

Improving User Interface Generation Models from Designer Feedback

Jason Wu, Amanda Swearngin, Arun Krishna Vajjala, Alan Leung, Jeffrey Nichols, Titus Barik
Apple

Seattle, USA

jason_wu8,aswearngin,a_krishnavajjala,alleu,jwnichols,tbarik@apple.com

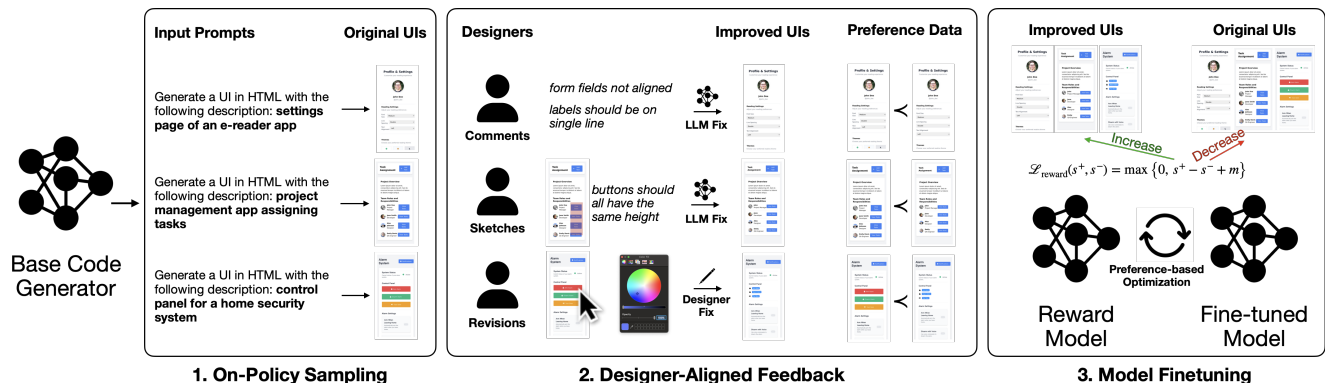


Figure 1: An overview of our approach for improving UI generation models from designer feedback. First, we conduct on-policy sampling by generating UI samples using a code LLM. Next, designers employ familiar workflows such as commenting, sketching, and direct revisions to fix design flaws and make improvements. Using the resulting preference dataset, we fine-tune the code LLM to increase the relative probability of generating code that resembles the improved UIs.

Abstract

Despite being trained on vast amounts of data, most LLMs are unable to reliably generate well-designed UIs. Designer feedback is essential to improving performance on UI generation; however, we find that existing RLHF methods based on ratings or rankings are not well-aligned with designers’ workflows and ignore the rich rationale used to critique and improve UI designs. In this paper, we investigate several approaches for designers to give feedback to UI generation models, using familiar interactions such as commenting, sketching and direct manipulation. We first perform a study with 21 designers where they gave feedback using these interactions, which resulted in ~1500 design annotations. We then use this data to finetune a series of LLMs to generate higher quality UIs. Finally, we evaluate these models with human judges, and we find that our designer-aligned approaches outperform models trained with traditional ranking feedback and all tested baselines, including GPT-5.

CCS Concepts

• **Computing methodologies** → *Learning from demonstrations; Learning from implicit feedback; Learning from critiques*; • **Human-centered computing** → *Graphical user interfaces*.

Keywords

UI modeling, reinforcement learning from human feedback, UI generation, UI assessment

ACM Reference Format:

Jason Wu, Amanda Swearngin, Arun Krishna Vajjala, Alan Leung, Jeffrey Nichols, Titus Barik. 2018. Improving User Interface Generation Models from Designer Feedback. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym ’XX)*. ACM, New York, NY, USA, 16 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Despite being trained on vast amounts of data, today’s large language models (LLMs) are unable to reliably generate well-designed user interfaces (UIs) [47]. This suggests that current datasets do not capture what constitutes “good” UI design, much of which resides in tacit domain knowledge [38]. Although general heuristics exist [27, 37], applying them effectively requires experience to understand and navigate subtle trade-offs. As both prospective users and domain experts, designers are well positioned to inform the training of UI generation models. How can models effectively learn from designers’ expertise?

Machine-learning approaches, such as reinforcement learning from human feedback (RLHF) [30], have been developed to steer

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 the author(s) 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.

Conference acronym ’XX, Woodstock, NY

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/2018/06

<https://doi.org/XXXXXXX.XXXXXXX>

models toward human-preferred responses; however, collecting data required to represent design knowledge is difficult. For example, prior work applied rubric-guided ratings and rankings to collect designer feedback [9, 25] but found that subjectivity within the guidelines and differing preferences among designers led to a noisy learning signal for model training [9, 46]. Moreover, these interfaces elicit judgments about existing UIs rather than the changes designers would make, such as concrete edits to layout, colors, and typography.

To understand which types of data are most effective for training models, we investigate designers' common workflows. For example, previous work [15] has shown that designers are already effective at communicating design knowledge as part of their day-to-day job through activities such as design reviews, white-boarding, and using design software (e.g., Figma, Sketch). These activities produce artifacts such as i) natural language comments, ii) visually-grounded annotations, and iii) low-level design revisions (e.g., before and after a designer's edits) that might be useful for model training. To this end, we develop interfaces for model training based on these activities and introduce techniques that transform them into machine-learnable preference data for UI generation models.

To validate our approach, we performed a designer feedback study with 21 professional designers at a large technology company, which led to a dataset of ~1500 training examples. To measure the quality of this generated data, we conducted a study where we found that data collected from designer-aligned interfaces led to higher agreement rates than data collected from conventional ranking interfaces. To demonstrate the utility of our dataset, we fine-tuned several existing code generation models. First, we found that supervision derived from designers' sketch and revision feedback yielded significant gains over both untuned baselines and models tuned on conventional ranking feedback. Second, applying designer supervision to a strong open-source model enables it to outperform all tested baselines, including a larger, state-of-the-art, proprietary reasoning model, GPT-5.

To summarize, this paper makes the following contributions:

- (1) Techniques for transforming designer comments, sketches, and revisions into machine-learnable preference pairs.
- (2) A dataset of ~1500 UI screens associated with design feedback collected from twenty one designers. An analysis of this data shows UI preference pairs generated from natural designer feedback workflows have lower levels of disagreement than conventional ranking feedback.
- (3) A validation of our approach that shows how designer feedback improves open-source UI generation models. An arena-style evaluation showed that some forms of designer feedback, such as sketches and revisions, were highly effective in improving open-source models. Our best-performing model outperformed all tested baselines, including a larger proprietary model, GPT-5.

2 Related Work

To contextualize our work, we review literature in *i*) UI generation, *ii*) annotation interfaces for collecting human feedback, and *iii*) and tools to support designer workflows.

2.1 UI Generation

There have been numerous approaches developed to dynamically generate UI. Model-based user interface development (MBUI) is a development approach for user interface applications, where developers first describe desired UIs in abstract specifications [31, 40], which an MBUI environment then uses to generate concrete code implementations. Mobi-D [32] is an example of a MBUI system that separated UI development into different models corresponding to application data, user interactions, and presentation strategy. The Personal Universal Controller is another example that was developed to generate personalized interfaces for different appliances, using a unified remote control device [26].

Artificial intelligence (AI) approaches have also been used to make dynamic, data-driven decisions on how this model to implementation should occur. SUPPLE is a toolkit developed to dynamically personalize UIs based on a set of constraints and objective functions that represent device affordances [11, 12], user preferences [13], and user ability [14]. More recently, machine learning approaches, especially those using data driven neural networks, have been applied to code generation. These large models [1, 24, 34] are promising in that they are trained on large amounts of data, which allows them to learn the distribution of text and code. Several works have finetuned LLMs for layout and UI-related tasks [35, 39, 47, 51], often using custom datasets. UICoder is an LLM that generates SwiftUI code using automated tools such as code compilers and vision-language models as training signals [47]. While this training approach was effective at increasing the syntactic understanding of LLMs, improvement was more difficult for design-related aspects, and the authors showed that models still made numerous design-related errors. Overall, data-driven approaches to UI generation would benefit from greater volumes of human feedback, specifically from domain experts such as designers.

2.2 Interfaces for Collecting Human Feedback

One way that machine learning models have improved over time is by using human expertise to provide labeled data for training. Both domain experts and non-experts have used annotation interfaces to provide labels for model training. A body of work studies the design of annotation interfaces for humans to provide their knowledge to models. For example, the Teachable Machine is a system where people interactively provide a small number of demonstrations to train a machine learning model [6].

Researchers have designed some of these interfaces to maximize the speed in which people can provide labels and enable them to collaborate to solve tasks [23]. Other work has designed interfaces to make annotation more enjoyable, such as Games with a Purpose (GWAP), which let humans collaborate to solve real-world, computationally difficult tasks [42]. In contrast to our goal for this paper, most of these works are not targeted towards enabling experts to provide feedback through their existing workflows.

