HelpViz: Automatic Generation of Contextual Visual Mobile Tutorials from Text-Based Instructions

Mingyuan Zhong*
University of Washington
Seattle, WA
myzhong@cs.washington.edu

Peggy Chi
Google Research
Mountain View, CA
peggychi@google.com

Gang Li
Google Research
Mountain View, CA
leebird@google.com

Yang Li
Google Research
Mountain View, CA
liyang@google.com

ABSTRACT

We present HelpViz, a tool for generating contextual visual mobile tutorials from text-based instructions that are abundant on the web. HelpViz transforms text instructions to graphical tutorials in batch, by extracting a sequence of actions from each text instruction through an instruction parsing model, and executing the extracted actions on a simulation infrastructure that manages an array of Android emulators. The automatic execution of each instruction produces a set of graphical and structural assets, including images, videos, and metadata such as clicked elements for each step. HelpViz then synthesizes a tutorial by combining parsed text instructions with the generated assets, and contextualizes the tutorial to user interaction by tracking the user's progress and highlighting the next step. Our experiments with HelpViz indicate that our pipeline improved tutorial execution robustness and that participants preferred tutorials generated by HelpViz over text-based instructions. HelpViz promises a cost-effective approach for generating contextual visual tutorials for mobile interaction at scale.

CCS CONCEPTS

• **Human-centered computing** → **Interactive systems and tools**.

KEYWORDS

mobile tutorials, automatic generation, contextual help

ACM Reference Format:

Mingyuan Zhong, Gang Li, Peggy Chi, and Yang Li. 2021. HelpViz: Automatic Generation of Contextual Visual Mobile Tutorials from Text-Based Instructions. In *The 34th Annual ACM Symposium on User Interface Software and Technology (UIST '21)*, October 10–14, 2021, Virtual Event, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3472749.3474812>

*This work was completed while the author was an intern at Google Research.

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

UIST '21, October 10–14, 2021, Virtual Event, USA
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8635-7/21/10.
<https://doi.org/10.1145/3472749.3474812>

1 INTRODUCTION

Mobile phones today provide a myriad of functionalities for a variety of everyday tasks. However, it is not always clear to a user how to access specific functionalities, such as dimming the screen or enabling tethering on the phone. It is thus common for mobile users to look for instructions on how to perform a specific task on the web.

Existing online tutorials are typically created by professionals or domain experts. One of the popular formats is a web document, which presents a sequence of steps where each step includes text instructions that are often accompanied by one or more images [3].

To ease the generation of high-quality tutorials, previous work has proposed techniques to automate a creation process through demonstrations by expert users [3, 13] or composing by a community [26, 36]. However, the process of tutorial creation still needs a significant amount of human effort. As software is frequently updated, a tutorial needs to be constantly revised given that a task procedure might contain missing or mismatched details or its supporting screenshot images become outdated.

Thus, it is important to further reduce the cost of tutorial creation. In this work, we focus on utilizing the vast corpus of text instructions that are already available on the web. We tailor their content and presentation to the user's specific interaction context based on an understanding of the device infrastructure. Prior work has suggested visualization techniques for tutorials that guide a mobile user through a task workflow by highlighting UI targets on a mobile device [35]. In addition, recent research demonstrated that by enhancing a text article with animated images, readers could better follow and comprehend the content [22]. Our goal is to convert text-based instructions to a visual tutorial that is seamlessly integrated into the user's interaction context on the mobile device.

