers to penalty functions from additional constraints the designer hopes to enforce, e.g. ensuring two particular elements are aligned. The penalty constants (e.g. *penalty_constant*) are constants that add high values to the objective function if their corresponding penalty functions are positive. Furthermore, at each update step, if two elements end up overlapping, their locations are swapped if their gradients indicate that may lower the objective function.

Penalty Functions

Since each UI element’s layout features are updated independently, elements may inevitably overlap with each other or go out the boundary of the interface after updates. Penalty functions can prevent these undesirable conditions by adding a large value to the objective function when these conditions are detected. The penalty function would then steer the gradients towards eliminating these conditions so as to remove the large penalty value. Since penalty functions must contribute to gradients of the objective function to have an effect, they must be differentiable. We use the rectified linear (*ReLU*) function, which is defined as

$$\text{ReLU}(x) = \max(0, x) \quad (6)$$

ReLU(x) is 0 when x is negative and is equal to x otherwise. The gradient of *ReLU*(x) is 1 for positive x and is 0 otherwise. This satisfies the conditions for our penalty function, since it will not contribute to the objective function nor the gradients when it is not activated (equal to 0), and will steer the gradients away from overlaps or other undesirable situations when activated.

The overlap and boundary penalty functions we used are provided, along with derivations, in the supplementary materials. In addition to the necessary overlap and boundary penalty functions, designers can add more constraints to improve the output of the optimization. For instance, penalty functions can be enforced for two UI elements to have the same size, or to introduce a minimum size limit for an element. Furthermore, if a designer wants to group two elements together, such as the sticker button group (Figure 2, Section D) with the stickers (Section E), they can introduce a penalty function to ensure the two elements are in close proximity, aligned horizontally or vertically, and have the same dimension size along their bordering sides (e.g. the sticker button and stickers in Figure 2 should have the same *widths*). The equations for all these penalty functions are detailed in the supplementary materials. In sum, penalty functions can be added to ensure any desired characteristics in the optimized layout. The penalty constant of each constraint can also be tuned based on importance, with more important constraints being given a larger constant. This is especially useful in the case of conflicting constraints.

Swapping Locations of UI Elements

Since the overlap penalty function will push UI elements away from each other if they overlap, elements will be restricted to their initial regions throughout the optimization. For instance, the sticker group (Figure 2, Section E) is constrained to be below the sticker button group (Section D) because whenever the

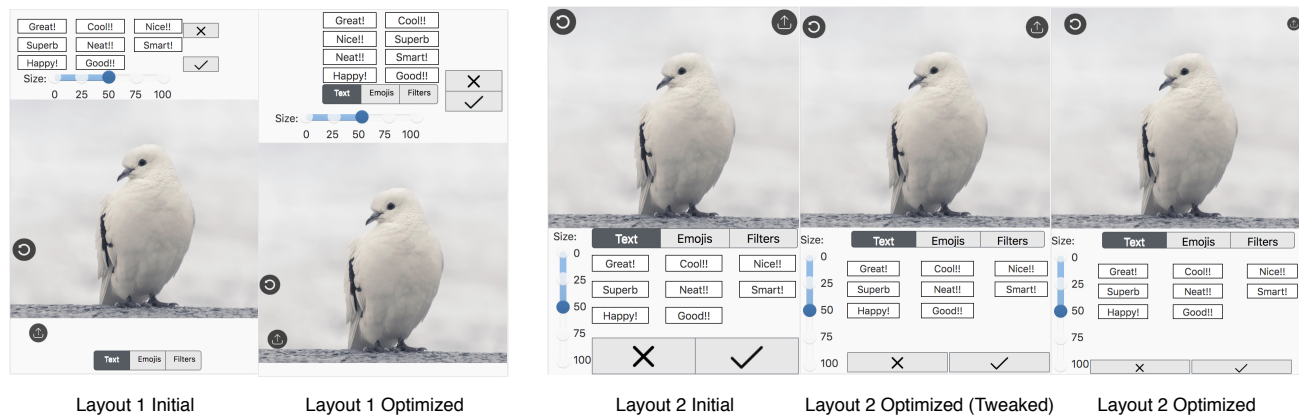


Figure 5. The initial and optimized layouts of two different initial layouts. **Layout 1** was optimized with strict constraints, and the predicted task performance improved by 6.3 percent. "Layout 1 Optimized" shows algorithm's output. **Layout 2** was optimized using relaxed constraints. The layout directly outputted by the algorithm ("Layout 2 Optimized") had a predicted improvement of 7.1 percent, and the optimized layout after some minor tweaking ("Layout 2 Optimized (Tweaked)") had an improvement of 4.7 percent.

sticker group overlaps with the sticker button group, the sticker group will be pushed back down. To enable exploration of layout alternatives with a much larger location space, we introduce location swapping of overlapping elements. Specifically, when two elements overlap, we look at their gradients of the objective function (Equation 5) without the overlap penalty. If the two elements' gradients indicate that swapping their locations may lower the objective function, the swap is performed. A rigorous definition of this swapping condition and details of the location swap can be found in the supplementary materials.