More recently, research has introduced interfaces to allow people to provide feedback to LLMs, which are unique in their breadth of supported tasks and expressive output space. Some of these interfaces allow people to rank LLMs output which is later used for training and evaluation. The LMSYS arena is an interface for crowdsourcing LLM evaluation by having people provide pairwise

evaluations [8]. The LMSYS arena has several variations for people to provide feedback for specific applications like design, code generation, and multi-modal conversations. Beyond arena-style interfaces, other work has investigated collecting richer types of human feedback beyond binary signals such as natural language explanation, demonstrations [36], and sequential design revisions [48]. Natural language feedback is generalizable and easy to author, but it may not be the most efficient or informative for some domains with visual components like UI design. Thus, we investigate alternative interfaces for collecting designers’ feedback on UI designs going beyond natural language annotations to visually grounded annotations and revisions. We also investigate methods that would enable designers to provide UI design feedback through their everyday activities in contrast to work that targets collecting annotations and feedback through a separate tool.

2.3 Tools for Designers

Since our work focuses on types of natural feedback that can be captured from designers’ workflows, we review computational tools built to support designers’ workflows. In the book, *Sketching User Experiences*, Buxton highlights the need for tools that facilitate iterative UI design and development [5]. Several lines of research have focused on building tools that process low-fidelity designer artifacts, such as sketches and mockups, and transform them into higher-fidelity UIs with code implementations [3, 17, 22]. For example, SILK was a system that translated hand-drawn sketches of UIs into interactive, testable prototypes using gesture-recognition software [21]. Swire uses a neural network encoder to retrieve relevant high-fidelity examples from a UI database for design inspiration [17]. This body of work suggests that artifacts already produced by designers (e.g., sketches) can encode information valuable for producing high-quality UIs.

Other research investigates how designers critique and revise UIs and provides tools to support this process. d.note helped designers revise UIs through change tracking, annotations, and supporting design revisions [15]. Compared to traditional workflows involving sketching on static images, d.note allowed designers to more efficiently implement suggested changes with fewer clarifications. Charrette supported designers in giving feedback during design reviews and meetings [28]. Charrette classified the types of annotations that occur on *artboards*, digital canvases used by designers to create, edit, and present designs, and the authors developed a web-based application that facilitates iteration and discussion of UI designs. Insights from these tools suggest that in many cases, designer workflows can be computationally supported and encoded.

3 Background

Our paper focuses on a formulation of UI generation where an LLM model is provided with a textual prompt (e.g., a natural language description), and generates a UI represented as code (e.g., HTML). Since we aim to improve their UI design capabilities, we first provide technical background on *i*) the data needed to train LLMs and *ii*) opportunities for designer input to improve this data.

3.1 Data Format

LLMs are typically trained in multiple stages and can incorporate different types of data in each [30]: *i*) unsupervised pretraining, *ii*) supervised finetuning, and *iii*) model alignment.

Most of an LLM’s training time is spent during the *pretraining* stage, where the model uses an unsupervised objective to learn the distribution of text and other data from large, unstructured data, such as web dumps. During the *supervised finetuning* stage, the model is trained to replicate a human-authored output given a distribution of input prompts, which allows the model to follow instructions rather than simply predicting continuations. Finally, the *model alignment* stage, focuses on fine-tuning models so that they more closely match human preference data. Instead of training on a single “ground truth” response, as in the previous stage, alignment training relies on numerical ratings or rankings of output candidates, which is more useful for subjective tasks.

Of the two stages of LLM training where it is possible to incorporate human input, we choose to focus on the *model alignment* stage, since designers tend not to directly write code implementations needed for constructing input/output examples and there might be numerous possible “ground truth” responses for a single prompt. Therefore, **our paper focuses on collecting paired comparison data from designers to model their design preferences.**

$$\mathcal{D}_{\text{pref}} = \{(x, y^+, y^-) : y^+ \succ y^-\} \quad (1)$$

Formally, the dataset of design preferences $\mathcal{D}_{\text{pref}}$ consists of triplets of an input natural language UI description x and two possible UI designs y^+ and y^- , where y^+ (“chosen” sample) is rated by a designer to be preferred over y^- (“rejected” sample).

3.2 Data Quality

Previous research has shown that, in practice, the quality of data has a significant impact on model performance [20, 49, 52]. It may seem straightforward to generate paired comparison data entirely synthetically. For example, a simple strategy would be to use a large, strong model (e.g., GPT-5) to generate chosen responses and a small, weak model (e.g., GPT-1) to generate rejected responses. However, these “trivial” pairs may be less effective, since during training, they cannot help the model assess and choose between outputs with more subtle differences.

Due to designers’ valuable expertise, there have been several attempts to include them in model training. For example, previous work [10] collected natural language critiques from designers who were asked to follow a pre-defined rubric. However, the authors found that they disagreed with a significant portion of critiques generated by models trained on this data and even other designers. Other work [46] employed a similar approach where designers were asked to follow a high-level rubric and provided binary comparison ratings of sampled UI pairs. However, this also led to low inter-rater reliability and significant amounts of label noise, which negatively impacted model performance. We hypothesize that a key reason is a mismatch between the subjective nature of UI design and the rigidity of rubric-based rating tasks. When evaluating synthetic outputs for model alignment, designers often must choose between two flawed options, each with different strengths and weaknesses. This process reduces their nuanced expertise to coarse labels and

offers little opportunity to propose concrete improvements. Instead, in this paper, **we show that collecting data from designers models around their existing workflows leads to higher quality data and better model performance.**

4 Collecting Designer Feedback

In this section, we describe the construction of a large dataset of designer-annotated UIs that we collected using designer-aligned feedback strategies. We first generate an initial dataset of UI code and screenshots using a code LLM. Based on common workflows identified in the literature, we developed four interfaces that allow designers to critique and revise the code LLM's output. Finally, we use these interfaces in a data collection with twenty one designers at a large technology company to collect annotations of our data.

4.1 Initial Data Generation

We first generated a large UI dataset to provide materials for designers to annotate using on-policy sampling of a base model. At a high-level, our approach consisted of *i*) generating a list of textual descriptions of UIs, *ii*) using a code LLM to generate UI programs for the textual descriptions, and *iii*) rendering the UI programs into screenshots.

Instead of using existing datasets of UI descriptions [44], we synthetically generated a list of diverse and detailed descriptions to seed our UI generation. We prompted an LLM to generate a large list of approximately 100,000 natural language descriptions. To generate this list, we first prompted the LLM with a set of manually-authored example descriptions that described UI functionality, layout, and content, and then asked the LLM to generate 10 more unique descriptions. We evaluated this prompt using a high temperature value, which increased the probability of the LLM generating a new set of descriptions. We merged descriptions from each evaluation into a global set until the number of descriptions reached our target.

To generate UI programs corresponding to these prompts, we used a publicly released code LLM, Qwen2.5-Coder 32B. At the time we started this project, Qwen2.5-Coder was the strongest publicly available code LLM that could fit on a single GPU. We sampled a random description from the list of UI descriptions and prompted the code LLM to generate a well-designed web page using a limited set of web libraries (HTML, Tailwind, and Font Awesome). We used each randomly-selected description to generate 32 different possible web pages via temperature-based sampling of the LLM.

We rendered each UI program into a screenshot using an automated pipeline. First, a HTML parser extracts all referenced images in the code (i.e., `` tags). To generate plausible placeholder assets for rendering images, we fed each image tag's `alt` attribute into an off-the-shelf text-to-image model called Flux Schnell. The HuggingFace ID for this model is `black-forest-labs/FLUX.1-schnell`. Our code generation prompt contained explicit instructions to include `alt-text` for all images; however, if the `alt` tag was still missing, we used the `src` attribute value instead. We developed an automated script that *i*) staged each UI's HTML code and required assets (e.g., library files and image assets) into a server then *ii*) used a headless browser to visit the hosted URL and take a screenshot. In addition to the browser-based screenshot rendering, we developed a program based on the open-source `html2sketch` library to convert each

web page into an editable Sketch file using the computed positions and styles from the browser DOM. To further improve this dataset, we used the UIClip [46] base model to compute the quality scores of UIs. For the 32 original UIs generated for each description, we kept the top 8 outputs according to their computed quality score. This allowed many "obviously bad" outputs to be filtered out automatically, which allows designers to give more subtle feedback on plausible screens.

To summarize, we generated a dataset of ~6400 UIs corresponding to ~200 unique textual descriptions. Each UI contains a natural language description, HTML source code, a screenshot rendered by a browser engine, and a Sketch file.

4.2 Data from Designer Feedback

To collect feedback from designers on the generated UIs, we designed four annotation interfaces inspired by the design principles identified by Hartmann et al. [15] (Figure 2). Hartmann et al. conducted a broad literature review of how designers revise artifacts in textual documents, source code, movies, games, and user interfaces. Based on their review of these domains, they proposed four principles important for UI revision, of which three correspond to actions by designers themselves: *i*) commenting, *ii*) sketching, and *iii*) revising. We developed annotation interfaces for *i-iii*, and included a baseline ranking interface inspired by current practices for incorporating user feedback into LLMs through pairwise comparisons.

4.2.1 Ranking. We designed the baseline ranking interface to mimic strategies used by LLM chat interfaces to elicit user feedback, e.g., asking users to select the better of two generated responses. Some other common interactions used to collect rankings include "thumbs up" or "thumbs down" controls, and implicitly ranking responses through user-initiated response regenerations (e.g., user is not satisfied with the first response). Our ranking interface displays a textual description and two candidate UIs arranged side-by-side. The interface (Figure 2 Far Left) asks designers to select the UI that they feel is better-designed. If both UIs are poorly designed (e.g., both contain many design flaws), the interface asks designers to choose the UI that they feel is a better starting point for fixing.