We present HelpViz, a tool that automatically transforms a batch of text instructions to contextual visual tutorials that can be interactively navigated on a mobile device (see Fig. 1). Our pipeline (1) parses the instructional text to extract actions described in each step using a pre-trained deep model [25], (2) segments text instructions into steps based on the extracted actions, (3) executes the actions in a simulation infrastructure that manages an array of Android emulators, (4) synthesizes visual tutorials by combining segmented text and assets generated from execution, and finally (5) serves these tutorials in situ in the context of mobile interaction. Because such tutorials are equipped with screen assets and metadata captured

Figure 1: HelpViz converts text instructions that can be found abundantly on the web to visual tutorials and then serves these tutorials in the context of mobile interaction. (1)–(4): Our pipeline converts text instructions to visual tutorials. HelpViz generates the assets such as screenshots and interaction metadata of each instruction by actually executing it on a simulation infrastructure. It then synthesizes visual tutorials by combining these assets with segmented text instructions on a mobile device. (5): HelpViz Viewer serves visual tutorial in the interaction context of the user.

from automatic execution of instructions, HelpViz is capable of tracking a user's progress and automatically anchoring the tutorial to the current step.

Our evaluation with 10 participants indicated that HelpViz received positive feedback from participants, for its rich graphical representation and progress tracking features, over text-based instructions. Our paper makes the following contributions:

- The first pipeline that automatically generates visual tutorials for mobile tasks from raw text instructions, which can be those directly crawled on the web. Such a synthesized tutorial is served in the specific context of user interaction on the mobile device that automatically advances the focus of the tutorial by matching the user's progress.
- A set of algorithmic techniques and interaction design to address errors and incompatibility from automatic tutorial generation based on an AI model: *beam-search* that addresses inaccuracy of the instruction parsing model and related interaction design to present alternatives in a tutorial, and *look-ahead* that addresses the incompatibility between an instruction and an emulator during batch execution—these manifest how the imperfection of an automatic process can be addressed by closely coupling algorithmic and interaction design. These techniques were shown to substantially improve the robustness of tutorial synthesis.
- A set of experiments and findings that investigate the feasibility of visual tutorial generation and its impact on end-user experiences.

2 RELATED WORK

Our work is related to three research areas in the literature: tutorial creation, contextual tutorials, and the general topic of incorporating an AI model in an interactive system.

2.1 Structural Tutorial Creation

The creation of tutorials for software has traditionally been completed by either user demonstration [3, 7, 13, 28] or crowdsourcing [12, 20]. Tutorials generated by *demonstration* usually require tracing a user's interaction flow while recording a screencast video of the user interface. Given a demonstration, MixT detects the most relevant areas in a step, and generates both screenshots and edited videos in a mixed-media document [3]. Torta allows creators to specify validation criteria for specific steps through file matching or explicit scripting [28]. PPTutorial adds overlays on UI elements to facilitate locating them [23]. On the other hand, *crowdsourcing* can be an effective method for tutorial creation with the advantage of better scalability. ShowMeHow presents command strings in different applications and utilizes the crowd to match similar functionalities in order to translate commands between apps [31]. Codepourri enables crowd learners to collaboratively annotate each step in a tutorial for a piece of code [20], while ToolScape allows learners to apply step-by-step annotations to instructional videos [12].

Previous approaches for tutorial creation commonly require human participation, in the form of user demonstration or crowdsourcing. In contrast, HelpViz automatically converts text instructions that exist on the web to visual tutorials by parsing and executing these instructions. Instead of logging the steps of a human user, or

requiring manual collection and editing of assets, our system automatically gathers necessary assets generated from the execution of instructions, and synthesizes visual tutorials by combining the assets and text instructions.

2.2 Context-Aware Tutorial Following

The analysis of task executions allows users to discover relevant or more efficient workflows. Delta visualizes software workflows at multiple granularities and allows users to compare the trade-offs between them [21]. Collective navigational traces within an instructional video have been used to facilitate the navigation within a tutorial [19]. To recommend new workflows, Wang et al. [36] extracted workflow topics and frequent command patterns. Khan et al. personalized the recommendations by utilizing web documentations and a user's command history [18]. Recent work has shown that by presenting its two-level hierarchy, a tutorial can be more easily consumed, especially for longer sequences [34].