Optimization Results

This section contains representative samples of optimization results for the photo editing interface and a new UI that the model has not been trained on. Additional examples of initial and optimized layouts for both interfaces can be found in the supplementary materials. We also crowdsourced human performance data to verify actual improvement in task performance for the optimized layouts.

Optimization Set-up

While optimizing the layouts of both user interfaces, we used values of 10000 for the overlap and boundary penalty constants and smaller values for less serious penalties (e.g. grouping of two elements.) We used a learning rate of 0.05 for the gradient descent. We also clipped gradients to have norms at most 0.5 to prevent extreme shifts in the layout from a single update. We ran gradient descent for 500 steps, and saved a CSS file of the layout at each step. We then took the layout with the best predicted task performance. In order to generate layouts with better task performance across the entire population, we passed in the average age of all adult iPhone users (37.7) [1] and the average fraction of left-handed people in the population (0.1) [11] to our model.

Photo Editing UI

For the photo editing UI, we first experimented with optimization using several stringent penalties to ensure good output layouts. These penalties include large minimum size constraints for all UI elements, identical sizes between the undo

and upload icons (Sections A and B of Figure 2), and grouping the sticker button with the stickers (Sections D and E of Figure 2). These penalty functions are described in "Additional Penalty Functions". Figure 5 shows an example of the initial (Layout 1) and optimized (Layout 1 Optimized) layout, along with the predicted task performance improvement. An improvement in task performance refers to a decrease in value of the task completion time with error penalty metric. The improvement in task performance for Layout 1 is likely due to the enlargement of the stickers and other UI elements, which makes tapping easier according to Fitts' Law. Furthermore, the sticker buttons (Figure 2, Section D) in the optimized layout are adjacent to the stickers (Figure 2, Section E), making the layout more intuitive. This is because sticker buttons control which stickers are displayed, so their relationship is more evident in the optimized layout. Furthermore, this should make tasks where users must first open the appropriate set of stickers before selecting a sticker to add (Task Type 2) more efficient since the two elements are closer together.

We also applied our technique on a layout that was initially good (Figure 5, Layout 2). We did not see much improvement when we performed optimization with all the stringent constraints described earlier. Since these strict constraints limit exploration, we removed those constraints and just set low minimum size requirements. With this, we were able to achieve a predicted improvement of 7.1 percent in the total task performance metric for the optimization's output layout (see Figure 5, "Layout 2 Optimized"). However, some of the elements are too small (e.g. the save and cancel buttons) and misaligned. We can manually enforce the original stringent constraints and align the elements, resulting in a more visually appealing layout (see "Layout 2 Optimized (Tweaked)"). These alignment and resize tweaks to the output are very simple and straightforward for the designer to apply, and the tweaked output still had an improvement of 4.7 percent. The slight decline in predicted task performance after tweaking (compared to the algorithm's output layout) could be due to the fact that misaligned elements stand out and are hence, easier to find [6]. The increased spacing between some elements (e.g. the stickers and exit buttons (Sections E and F of Figure 2))



Figure 6. The initial layouts and the optimized layouts after minor tweaking for the recipe planning UI. Layout 1 had a predicted task performance improvement of 3.2 percent, and Layout 2 had a predicted improvement of 5.4 percent. The algorithm's output layouts can be found in Figure 7.

in the tweaked output layout (compared to the initial layout) likely caused the improvement in predicted task performance.