4.2.2 Commenting. Multiple studies suggest that designers often write high-level comments as critiques, propose changes, or "todo items." [15, 28] While designers use a wide variety of tools (e.g., text editors, note-taking software), for the purposes of our experiment, we developed a commenting interface that allows designers to provide natural language feedback on UIs. The interface (Figure 2 Center Left) presents designers with a UI screenshot and its textual description, and asks them to write a list of natural language critiques. Designers type each critique or suggestion into a text-field then hit the Enter key to add it to a list.

4.2.3 Sketching. The sketching interface (Figure 2 Center Right) is similar to the commenting interface; however, the sketching interface allows designers to provide visually grounded feedback (e.g., annotations). The interface presents designers with a UI screenshot and asks them to identify areas for improvement within a screenshot. The interface allows drawing bounding box and point annotations on top of the UI. After drawing an annotation, the

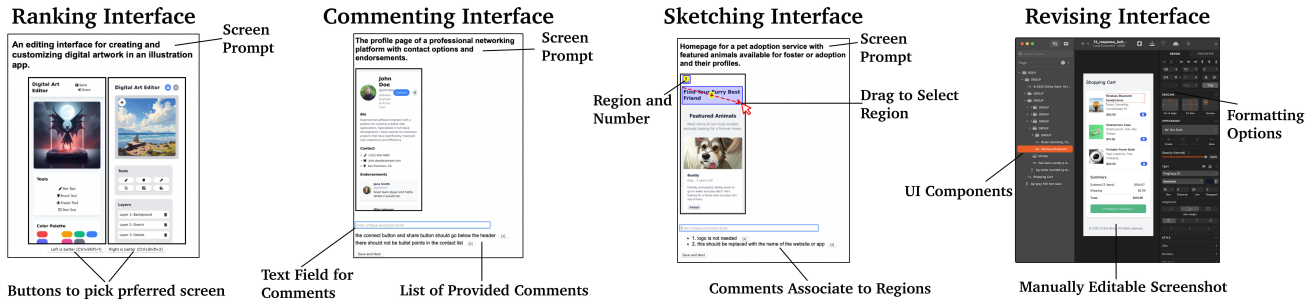


Figure 2: Figure shows the four interfaces we developed to collect feedback from designers. The ranking interface (Far Left) allows users to select the better of two UI screenshots through a binary response. The commenting interface (Center Left) allows users to write a list of natural language critiques or comments for a UI screenshot. The sketch interface (Center Right) allows users to draw annotations (boxes and points) on a UI screenshot and associate them with textual comments. Designers used the Sketch design software (Far Right) to make direct edits to model-generated UIs, which were first converted into the appropriate format. The commenting, sketching, and revising interfaces and inspired by interactions identified by Hartmann et al [15].

interface asks the designer to provide a natural language feedback for that region. Designers can draw multiple annotations on the UI screenshot which it displays as a list on the interface.

4.2.4 Revising. In addition to providing critiques or feedback, designers often use direct manipulation to revise UI designs in software such as Figma and Sketch. Previous work suggests that in many cases, designers may prefer this revision to occur in the output domain (e.g., rendered graphics) rather than the source domain (e.g., source code) [15].

We developed an interface that allows designers to revise LLM-generated UIs using Sketch (Figure 2 Far Right). The interface first displays a UI’s description and screenshot to allow designers to think about possible flaws and potential fixes. The interface asks designers to download the corresponding Sketch file, which we generated from our initial data generation process. Designers modified the Sketch file to improve the UI, and then uploaded it back to the interface.

4.3 Designer Feedback Study

To obtain examples of designer feedback, we recruited twenty one designers at our institution, a large technology company.

4.3.1 Participants. We recruited participants by posting messages in company message boards and through word-of-mouth. The recruited participants had varying levels of professional design experience, ranging from 2 to over 30 years. Participants also worked in different areas of design, such as UI/UX design, product design, and service design. Participating designers also noted the frequency of conducting design reviews (both formal and informal) in job activities: ranging from once every few months to multiple times a week.

4.3.2 Procedure. During the study, participants joined a video call where a member of the research team first gave an overview of the study and asked for their informed consent. Participants used each of the four feedback interfaces in randomized order, where they first watched a pre-recorded tutorial video (1-2 minutes) for each feedback interface then spent 10.5 minutes giving feedback

to UIs using that feedback interface. In total, the session lasted approximately one hour. As a thank-you for participating in our study, we offered participants a meal voucher.

4.3.3 Post-study Interview. Since the primary goal of this work is to develop designer-aligned interactions for providing feedback to ML models, we first validate interactions a set of designer interactions identified by pervious work [15]

In our post-study interview, we asked participants to estimate the amount of time that they spent doing similar tasks. Overall, participants estimated that they spent the most time on revision-like tasks (average of 33%), and participants estimated they spent the least time on ranking-like tasks (9%). Commenting (26%) and sketching (26%) also constituted a significant amount of participants’ design activities, validating that our chosen interaction strategies are representative of the types of tasks that designers normally do. Since designers in our study report doing other types of activities in their jobs, these percentages do not add up to 100. Most participants who were primarily UI/UX designers felt that editing and revising designs (e.g., direct manipulation in design editing software), took a significant amount of their time, and sometimes referred to these activities as “hands-on design work.”

4.3.4 Dataset Overview. In total, we collected 1460 annotations from twenty one designers, where one “annotation” refers to a single UI screenshot paired with all the feedback for that screen, e.g., multiple comments. Designers collected the most annotations using the ranking interface (1063), and designers, on average, were able to generate 4.8 rankings per minute. Designers collected the least annotations using the revision interface (64), on average requiring 3.45 minutes per revision. We expected this discrepancy in samples, since this condition required participants to attempt to fix UIs rather than just evaluate them.

The sketching interface (181) led to more annotations than the commenting interface (152). We hypothesize this is because the commenting interface required designers to type longer text (87.1 characters on average) than the sketching interface (42.2 characters on average), since the designers needed to textually describe UI

elements instead of being able to annotate the element directly on the image. Furthermore, the sketching interface generally led to more feedback (e.g., number of comments) per UI (2.7 on average) than the commenting interface (1.9 on average). Examining the annotations and comments, we found that comments from the sketching interface typically contained lower-level feedback about particular regions (e.g., “make this text larger”) while the commenting interface typically contained higher-level feedback (e.g., “the screen has poor information hierarchy”).

The annotation interface designers used strongly impacted the number of data samples that could be collected in a fixed duration. On average, designers collected over 15x more annotations using the ranking interface than the revision interface.

4.3.5 UI Preference Pair Generation. We used several strategies to convert each type of designer feedback into usable preference data for model training (i.e., UI preference pairs).

- **Ranking.** As a baseline, we use designers’ rankings to form preference pairs out of the UI screenshots.
- **Commenting.** Designers produced a list of natural language comments corresponding to a UI screenshot. For each screenshot, we prompt an LLM to improve the UI’s HTML code using the list of provided comments. We then re-rendered the code into a new screenshot, which we labeled as preferred over the original UI screenshot.
- **Sketching.** Designers produced a list of visually-grounded annotations corresponding to a UI screenshot. We adopted a similar approach to processing commenting data, however we associated each textual comment with an HTML code snippet of the DOM element with the highest overlapping IoU score with the drawn box.
- **Revising.** Designers downloaded a Sketch file generated from the code used to render a UI screenshot then modified and re-uploaded an improved Sketch file. We formed a preference pair using the rendered preview of the improved Sketch file as the preferred sample and a rendered preview of the original Sketch file as the alternative.

4.3.6 UI Preference Pair Quality Assessment. We conducted a quality assessment with six members of the research team (who are HCI experts) to estimate the quality of UI preference pairs generated from each condition of the designer feedback study. Since there is no standardized rubric or rating system to evaluate the quality of UI preference pairs, we operationalize quality as the percent agreement between the researchers’ choice of the best screen in each UI preference pair, and the “improved” UI from the UI preference pairs generated from designer feedback.

In other words, if both an HCI expert and a design professional independently agree on the best screen for a UI preference pair, that UI preference pair ranking is more likely to be accurate.

During this study, we sampled UI preference pairs from the designer-generated dataset, and stratified them by the feedback interface used to collect them. In a web interface, we showed example pairs to a researcher, who did not have knowledge of *i*) the type of

feedback used to generate the pair and *ii*) which screenshot in the pair was “preferred” by the designer. The interface told researchers to choose the UI from each pair that they felt was better designed, and we measured the rate at which the researchers’ choice of the best screen corresponded to either *i*) the UI preferred by designers during the ranking task, *ii*) the improved UI from the designer’s edits during the revision task, or *iii*) the improved UI generated by running an LLM on the designer’s comments or sketch annotations from the commenting and sketching task, respectively. Each researcher spent around 30 minutes on this task, which led to 695 ratings in total.