In order to understand the current context, existing systems typically utilize accessibility APIs [11, 32] or image-based analysis [29, 35] to determine the user action and match it against the tutorial. Interaction designs that previous work often use for desktop applications include direct overlays [17] and tooltips [14], which can be linked with relevant images or videos. Previous work also focused on enhancing video-based tutorials, e.g., pausing video when the user fell behind [29] or presenting relevant videos to the user task [11]. Torii [15] enables flexible programming tutorials by linking code copies and edits with a reference implementation, preserving the context for code snippets. InterTwine [10] bridges the web browser and desktop application through an explicit link, enabling shared context between the applications and easier help seeking.

On mobile systems, EverTutor [35] detects user context through matching screenshots with pre-recorded templates. AidMe [32] allows voice-based tutorial creation and playback tailored to visually impaired users through an accessibility service. Similar to previous work, a tutorial generated by HelpViz is presented to a user contextually. Specifically, HelpViz utilizes the accessibility APIs to match the user's current screen with a step in the tutorial. Importantly, the assets and metadata that enable such context awareness are automatically captured by the HelpViz pipeline, without any human effort.

2.3 Incorporating AI Models into Interactive Systems

Our work is broadly related to the topic of incorporating the output of an AI model into an interaction task. Because the imperfection of the instruction parsing model or the inherent mismatch between an instruction and an emulator, there can be errors for generating tutorial assets during the automatic execution. Several strategies have emerged in the literature for integrating an AI model. Previous work has proposed interactive machine learning that iteratively refines an AI model using user input (e.g., [8, 9]). Previous work has also developed interaction techniques to allow users to easily navigate alternative model results (e.g., [1, 30]).

In our task, we aim to generate visual tutorials from a batch of text instructions in parallel, without human intervention. Our intended users are help seekers who are likely unfamiliar with

the issues themselves. As a result, interactively refining tutorial generation is not feasible in HelpViz. Instead, we choose to reduce tutorial generation failures by capturing model and emulator execution uncertainty, via beam search [27] and look-ahead, and then presenting the captured uncertainty and alternatives to the user via appropriate interaction designs.

3 HELPviz

HelpViz consists of a pipeline that converts text instructions to visual tutorials (see Figure 1.1–1.4), and a mobile tool, HelpViz Viewer, that serves a visual tutorial in the interaction context of the user (Figure 1.5). The web offers a rich corpus of text instructions, e.g., online help documents for how to use Android¹ or iOS². For example, to use less mobile data, users can find instructions that guide them through a workflow: “*Open your device's settings app. Tap network & internet. Click data usage > data saver. Turn data saver on.*” It requires a user matching the information with their device screen, one action at a time. To generate visual tutorials for such instructions with useful graphical assets, we propose an automatic approach to parse the text instructions, execute the actions on a device emulator, and interactively present on user's device. Below we describe the components in our pipeline.

3.1 Instruction Parsing & Segmentation

We first use an instruction parsing model to extract the action phrase tuples that describe each *step* of the text *instruction*³, using a pre-trained deep model [25]. The model takes the raw text of an instruction as input. It outputs a sequence of action tuples, which consists of the operation (e.g., “click”, “turn on”) and target object (the element on which to apply the operation) at each step (see Fig. 1.1).

The extracted action tuples from instruction parsing are used in two separate processes, instructional text segmentation and automatic execution. We split the raw text of an instruction into multiple steps by matching each phrase tuple to its respective sentence(s) in the instruction (see Fig. 1.2). Our pipeline generates visual assets for each step of the task by executing the extracted actions on an Android emulator, similar to prior work that mines UI layouts [5, 6].

3.2 Batch Execution on Emulators & Tutorial Synthesizing

We developed an infrastructure that can automatically execute a batch of parsed instructions in parallel (see Fig. 1.3), by running each instruction on a separate emulator, which is a standard Android Emulator spawned on a distributed computing platform controlled through ADB commands (e.g., sending a tap command). Given an action tuple extracted from an instruction, our instruction *executor* identifies and acquires a target element on the current UI screen of the emulator, which triggers a transition to the next screen. However, it is likely that the target element described in the tuple cannot be immediately located, e.g., when the screen is still loading or the element is outside of the screen. The executor will either

¹Android Help, <https://support.google.com/android>

²iPhone - Official Apple Support, <https://support.apple.com/iphone>

³In this paper, an *instruction* refers to a text description of a sequence of steps that complete a task, while a *step* refers to a single executable action from the instruction.

wait for the screen to be loaded or attempt to scroll the screen to find the element. Note that these dynamic aspects are not addressed in previous work [25], which assumes a static setup for supervised sequence modeling.