To verify that our optimized layouts have better human performance, we crowdsourced task completion times and error rates of the initial and optimized layouts for both layouts in Figure 5. For Layout 2, we collected data for the "Optimized Tweaked" layout instead of the "Optimized" layout. We assigned at least 10 mTurk workers to each layout. The crowdsourced results confirm improvement in task performance in the optimized layouts. The observed performance increase (decrease in the value of the task performance metric) is 8.9 percent for Layout 1 (6.3 percent predicted) and 2.0 percent for Layout 2 (4.7 percent predicted). Layout 2 likely had a smaller improvement than Layout 1 because Layout 2 was initially quite good. To verify statistical significance of the observed performance improvement, we computed a two-sample t-test for the observed task performance metric for the initial and optimized versions of both layouts. For Layout 1, $t(13) = -16.674$, $p < 1e-13$ and for Layout 2, $t(10) = -2.780$, $p < 0.05$, which indicate statistical significance of the observed improvements.

New UI: Recipe Planner

The recipe planner is an interface where users drag and drop ingredients that they like and dislike to the appropriate boxes (see Figures 6 and 7). Once users have made all their selections, they tap the "Get Recipe" button to get a list of recipes they may like based on their ingredient preferences. All element and interaction types in this recipe planning UI are found in the photo editing UI. This is necessary in order to utilize the model trained on the photo editing UI for optimization on this new UI, as the model would probably not be able to make accurate predictions on new element and interaction types. However, the UI element and interaction types in the photo editing UI are very general and can be used to build a wide variety of user interfaces. We created a comprehensive task sequence for the recipe planning interface with which we will optimize for task performance. The task types include tapping the "Get Recipe", undo, cancel, or one of the ingredients buttons ("Grains", "Fruits", "Veg."), dragging and dropping a specified ingredient sticker to the "Ingredients You Like" or "Ingredients You Dislike" box, and a two-step task, where the user must first open the appropriate set of ingredients stick-

ers via tapping the corresponding ingredients button and then dragging and dropping the target ingredient to the appropriate box. This task sequence for the recipe planner UI had 136 tasks.

We first tried optimizing the layout with strict constraints, such as large minimum size constraints for the UI elements. However, these strict constraints did not lead to an improvement in predicted performance of the task sequence. We then relaxed these constraints to include only low minimum size constraints. This resulted in improvements in predicted task performance of the task sequence, even after the optimization output has been tweaked to align elements and enlarge elements that were too small. The optimization results of an initially bad layout (Layout 1) are shown in Figures 6 and 7. For Layout 1, the optimization algorithm moved the ingredients stickers closer to the "Ingredients You Like" and "Ingredients You Dislike" drop targets and added more spacing between the undo and cancel icons in Layout 1 that were initially too close together. We also optimized a layout that was initially good (Layout 2), and Figures 6 and 7 show the results. For Layout 2, the algorithm decreased the distance users had to drag the ingredient stickers to the drop targets, and also moved the "Get Recipe" button and the undo icon farther to the right, which makes it more accessible to right-handed users, who make up the 90 percent of the user base. To verify that the optimized layouts have improved human task performance, we crowdsourced task completion times and error rates for Layouts 1 and 2 following the same procedure used for the photo editing UI. There is an observed improvement of 9.2 percent ($t(12) = -9.799$, $p < 1e-8$) and a 3.2 percent predicted improvement for Layout 1, and an observed improvement of 4.9 percent ($t(13) = -5.622$, $p < 1e-5$) and a 5.4 percent predicted improvement for Layout 2.

DISCUSSION

Our optimization algorithm was able to discover layout alternatives with better task performance for both user interfaces, which demonstrates the algorithm's generalizability. Hence, our trained model would likely be able to improve the layout of any interface containing UI elements and interactions found in the photo editing interface, which is a very large set of possible UI's. Furthermore, according to Arhippainen et. al. [3], the

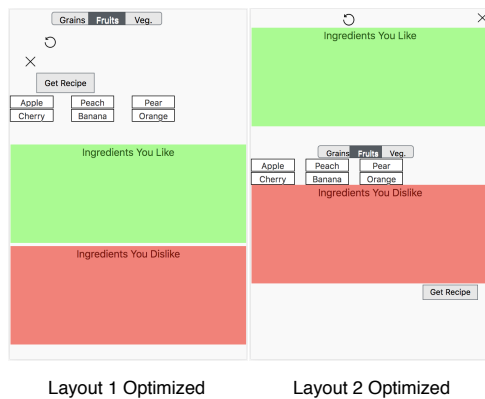


Figure 7. Layouts directly outputted by the algorithm. For Layout 1, the algorithm made a 5.8 percent improvement (predicted) in task performance, and for Layout 2, the algorithm made a 4.1 percent improvement. These output layouts were tweaked slightly for alignment and re-sizing. Figure 6 shows the tweaked layouts.

only layout-relevant factors affecting general user experience are usability and aesthetics. Since task performance is a good indicator of usability [20], the optimized layouts will likely result in better user experience if aesthetics are not worsened.