4.3.7 Results. Overall, researchers agreed with designers for 61.7% of the UI preference pairs. Of the four feedback conditions, researchers most often agreed with the designer on the best screen from pairs generated from the revision condition (76.1%), which suggests that allowing designers more control reduces label noise. On the other hand, researchers and designers had very low agreement on the best screen (49.2%, close to random chance) from pairs generated from the ranking condition. The low agreement for ranking-style data is consistent with prior work that compared designer pairwise rankings and also found low agreement among independent designers [46]. Compared to prior work, we expected lower agreements from our ranking experiment because the pairs used were already pre-filtered to reduce “easy” comparisons with obvious design flaws (Section 4.1).

In general, designers and researchers had higher agreement on their choice of the best screen from preference pairs generated from the comment and sketch conditions, likely because they provided opportunity for designers to give specific feedback on flaws and how to fix them. However, the quality of these UI preference pairs may have been limited by the LLM’s ability to operationalize some types of high-level designer comments, e.g., “improve the information hierarchy of this screen”. It is reasonable to expect that future, stronger LLMs will possess better code editing abilities which would in turn improve the quality of comment-derived data. Researchers and designers agreed more often on the UI preference pairs generated from the sketching condition (63.6%) than the UI preference pairs generated from the commenting condition (57.3%). We hypothesize this may be because including grounded annotations for each comment on a UI design enabled the LLM to better localize the area of the UI the comment applied to, and thus generate a better fix for it.

Researchers and designers had varied levels of agreement on UI preference pairs generated as output from designer feedback on different annotation interfaces. We use agreement as a proxy for data quality. *For our chosen model*, UI preference pairs generated from designers’ revisions led to the highest agreement while UI preference pairs generated from designers’ pairwise rankings led to the lowest agreement.

5 Learning from Designer Feedback

In this section, we described how we used our collected designer feedback to fine-tune and evaluate several code generation LLMs.

We first describe our model training technique that i) trains a reward model using preference pairs generated from designer feedback then ii) uses the reward model to finetune a code generation model. Using this approach, we trained various code generation models and evaluated them in two arena-style human evaluations [8]. We first trained four generator models using each type of designer feedback collected in our study. We found that both sketching and revision feedback led to model improvements, and sketch-based feedback performed the best overall. Finally, we used our sketch-based reward model to finetune two more publicly available code generation LLMs. In all tested cases, incorporating designer sketch data improved model performance and our best model outperformed GPT-5, a proprietary reasoning model from OpenAI.

5.1 Training Approach

To fine-tune code generation models, we adopt a preference-based alignment strategy described in Section 3. Our training set up is similar to existing RLHF architectures [30] that involves training two models: i) a reward model and ii) a generator model.

Our setup uses a two-step process: we first train the reward model using human feedback, then we use this reward model as a training signal for training the generator. Previous work [30] found that this two-step approaches provides advantages over directly using human feedback to train the generator, such as improved sample efficiency. Reward models are typically trained to assign a numerical score to an input/output pair from the generator (e.g., the concatenation of a textual prompt and response). The reward model is trained to assign higher scores to “preferred” responses marked by human labels over ones that are not. After training, the reward model is used to guide the generator model to produce outputs that maximize its score, either through traditional reinforcement learning or data generation. More details can be found in other papers [20, 30].

5.1.1 Reward Model. We chose to base our reward model off the multi-modal CLIP B/32 architecture, which we initialized from the publicly released UIClip model [46] to improve training efficiency. To isolate the effectiveness of our training data, we used the UIClip checkpoint which was not trained with any human preference pairs from their original paper. The HuggingFace ID of our initialization checkpoint is `biglab/uiclip_jitteredwebsites-2-224-paraphrased`. The reward model accepts i) a rendered image (a UI screenshot) and ii) a natural language description (a target description of the UI). These two inputs are fed into the model to produce a numerical score (reward), which is calibrated so that better-quality visual designs result in larger scores. To assign rewards to HTML code, we used the automated rendering pipeline described in Section 4.1 to first render code into screenshots using browser automation software.

To train the reward models using our designer feedback dataset, we used a margin-based variation of the original pairwise contrastive objective in UIClip, which tunes the model so that it assigns higher scores to “preferred” samples over “rejected” samples.

$$\mathcal{L}(s^+, s^-) = \max \{0, s^+ - s^- + m\} \quad (2)$$

In our loss function (Equation 2), s^+ and s^- refer to the reward model scores of the preferred and non-preferred UI screenshots,

respectively. These scores are computed by computing a scaled dot product between the encoded textual description and UI screenshots. m refers to an empirically determined margin value.

When finetuning, we freeze all but the last layer to prevent overfitting and employed an additional data augmentation technique. Because we only asked designers to give feedback on the top 8 out of 32 outputs (Section 4.1), it caused our dataset to exclude many examples of poor UIs, which are nevertheless needed to accurately represent the entire distribution of generated UIs. To resolve this imbalance, sampled additional UI pairs from the entire output distribution (i.e., from all 32 candidate outputs) and synthetically labeled them using UIClip’s score. The reward model was trained using both designer-labeled preference pairs and these UIClip-labeled pairs, which were sampled with a fixed probability.

All reward models using the same hyperparameters and for a fixed number of optimization steps. We determined the hyperparameters for the reward model training by manual inspection and experimentation, and we provide them in the appendix of this paper.

5.1.2 Generator Model. We finetuned generator models to generate outputs that maximized the expected reward using an optimization algorithm called ORPO [16]. The ORPO algorithm expects a dataset of triplets consisting of an input prompt (textual description) and two candidate outputs (HTML programs), where one of them is preferred over the other. Note that this is similar to the format of our design preference dataset (Equation 1); however, ORPO’s training data requires the output candidates to be code (e.g., HTML), not the UI images collected from designers. To produce the required data, we first use the generator model to synthesize a large batch of UI programs (HTML) from a list of descriptions, sampling multiple candidates (32) per description. In total, we generated ~400,000 HTML UIs from ~11,600 randomly sampled unique textual descriptions. Following this code generation, these HTML programs are assigned numerical scores by first rendering them into screenshots then running the reward model over them. With this dataset of scored HTML programs, we followed a procedure from previous work [41] to generate preference pairs by using the top response for an input description as the “chosen” response and a randomly selected HTML program from the same description as the “rejected” response. Finally, we ran the ORPO algorithm [16] for one epoch to learn on preference pairs labeled by the reward model. When finetuning, we employed parameter offloading [33] and mixed-precision training to improve efficiency.

All hyperparameters for the generator model training were determined by manual inspection and experimentation, and we provide them in the appendix of this paper.

5.2 Feedback Comparison Study

We conducted a study to empirically determine the most effective source of designer feedback for model training. While our data quality assessment showed that data collected from designers’ revisions of UIs led to the highest agreement rate, other forms of feedback (e.g., sketching), could be collected significantly faster, and the additional data volume could potentially offset the slightly higher noise when training models. Using our training approach, we trained several versions of code generation models from different types of designer feedback and two baseline conditions.

In total, we evaluated 6 conditions:

- Qwen2.5-Coder - The 32B variant of the Qwen2.5-Coder model, with the Ollama ID `qwen2.5-coder:32b-instruct-fp16`. This was the base generation model that generated UIs that were shown to designers. This baseline represents model performance without any additional UI-specific fine-tuning.
- Qwen2.5-Coder + UIClip - Qwen2.5-Coder trained using the base model of UIClip [46] as a reward model. This baseline has undergone additional UI-specific fine-tuning with an off-the-shelf reward model trained without any designer feedback.
- Qwen2.5-Coder + Ranking - Qwen2.5-Coder fine-tuned using a reward model derived from ranking data.
- Qwen2.5-Coder + Comment - Qwen2.5-Coder fine-tuned using a reward model derived from comment data.
- Qwen2.5-Coder + Sketch - Qwen2.5-Coder fine-tuned using a reward model derived from sketch data.
- Qwen2.5-Coder + Revision - Qwen2.5-Coder fine-tuned using a reward model derived from revision data.

Samples from these models can be found in the Appendix, Section C.

5.2.1 Procedure. To evaluate the performance of models trained under different conditions (e.g., types of designer feedback), we employed an arena-style evaluation with human judges, which has been used for general [8] and UI-specific [45, 47] LLM model evaluation. An arena-style evaluation compares the relative performance of several models by repeatedly sampling a pair of the pool of models, using both models to generate output for the same input prompt, then asking a human judge to select a “winner” from that pair. Following existing practices, we computed rating scores from judges’ responses. These scores, also referred to as Elo ratings, are numerical estimations of each model’s performance against other models in the same arena [8]. We use the same approach and parameters set by LMSYS Chatbot [8], using their publicly available code [53]. This approach scales Elo ratings to a standardized range, centered roughly at 1000, and calculates confidence intervals using bootstrap sampling.

We generated 210 descriptions for the evaluation, to align the size of our test set with those used by previous work [46, 47] and other coding-related benchmarks [7]. The list of descriptions were generated by prompting an LLM with one-sentence app screen descriptions (methodology described in Section 4.1) and are available in supplemental material. We examined the list of descriptions used for evaluation to remove duplicates and ensure they do not appear in any of the training data. Although we did not detect any exact matches in the remaining data, it is possible that some descriptions in the two splits are semantically similar, e.g., “a login screen for a banking app” and “a sign-in page for a financial planner.” We consider these types of similarities acceptable, since both login screens and finance related applications are common types of UIs.

We ran all models in the arena using their default sampling parameters. We converted models with downloadable weights to GGUF format, quantized them to 16-bit precision, and loaded them into Ollama¹, an open-source utility for managing and running

LLMs. For efficiency, we limited the maximum output length for all models to 4096 tokens, which covers the portion of code responsible for rendering the initial screen viewport seen and rated by annotators. We generated screenshots from the model code output using the pipeline described in Section 4.1.

5.2.2 Participants. To evaluate models, six members of our research team (who are HCI experts) participated in the model evaluation as human judges. We showed judges a randomly selected textual description and the generated UI screenshot output of two randomly selected models, and we instructed the judges to select the UI that they felt had a better design. Note that the evaluation interface shown to judges does not show which samples were generated by which models (i.e., a “blind” rating task), reducing the risk of bias towards any specific condition.

In total, judges conducted 405 pairwise comparisons between the UI output screenshots of two models. Since we tested 8 different models, this led to approximately 27 comparisons for each of the $\binom{6}{2} = 15$ possible model comparisons.

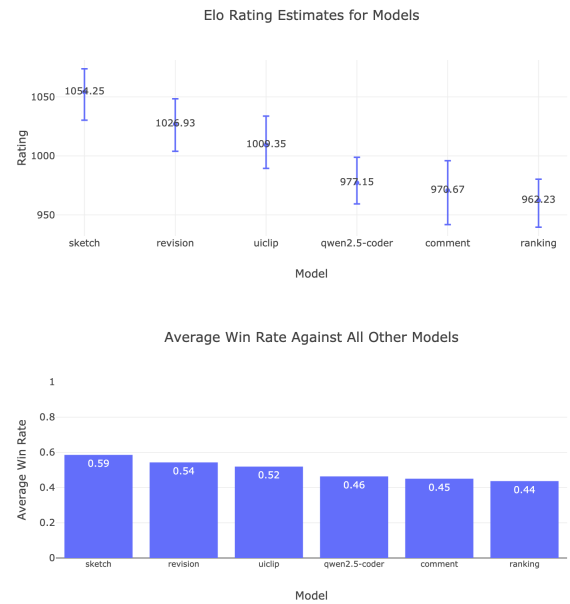


Figure 3: Rating scores (Top) and average win rate (Bottom) of models in our feedback comparison study. We computed the rating scores by using the LMSYS calculation methodology [8, 53], and higher scores indicate models that were more often preferred by human judges. Bars show the median score and 95% confidence intervals generated using bootstrap sampling.

5.2.3 Results. Figure 3 shows the results of our evaluation, which shows the rating scores and confidence intervals for each model using the LMSYS calculation methodology [8, 53].

Overall, our results show that our fine-tuning approach can improve a code generation model’s UI generation ability. First, even training with an off-the-shelf reward model trained without designer feedback [46] (i.e., UIClip condition) led to improvements

¹<https://ollama.com/>

over the un-tuned Qwen2.5-Coder model. Because some pairwise comparisons were decided on obvious design flaws (e.g., overlapping text), UIClip’s training on synthetically introduced artifacts helped the generator avoid these errors. We hypothesized that incorporating designer feedback would enable the model to better capture subtle design choices and assess trade-offs. While some types of designer feedback were effective at further improving performance, others were not.

Among the tested designer-trained conditions, the sketch and revision models performed better than the original base model, and the sketch-trained model performed the best overall. These scores roughly coincide with our result from our previous data quality assessment (Section 4.3.6) which showed these types of feedback resulted in the highest agreement rates between researchers and designers (63.6% and 76.1% , respectively). We hypothesize that the sketch-trained model’s higher score over the revision-trained model stems from the roughly three times larger sketch dataset, which offset the revision data’s quality advantage. In contrast, training on commenting or ranking data led to no change or even slight degradation in performance (Figure 3). We hypothesize that the noisier labels and lower agreement from these feedback conditions led to a weak or detrimental signal for model training.

The performance of UI generation models depends on both the quality and quantity of data used to train them. While some interactions require designers to spend more time and effort, they can also result in better data quality. *For our chosen base model*, we found that data collected from sketching and revision feedback led to model improvements, and training on sketching feedback led to the best model performance.

5.3 Model Generalization Study

Our feedback comparison study identified effective forms of designer feedback for model training; however, the prior experiment fine-tuned only the model that produced the designer-annotated data. Because different base models may generate UI code with distinct distributions of design flaws (in both type and frequency), we now assess whether feedback derived from Qwen2.5-Coder 32B generalizes beyond that source model. Therefore, in this section we i) evaluate generalizability by fine-tuning two additional open-source LLMs and ii) compare their performance against their untuned baselines and proprietary LLMs (e.g., GPT-5). We selected Qwen2.5-Coder 3B [18] and Qwen3-Coder [50] to test generalization, since they were the best-performing permissively-licensed coding models capable of running locally on an edge device and consumer GPU at the time of this experiment, respectively. From our visual inspection, their outputs qualitatively differ significantly from Qwen2.5-Coder 32B; however, because both are from the Qwen family, we acknowledge that this limitation potentially overstates our method’s generalizability.

To train these base code generation models, we chose to use the best-performing sketch-trained reward model, which led to the best performance in our previous experiment. It is possible that other

forms of feedback may also be effective or exhibit different generalization properties, but we exclude them to prevent a combinatorial explosion of model variants.

In total, we evaluated the following 6 conditions:

- Qwen2.5-Coder + Sketch - The Qwen2.5-Coder 32B model fine-tuned using a reward model derived from sketch data. This model performed the best in the previous feedback comparison arena, and we include it again as a point of reference.
- Qwen2.5-Coder 3B - The 3B variant of the Qwen2.5-Coder model, with the Ollama ID `qwen2.5-coder:3b-instruct-fp16`.
- Qwen2.5-Coder 3B + Sketch - The Qwen2.5-Coder 3B model fine-tuned using a reward model derived from sketch data.
- Qwen3-Coder - The 30B variant of the Qwen3-Coder model, with the Ollama ID `qwen3-coder:30b-a3b-fp16`.
- Qwen3-Coder + Sketch - The Qwen3-Coder model fine-tuned using a reward model derived from sketch data.
- GPT-5 - A multi-modal, reasoning-capable foundation model released by OpenAI [29]. We set reasoning effort to “medium,” which is the default value in the API.

Samples from these models can be found in the Appendix, Section C. For this experiment, we used the same evaluation procedure and participants as the previous feedback comparison study (Section 5.2).




Figure 4: Rating scores (Top) and average win rate (Bottom) of models in our model generalization study. We computed the rating scores by using the LMSYS calculation methodology [8, 53], and higher scores indicate models that were more often preferred by human judges. Bars show the median score and 95% confidence intervals generated using bootstrap sampling.

5.3.1 Results. Figure 4 shows the results of our evaluation. Among the tested models, Qwen3-Coder finetuned with our sketch reward model performed the best, while the base Qwen2.5-Coder 3B model performed the worst. The additional base models included in this experiment performed roughly as we expected, where Qwen2.5-Coder 3B performed worse than its 32B counterpart and the newer Qwen3-Coder performed better. Although the base Qwen3-Coder has roughly the same number of parameters (30B), it significantly outperformed the best-performing model from the feedback study, Qwen2.5-Coder + Sketch. The Qwen3 technical report documents a wide range of improvements to model architecture, training data, and training techniques [50]. For example, one improvement is that Qwen3 models were trained on roughly double (36 trillion tokens [50]) the amount of overall data as Qwen2.5 models (18 trillion tokens [18]). This observation suggests that strategies for improving the general performance of LLMs can also lead to improvements in specific domains, like UI generation.

Nevertheless, we showed that our reward model trained with only 181 feedback samples consistently improved all tested base models. Notably, the best-performing Qwen3-Coder + Sketch model generated UI designs that were preferred over those from GPT-5, a proprietary LLM which has been estimated to be several orders of magnitude larger [2]. This suggests that design expertise and specialized fine-tuning can greatly improve the efficiency of learning UI generation.

Our results show that fine-tuning with our sketch-based reward model consistently led to improvements in UI generation capabilities for all tested baselines, suggesting generalizability. We also show that a small amount of high-quality expert feedback can efficiently enable smaller models to outperform larger proprietary LLMs in UI generation.

6 Discussion

The results of our experiments show that UI code generation models can be improved with input from expert designers. However, the method used to elicit designer feedback and convert it into machine-learnable data has a significant impact on the resulting model performance. In this section, we discuss the implications of our work for data collection methodologies and learning from designer feedback. We conclude with limitations and avenues for future work.

6.1 Agreeing to Disagree

One major challenge of our work and other human-centered problems is handling subjectivity and multiple resolutions of design problems. Both phenomena can also lead to high variance in responses, which poses challenges for widely-used ranking feedback mechanisms.

This is especially true for UI design and design in general. Our experimental results validate this in multiple ways. First, in our UI preference pair quality assessment (Section 4.3.6), six researchers ranked UI preference pairs and had very low levels of agreement on this ranking with expert designers, similar to findings by prior work [46]. While the speed of the ranking condition resulted in

the largest volume of training data (1063 samples), this was not enough to offset the low data quality. In our model evaluation study, fine-tuning an LLM on data generated from the ranking condition degraded the model's performance, resulting in the worst outcome among all evaluated models.