The findings of our optimization experiments support those by Lomas et. al. [17]. Their paper emphasized optimizing for the right metric, and while our optimization technique returned layouts with better task performance, the layouts did not show much improvement in terms of aesthetics. Hence, our objective function could have been more comprehensive and incorporated metrics for the visual aspect of the layout. Furthermore, as shown in the layouts directly outputted by the optimization algorithm (Figures 6 and 7), the model made some elements very small to possibly reduce user error. Lomas et. al. warned that the optimization output may be unsatisfactory when metrics are optimized to the extreme.

Human-AI Collaboration

These limitations with our optimization algorithm can be fixed easily by humans, as shown in Figures 6 and 7. Both human judgement and the optimization algorithm can be misleading, so ideally humans and AI would collaborate in the design process [17]. The following illustrates how this hybrid workflow might work with our system. Designers can start by building an initial layout and specifying any constraints in the optimization output, such as elements that must be grouped together. Then, our algorithm returns a layout that the designer can refine for aesthetic quality and other features. The designer can then use the task performance predictor to compare the task performance between the refined and initial layouts. Furthermore, the designer can run the optimization again on the refined layout to further improve its task performance, and this iteration between AI optimization and human refinement can be continued until a satisfactory layout has been achieved. Furthermore, the optimization algorithm outputs the CSS file for the layout at every step, which generates a large set of layout alternatives. If the designer sees a layout they like better than the one with the best predicted task performance, they can work with their preferred layout instead. In sum, there are

many ways for our optimization algorithm to assist designers in creating a great layout for their user interface.

Limitations and Future Work

There are three main limitations to our work. First, our optimization algorithm focuses primarily on usability, whereas layouts have other important attributes, such as aesthetics. As explained earlier, less effort will be required from designers if the algorithm's output had better aesthetics. Sketchplore [27] accounted for both usability and aesthetics in their optimization, including aesthetics metrics for visual clutter, grid quality, and alignment. Furthermore, our method only modifies the sizes and locations of UI elements, while other layout optimizers, such as [27], also adjust the elements' colors. Second, many of the tasks in the dataset's task sequence are simple and require only one interaction. Although we did have 2-interaction tasks, including even more complex tasks could result in stronger emphasis on the layout's learnability in the task performance metric. Third, while optimization using our trained model would likely generalize to other mobile UIs with element and interaction types found in the photo editing UI, more training data would probably be necessary for interfaces with other UI element and interaction types (e.g. swiping) and for other platforms (e.g. computers). Also, the model's input features and architecture may need to be modified for these cases.

These limitations suggest promising opportunities for future work. There exist metrics for aesthetic quality, such as the Balinsky symmetry metric that measures the grid quality of the layout [5], which can be incorporated in the objective function to optimize for the appearance of the layout in addition to usability. Furthermore, since our system provides a platform for human-AI collaboration in UI design, it would be insightful to study the dynamics between human and AI during such a collaboration, such as disagreements that may arise between the two. It would also be interesting to compare the final layout after many iterations of this human-AI collaboration for a novice designer with the layout created by an expert designer.

CONCLUSION

Our model predicts task performance (a metric combining task completion time and error rate) of a general mobile UI that contains a variety of element types and supports many types of user interactions. We then developed an algorithm that uses this model to make iterative updates to a UI's layout using gradient descent. We used this optimization algorithm to improve layouts of two UI's, including one interface that our model has not been trained on. Our optimization algorithm was able to generate layout alternatives with better task performance for both interfaces, as confirmed by crowdsourcing studies comparing the initial and optimized layouts. This demonstrates the effectiveness of this algorithm in improving a layout's task performance, as well as its ability to generalize and make improvements to new interface layouts. The optimization's output layout do need minor human refinement in order to look aesthetically pleasing. However, collaboration between human and AI is recommended over relying solely on either.