One approach to addressing this problem is to develop more stringent rubrics to encourage more repeatable evaluation of UI screens [4]. However, our work advocates for a different approach: instead of presenting designers with a decision where disagreement can result in “noise,” (i.e., ranking UI screens) our approach can allow different designers to each give non-conflicting alternatives (i.e., UI design revisions) that are useful for representing for design assessment. Suppose there is a UI screen where an important piece of information is not highlighted, e.g., a checkout screen whose information hierarchy makes it difficult to locate the total shopping cost. Multiple designers might be able to diagnose the same problem but propose different fixes, e.g., highlighting the important element with color or moving it to a more isolated area to draw attention. Whereas a ranking task would force the designer to choose between two model-generated UI screens, where the alternative possibly does not address the problem (or potentially even introduces new ones), our approach is able to introduce a new data point that specifically targets the original UI screen's flaws. By “agreeing to disagree,” we train the model to consider multiple solutions for a problem, instead of a single canonical answer. Anecdotaly, informal conversations with designers revealed that they often generate and evaluate multiple alternative UI designs so incorporating multiple UI design revisions into model training aligns better with designer practice.

Other research in the machine learning literature has also explored feedback mechanisms outside of standard ranking, for example through online imitation learning of human demonstrations [36] or model alignment based on editing model-generated output [19]. These approaches are motivated by the observation that desirable generations often lie outside a base model's output distribution [36], while conventional pairwise ranking limits feedback to in-distribution outputs. This observation is especially relevant for UI generation, where excellent designs are often creative and uncommon, making them unlikely to emerge from a model's typical outputs.

6.2 Tradeoffs in Collecting Designer Feedback

To capture feedback from expert designers, our paper explored tradeoffs in collecting and learning from designer feedback.

Model-centric vs Designer-centric Interfaces. An early but important step is determining the requirements of feedback interfaces. RLHF finetuning approaches (Section 3) require datasets of ranked preference pairs to calibrate the probability of “good” and “bad” outputs, which motivates most conventional ranking interfaces. This *model-centric* approach to designing feedback interfaces primarily considers the requirements of machine learning algorithms; however, our work and previous work [46] found that these feedback interfaces are often ineffective at capturing designer preference and critiques. An alternative approach that we advance in this paper is to build interfaces around the existing practices and workflows of designers [15, 28], which typically focus

on improving a single UI screen in a non-comparative setting. This *designer-centric* approach for data collection involves converting artifacts from designers' day-to-day activities, such as design reviews, white-boarding, and direct edits, into machine-learnable data. Our UI preference pair quality and model evaluation experiments indicate that this approach is effective in some contexts (e.g., sketching and revision) but not in others (e.g., commenting), suggesting the need to better understand and balance the goals of designers and model training.

Quantity vs Quality. Our work investigates the tradeoff between data quantity and quality for training UI generation models. Previously, the effectiveness of ML annotation systems (e.g., feedback interfaces) has often been measured by the speed at which they could produce responses from human annotators [23, 42]. For example, one proposed metric for evaluating annotation interfaces is the number of training examples that could be collected through the interaction in a given time-frame [42]. While these measurements may be more reflective of utility for tasks where voting [23] or other forms of redundant checking [42, 43] is possible, we found that they were not useful for UI design feedback interfaces. In our work, we found that while some interactions require more time and effort to complete (e.g., design revision), they often result in better data quality. Our approach to evaluating UI designs involved collecting feedback across multiple UI design feedback interfaces to train UI generation models and evaluate those models with human judges. Our results ultimately showed that sketching exhibited the best performance and balance between data quality and quantity. Nevertheless, the development of more standardized and easily calculable heuristics for UI design evaluation at an earlier stage remains an important direction for future work, especially to inform the earlier stages of design.

6.3 Limitations & Future Work

In this paper, we show that our data collection and modeling approach leads to significant improvements to the quality of generated UI code by learning from designer feedback. We see multiple avenues for improvement for our work.

First, our current work provides limited validation of our data collection approach. Our designer feedback study with designers was relatively short (one hour) and designers evaluated synthetically-generated UIs, which are typically of lower quality than professional designs. Our study design was motivated by several practical factors (e.g., our desire to control study conditions) and data confidentiality (designers may normally work on proprietary, unreleased designs). A future study could investigate the use of more realistic training data consisting of professional designs to build a stronger reward model and could yield new insights into the real-world feasibility of the approach.

We see numerous opportunities to collect other types of designer feedback and developing machine learning approaches to learn from them. When asked about other types of feedback common in their work², designers in our study often referenced numerous methodologies in the HCI literature, such as usability studies, cognitive walkthroughs, and analyzing user engagement metrics. The

techniques we used to collect feedback would likely need to be adapted to infer complex interactive properties related to usability and accessibility.

Finally, the machine learning formulation to learn from these other types of design feedback data would need to be updated as well to translate these artifacts (e.g., recordings or transcripts) into machine-learnable datasets. We designed our training setup for a single-screen evaluation. For example, we employ a visual language model (i.e., CLIP) as a reward model to score a single image input. Evaluating higher-level aspects of UX and interaction design would necessitate a more complex evaluator module that could potentially crawl and interact with a running version of the application (e.g., a UI agent). The generator module must also learn from this feedback and use it to generate functionality for app navigation and handling interaction (e.g., a login flow). We expect numerous opportunities for incorporating design expertise in advancements in LLM model training and their HCI applications.

7 Conclusion

In this paper, we presented interactions for collecting designer feedback based on common workflows, such as commenting, sketching, and revising. We introduce methods for converting this designer feedback into machine learnable datasets for UI assessment and generation. We conducted a designer feedback study where we asked twenty one designers to give feedback on a set of synthetically generated UIs and developed techniques to generate a designer feedback dataset from them. We analyzed this data to validate the quality of its labels, showing much higher levels of label agreement than existing ranking methods for output comparison. To validate the usefulness of our data, we finetuned UI code generation models in a RLHF configuration with a reward model trained from different forms of designer feedback. We conducted an arena-style model evaluation study comparing the performance of our models and baselines through repeated blind evaluations of the generated UI design outputs. Our results showed that feedback data collected from natural designer interactions led to better model performance compared to ranking data, which is widely used for model training. In addition, our best-performing model outperformed all test baselines, including a larger proprietary model, highlighting the impact of a small amount of high-quality expert feedback. Our work suggests applying designer-aligned interaction techniques is beneficial for training models to generate UIs.

References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023).
- [2] Reed Albergo. 2023. *Microsoft pushes the boundaries of small AI models with big breakthrough*. Semafor. <https://www.semafor.com/article/11/01/2023/microsoft-pushes-the-boundaries-of-small-ai-models>
- [3] Tony Beltramelli. 2018. pix2code: Generating code from a graphical user interface screenshot. In *Proceedings of the ACM SIGCHI symposium on engineering interactive computing systems*. 1–6.
- [4] Param Biyani, Yasharth Bajpai, Arjun Radhakrishna, Gustavo Soares, and Sumit Gulwani. 2024. Rubicon: Rubric-based evaluation of domain-specific human ai conversations. In *Proceedings of the 1st ACM International Conference on AI-Powered Software*. 161–169.
- [5] Bill Buxton. 2010. *Sketching user experiences: getting the design right and the right design*. Morgan kaufmann.

²The question was aimed at understanding different mechanisms for recording feedback; however, many participants responded with methodologies of collecting it.

- [6] Michelle Carney, Barron Webster, Irene Alvarado, Kyle Phillips, Noura Howell, Jordan Griffith, Jonas Jongejan, Amit Pitaru, and Alexander Chen. 2020. Teachable machine: Approachable Web-based tool for exploring machine learning classification. In *Extended abstracts of the 2020 CHI conference on human factors in computing systems*. 1–8.
- [7] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374* (2021).
- [8] Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Banghua Zhu, Hao Zhang, Michael Jordan, Joseph E Gonzalez, et al. 2024. Chatbot arena: An open platform for evaluating llms by human preference. In *Forty-first International Conference on Machine Learning*.
- [9] Peitong Duan, Chin-Yi Cheng, Gang Li, Bjoern Hartmann, and Yang Li. 2024. UICrit: Enhancing Automated Design Evaluation with a UI Critique Dataset. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–17.
- [10] Peitong Duan, Jeremy Warner, Yang Li, and Bjoern Hartmann. 2024. Generating automatic feedback on ui mockups with large language models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–20.
- [11] Krzysztof Gajos, David Christianson, Raphael Hoffmann, Tal Shaked, Kiera Henning, Jing Jing Long, and Daniel S Weld. 2005. Fast and robust interface generation for ubiquitous applications. In *UbiComp 2005: Ubiquitous Computing: 7th International Conference, UbiComp 2005, Tokyo, Japan, September 11-14, 2005. Proceedings* 7. Springer, 37–55.
- [12] Krzysztof Gajos and Daniel S Weld. 2004. SUPPLE: automatically generating user interfaces. In *Proceedings of the 9th international conference on Intelligent user interfaces*. 93–100.
- [13] Krzysztof Gajos and Daniel S Weld. 2005. Preference elicitation for interface optimization. In *Proceedings of the 18th annual ACM symposium on User interface software and technology*. 173–182.
- [14] 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.
- [15] Björn Hartmann, Sean Follmer, Antonio Ricciardi, Timothy Cardenas, and Scott R Klemmer. 2010. D. note: revising user interfaces through change tracking, annotations, and alternatives. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 493–502.
- [16] Jiwoo Hong, Noah Lee, and James Thorne. 2024. Orpo: Monolithic preference optimization without reference model. *arXiv preprint arXiv:2403.07691* (2024).
- [17] Forrest Huang, John F Canny, and Jeffrey Nichols. 2019. Swire: Sketch-based user interface retrieval. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–10.
- [18] Binyuan Hui, Jian Yang, Zeyu Cui, Jiaxi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, et al. 2024. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186* (2024).
- [19] Jiaming Ji, Boyuan Chen, Hantao Lou, Donghai Hong, Borong Zhang, Xuehai Pan, Tianyi Alex Qiu, Juntao Dai, and Yaodong Yang. 2024. Aligner: Efficient alignment by learning to correct. *Advances in Neural Information Processing Systems* 37 (2024), 90853–90890.
- [20] Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, et al. 2024. T"ulu 3: Pushing frontiers in open language model post-training. *arXiv preprint arXiv:2411.15124* (2024).
- [21] James A Landay. 1996. SILK: sketching interfaces like crazy. In *Conference companion on Human factors in computing systems*. 398–399.
- [22] James A Landay and Brad A Myers. 1995. Interactive sketching for the early stages of user interface design. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 43–50.
- [23] Greg Little, Lydia B Chilton, Max Goldman, and Robert C Miller. 2010. TurkIt: human computation algorithms on mechanical turk. In *Proceedings of the 23rd annual ACM symposium on User interface software and technology*. 57–66.
- [24] Anton Lozhkov, Raymond Li, Loubna Ben Allal, Federico Cassano, Joel Lamy-Poirier, Nouamane Tazi, Ao Tang, Dmytro Pykhtar, Jiawei Liu, Yuxiang Wei, et al. 2024. Starcoder 2 and the stack v2: The next generation. *arXiv preprint arXiv:2402.19173* (2024).
- [25] Kurt Luther, Amy Pavel, Wei Wu, Jari-lee Tolentino, Maneesh Agrawala, Björn Hartmann, and Steven P Dow. 2014. CrowdCrit: crowdsourcing and aggregating visual design critique. In *Proceedings of the companion publication of the 17th ACM conference on Computer supported cooperative work & social computing*. 21–24.
- [26] Jeffrey Nichols, Brad A. Myers, Michael Higgins, Joseph Hughes, Thomas K. Harris, Roni Rosenfeld, and Mathilde Pignol. 2002. Generating remote control interfaces for complex appliances. In *Proceedings of the 15th Annual ACM Symposium on User Interface Software and Technology (Paris, France) (UIST '02)*. Association for Computing Machinery, New York, NY, USA, 161–170. doi:10.1145/571985.572008
- [27] Jakob Nielsen and Rolf Molich. 1990. Heuristic evaluation of user interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 249–256.
- [28] Jasper O'Leary, Holger Winnemöller, Wilmo Li, Mira Dontcheva, and Morgan Dixon. 2018. Charrette: Supporting In-Person Discussions around Iterations in User Interface Design. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [29] OpenAI. 2025. *GPT-5 System Card*. Technical Report. OpenAI. <https://cdn.openai.com/gpt-5-system-card.pdf>
- [30] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems* 35 (2022), 27730–27744.
- [31] Angel Puerta and Jacob Eisenstein. 1998. Towards a general computational framework for model-based interface development systems. In *Proceedings of the 4th international conference on Intelligent user interfaces*. 171–178.
- [32] Angel R Puerta and David Maulsby. 1997. MOBI-D: a model-based development environment for user-centered design. In *CHI'97 Extended Abstracts on Human Factors in Computing Systems*. 4–5.
- [33] Jie Ren, Samyam Rajbhandari, Reza Yazdani Aminabadi, Olatunji Ruwase, Shuangyan Yang, Minjia Zhang, Dong Li, and Yuxiong He. 2021. {Zero-offload}: Democratizing {billion-scale} model training. In *2021 USENIX Annual Technical Conference (USENIX ATC 21)*. 551–564.
- [34] Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Romain Sauvestre, Tal Remez, et al. 2023. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950* (2023).
- [35] Jaeyung Seol, Seojun Kim, and Jaeyun Yoo. 2024. PosterLlama: Bridging Design Ability of Language Model to Content-Aware Layout Generation. In *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part LXXXII* (Milan, Italy). Springer-Verlag, Berlin, Heidelberg, 451–468. doi:10.1007/978-3-031-73007-8_26
- [36] Omar Shaikh, Michelle S. Lam, Joey Hejna, Yijia Shao, Hyundong Justin Cho, Michael S. Bernstein, and Diyi Yang. 2025. Aligning Language Models with Demonstrated Feedback. In *The Thirteenth International Conference on Learning Representations*. <https://openreview.net/forum?id=1qGkuxI9UX>
- [37] Ben Shneiderman, Catherine Plaisant, Maxine S. Cohen, Steven M. Jacobs, Niklas Elmqvist, and Nicholas Diakopoulos. 2016. *Eight Golden Rules of Interface Design* (6 ed.). Pearson, Chapter 3.3.4. <https://dl.acm.org/doi/10.5555/3033040>
- [38] Kihoon Son, DaEun Choi, Tae Soo Kim, and Juho Kim. 2024. Demystifying tacit knowledge in graphic design: Characteristics, instances, approaches, and guidelines. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [39] Zecheng Tang, Chenfei Wu, Juntao Li, and Nan Duan. 2024. LayoutNUWA: Revealing the Hidden Layout Expertise of Large Language Models. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=qCUWVT0Ayy>
- [40] Hallvard Trætteberg. 2002. Model-based user interface design. (2002).
- [41] Lewis Tunstall, Edward Beeching, Nathan Lambert, Nazneen Rajani, Kashif Rasul, Younes Belkada, Shengyi Huang, Leandro Von Werra, Clémentine Fourrier, Nathan Habib, et al. 2023. Zephyr: Direct distillation of lm alignment. *arXiv preprint arXiv:2310.16944* (2023).
- [42] Luis Von Ahn. 2006. Games with a purpose. *Computer* 39, 6 (2006), 92–94.
- [43] Luis Von Ahn, Benjamin Maurer, Colin McMillen, David Abraham, and Manuel Blum. 2008. recaptcha: Human-based character recognition via web security measures. *Science* 321, 5895 (2008), 1465–1468.
- [44] Bryan Wang, Gang Li, Xin Zhou, Zhouong Chen, Tovi Grossman, and Yang Li. 2021. Screen2words: Automatic mobile UI summarization with multimodal learning. In *The 34th Annual ACM Symposium on User Interface Software and Technology*. 498–510.
- [45] WebDev Arena [n. d.]. WebDev Arena: AI Battle to build the best website. <https://web.lmarena.ai/>. Accessed: 2025-04-08.
- [46] Jason Wu, Yi-Hao Peng, Xin Yue Amanda Li, Amanda Swearngin, Jeffrey P Bigham, and Jeffrey Nichols. 2024. UICLIP: a data-driven model for assessing user interface design. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–16.
- [47] Jason Wu, Eldon Schoop, Alan Leung, Titus Barik, Jeffrey P Bigham, and Jeffrey Nichols. 2024. Uicoder: Finetuning large language models to generate user interface code through automated feedback. *arXiv preprint arXiv:2406.07739* (2024).
- [48] Amber Xie, Chin-Yi Cheng, Forrest Huang, and Yang Li. 2024. Leveraging Human Revisions for Improving Text-to-Layout Models. *arXiv preprint arXiv:2405.13026* (2024).
- [49] Shusheng Xu, Wei Fu, Jiaxuan Gao, Wenjie Ye, Weilin Liu, Zhiyu Mei, Guangju Wang, Chao Yu, and Yi Wu. 2024. Is dpo superior to ppo for llm alignment? a comprehensive study. *arXiv preprint arXiv:2404.10719* (2024).
- [50] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388* (2025).

- [51] Ning Yu, Chia-Chih Chen, Zeyuan Chen, Rui Meng, Gang Wu, Paul Josel, Juan Carlos Niebles, Caiming Xiong, and Ran Xu. 2024. LayoutDETR: Detection Transformer Is a Good Multimodal Layout Designer. In *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part XX* (Milan, Italy). Springer-Verlag, Berlin, Heidelberg, 169–187. doi:10.1007/978-3-031-72661-3_10
- [52] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging LLM-as-a-judge with MT-bench and Chatbot Arena. In *Proceedings of the 37th International Conference on Neural Information Processing Systems* (New Orleans, LA, USA) (NIPS '23). Curran Associates Inc., Red Hook, NY, USA, Article 2020, 29 pages.
- [53] Lianmin Zheng, Ying Sheng, Wei-Lin Chiang, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Chatbot Arena: Benchmarking LLMs in the Wild with Elo Ratings. <https://lmsys.org/blog/2023-05-03-arena/>. LMSYS Org Blog; accessed 2025-09-04.

A Model Hyperparameters

We provide all hyperparameters used for our model training experiments in Table 1.

B Model Prompts

We used several prompts in our experiments to train and run large language models.