REFERENCES

- [1] 2019. Android vs iPhone mobile owners in the U.S. 2013, by age. (2019). <https://www.statista.com/statistics/271228/android-vs-iphone-mobile-owners-age/>
- [2] Khalid Alemerien and Kenneth Magel. GUIEvaluator: A Metric-tool for Evaluating the Complexity of Graphical User Interfaces.
- [3] Leena Arhippainen and Marika Tähti. 2003. Empirical evaluation of user experience in two adaptive mobile application prototypes. In *MUM 2003. Proceedings of the 2nd International Conference on Mobile and Ubiquitous Multimedia*. Linköping University Electronic Press, 27–34.
- [4] Gilles Bailly, Antti Oulasvirta, Duncan P Brumby, and Andrew Howes. 2014. Model of visual search and selection time in linear menus. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 3865–3874.
- [5] Helen Balinsky. 2006. Evaluating interface aesthetics: measure of symmetry. In *Digital Publishing*, Vol. 6076. International Society for Optics and Photonics, 607608.
- [6] Helen Y Balinsky, Anthony J Wiley, and Matthew C Roberts. 2009. Aesthetic measure of alignment and regularity. In *Proceedings of the 9th ACM symposium on Document engineering*. ACM, 56–65.
- [7] Rachel Eardley, Anne Roudaut, Steve Gill, and Stephen J Thompson. 2017. Understanding Grip Shifts: How Form Factors Impact Hand Movements on Mobile Phones. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 4680–4691.
- [8] 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.
- [9] Krzysztof Gajos and Daniel S Weld. 2004. SUPPLE: automatically generating user interfaces. In *Proceedings of the 9th international conference on Intelligent user interfaces*. ACM, 93–100.
- [10] Todd M Gureckis, Jay Martin, John McDonnell, Alexander S Rich, Doug Markant, Anna Coenen, David Halpern, Jessica B Hamrick, and Patricia Chan. 2016. psiTurk: An open-source framework for conducting replicable behavioral experiments online. *Behavior research methods* 48, 3 (2016), 829–842.
- [11] Curtis Hardyck and Lewis F Petrinovich. 1977. Left-handedness. *Psychological bulletin* 84, 3 (1977), 385.
- [12] Apple Inc. 2019. Adaptivity and Layout. (2019). <https://developer.apple.com/design/human-interface-guidelines/ios/visual-design/adaptivity-and-layout/>
- [13] Mohammad M Khajah, Brett D Roads, Robert V Lindsey, Yun-En Liu, and Michael C Mozer. 2016. Designing engaging games using bayesian optimization. In *Proceedings of the 2016 chi conference on human factors in computing systems*. ACM, 5571–5582.
- [14] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [15] Christophe Leys, Christophe Ley, Olivier Klein, Philippe Bernard, and Laurent Licata. 2013. Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. *Journal of Experimental Social Psychology* 49, 4 (2013), 764–766.
- [16] 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*. ACM, 29.
- [17] J Derek Lomas, Jodi Forlizzi, Nikhil Poonwala, Nirmal Patel, Sharan Shodhan, Kishan Patel, Ken Koedinger, and Emma Brunskill. 2016. Interface design optimization as a multi-armed bandit problem. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 4142–4153.
- [18] Risto Näsänen, Helena Ojanpää, and Ilpo Kojo. 2001. Effect of stimulus contrast on performance and eye movements in visual search. *Vision Research* 41, 14 (2001), 1817–1824.
- [19] Jakob Nielsen. 1995. 10 usability heuristics for user interface design. *Nielsen Norman Group* 1, 1 (1995).
- [20] Jakob Nielsen. 2001. Usability Metrics. (Jan 2001). <https://www.nngroup.com/articles/usability-metrics/>
- [21] Adam Paszke, Sam Gross, Soumith Chintala, and Gregory Chanan. 2017. Pytorch: Tensors and dynamic neural networks in python with strong gpu acceleration. *PyTorch: Tensors and dynamic neural networks in Python with strong GPU acceleration* 6 (2017).
- [22] Ken Pfeuffer and Yang Li. 2018. Analysis and Modeling of Grid Performance on Touchscreen Mobile Devices. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 288.
- [23] Sergey Polyuk. 2019. A Guide to Interface Design for Older Adults. (Jun 2019). <https://www.toptal.com/designers/ui/ui-design-for-older-adults>
- [24] Juan C Quiroz, Sushil J Louis, Anil Shankar, and Sergiu M Dascalu. 2007. Interactive genetic algorithms for user interface design. In *2007 IEEE congress on evolutionary computation*. IEEE, 1366–1373.
- [25] Andrew L Sears. 1995. AIDE: A step toward metric-based interface development tools. In *Proceedings of the 1995 8th Annual Symposium on User Interface Software and Technology, UIST'95*. ACM, 101–110.

- [26] Amanda Swearngin and Yang Li. 2019. Modeling Mobile Interface Tappability Using Crowdsourcing and Deep Learning. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 75.
- [27] 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*. ACM, 543–555.