We used the the following prompt for generating HTML web pages from a short natural language description. We also fine-tuned our model using prompts formatted in this template. This prompt was determined by trial and error and manual inspection on a small number of test cases.

```
provide the complete HTML code for a web page
implemented with only tailwind CSS and font
awesome icons. do not use any templating languages
like jinja. the result should resemble an award-
winning iOS app. include realistic and complete
placeholder data. do not treat this as the
starting point for an app - it should be the
mockup of a final complete UI. remember to include
alt text for all images. do not use javascript.
do not use SVGs. here is a description of the
webpage: <natural language description>
```

We use several prompts for evaluating screenshots. Our approach was to compute an augmented text embedding using a mix of positive and negative prompts.

We use the original prompt format described in UIClip [46] as the positive prompt.

```
ui screenshot. well-designed. <natural language
description>
```

We also computed a negative prompt to represent a poorly designed version of the screen.

```
ui screenshot. poor design. <natural language
description>
```

Finally, we computed a second negative prompt to represent empty or overly simple UIs.

```
ui screenshot. poor design. empty screen
```

Text embeddings were computed for each of these prompts and combined together using the following equation.

$$\mathbf{v}^* = \mathbf{v}_{pos} - 0.5 \cdot (0.9 \cdot \mathbf{v}_{neg} + 0.1 \cdot \mathbf{v}_{empty}) \quad (3)$$

\mathbf{v}^* refers to the final text embedding used for UIClip score calculation. \mathbf{v}_{pos} refers to the positive text embedding. \mathbf{v}_{neg} refers to the negative text embedding. \mathbf{v}_{empty} refers to the empty text embedding. The weighting between the positive and negative embeddings were determined by trial and error and manual inspection on a small number of UI screenshots.

We prompted Qwen2.5-Coder 32B to improve UIs given a list of designer-authored comments.

```
i have implemented a website using only html,
tailwind css, and font awesome icons.
```

```
```html
```

**Table 1: Hyperparameters used for our modeling training experiments.**

Algorithm	Hyperparameters	Value
Reward Models	max optim steps	100
	batch size	32
	weight decay	0.2
	learning rate	1e-3
	margin	1e-2
	UIClip aug prob	0.5
Generator Models	odds ratio weight	1.0
	effective batch size	8
	context length	4096
	learning rate	5e-6

```
<original UI HTML code>
```
```

```
a designer has wrote some notes and feedback:
"<list of comments>"
```

```
incorporate this feedback into the website code.
you must respond with the entire code
implementation. do not use comments that are
placeholders for the original code.
```

We used a similar prompt to improve UIs given a list of region-grounded annotations provided by designers.

```
i have implemented a website using only html,
tailwind css, and font awesome icons.
```

```
```html
```

```
<original UI HTML code>
```
```

```
a designer has wrote some notes and feedback for
several regions of the HTML:
"<list of comments paired with HTML snippets>"
```

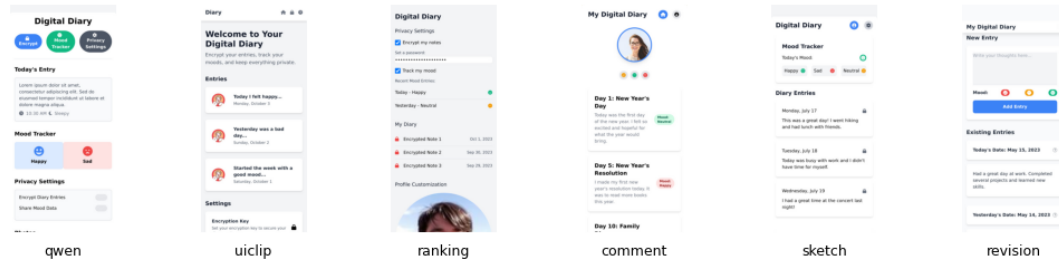
```
incorporate this feedback into the website code.
you must respond with the entire code
implementation. do not use comments that are
placeholders for the original code.
```

C Example Outputs

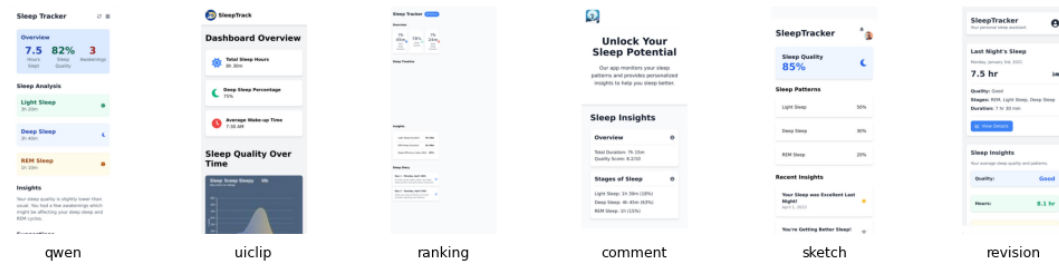
We provide some example outputs from our model evaluation experiment in Figure 5.

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009

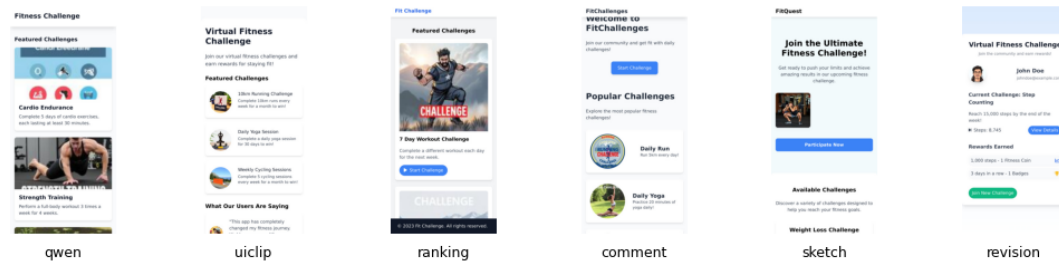
A digital diary app with encryption, mood tracking, and privacy settings.



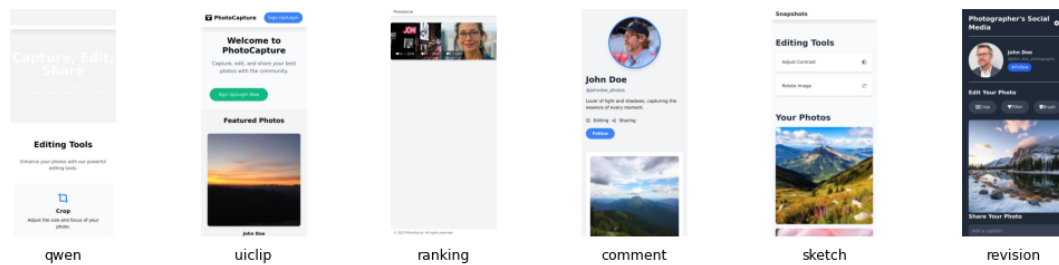
A sleep tracking app that monitors sleep patterns, quality, and provides personalized insights.



A virtual fitness challenge app that connects users with fitness challenges and rewards for participation.



A social media app for photographers with editing tools, sharing, and community forums.



'A project management app with task boards, Gantt charts, and team collaboration tools.'

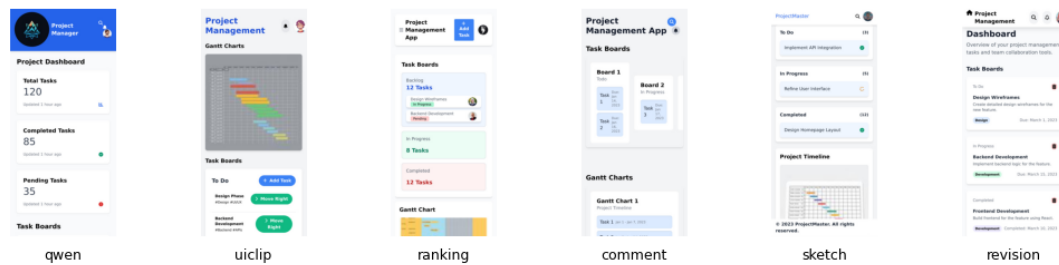




Figure 5: Figure shows rendered output of six models tested in the feedback comparison study. We rendered model outputs for five randomly sampled text descriptions from our evaluation set.

A virtual travel companion app with real-time travel updates, local recommendations, and itinerary adjustments.




qwen2.5+sketch




qwen2.5-3b




qwen2.5-3b+sketch



qwen3




qwen3+sketch



gpt5

A mindfulness and breathwork app that guides you through various breathing exercises for relaxation.




qwen2.5+sketch




qwen2.5-3b



qwen2.5-3b+sketch



qwen3




qwen3+sketch



gpt5


A language learning app with immersive stories, cultural insights, and interactive dialogues.




qwen2.5+sketch




qwen2.5-3b




qwen2.5-3b+sketch



qwen3




qwen3+sketch




gpt5


A time management app that organizes daily tasks, sets priorities, and tracks productivity.




qwen2.5+sketch



qwen2.5-3b




qwen2.5-3b+sketch



qwen3




qwen3+sketch




gpt5


A dream journal app that helps you record and analyze your dreams with tips for better sleep.




qwen2.5+sketch




qwen2.5-3b




qwen2.5-3b+sketch



qwen3



qwen3+sketch



gpt5

Figure 6: Figure shows rendered output of six models tested in the model generalization study. We rendered model outputs for five randomly sampled text descriptions from our evaluation set.