While executing an instruction, the executor collects a rich set of assets for each action taken, including UI screenshots, screencast video clips that capture transitional UI animations, and metadata such as UI view hierarchies and elements being clicked. These assets are necessary for synthesizing visual tutorials and enabling step matching to serve a tutorial contextually that we will describe in the following sections. Our pipeline then combines the assets captured by the executor and the segmented instruction steps from the previous stage to compose visual tutorials (see Fig. 1.4). HelpViz structures a tutorial as a sequence of steps where each step consists of text instructions and its corresponding screenshot overview and close-up view as well as a video segment.

There are two challenges in automatically executing instructions crawled from the web on an emulator. One is due to the inaccurate inference results of the parsing model. The other issue is rooted in the mismatch between the instruction and the actual emulator interfaces. Both issues can cause failures for automatic execution. We address these issues as following.

3.2.1 Addressing Model Parsing Errors. By default, the instruction parsing model [25] extracts a best sequence of operation-object tuples for actions from the text instructions by greedy decoding. However, with the presence of multiple steps in an instruction, such greedy decoding may generate a sub-optimal sequence if there are errors in the decoded actions, because for each step it only selects the best action depending on previous actions, which can be a local optimum considering the full sequence. To address this issue, we use *beam search* to track the top k decoded sequences at each step to improve the overall parsing result.

Specifically, the instruction parsing model parses an instruction and generates multiple possible action sequences (beams). In generating the action sequences, the model aims to maximize the probability of the overall action sequence, rather than the probability of individual actions. In our experiments, we chose to execute the top 3 sequences (i.e., $k = 3$) of an instruction. Thus, our emulator infrastructure launches three emulator instances for each instruction, and each of the decoded sequences is executed independently on a separate emulator instance. To merge the execution results into a consistent form that can be served to the user, HelpViz first discards the results from the executions that cannot reach the final step. Then, for executions that reach the final step, HelpViz merges the steps shared by executions and preserves steps that differ as alternatives. Finally, HelpViz sorts these alternatives based on the decoding score of the sequence they belong to. This approach allows the model to output several slightly different parsing results and reduces the likelihood that the greedy-decoded sequence contains errors and leads an dead end during execution.

Although beam search is a common technique in decoding models, it has special implications in our case for both system and interaction design. For the system design, our executor needs to handle multiple alternatives of executing an instruction by spawning multiple threads and then merging the results of the execution of these threads, as discussed above. For the interaction design, we

need to design appropriate interfaces for the user to examine these alternatives, which we will discuss later.

3.2.2 Addressing Instruction-Emulator Mismatch. It is possible that an instruction does not exactly match the app UIs on the emulator. For example, the instruction can be outdated as the UI of an app running on the emulator has been recently updated. In order to mitigate the issue, we implemented *look-ahead* in the executor to maximally execute an instruction by skipping steps that cannot be realized. Specifically, the executor allows a maximum of five unsuccessful actions (tapping or scrolling) for each task. If the executor still could not perform the action described in the current step after five attempts (i.e., the target element in the current step does not appear on the screen), it then skips the current step, looks for the target element described in the next step, and attempts to execute the next step's action if found. For example, an instruction may assume that an expandable view is folded and include a step of "tap *show all items*". During execution, the view might have already been expanded on the emulator (resulting from a prior action or a difference in app versions) and skipping the corresponding step in the tutorial would be warranted. In this case, after failing to find the button "*show all items*" after five scrolls, the executor would look for an item in the list that corresponds to the next action in the instruction.

3.3 Serving a HelpViz Tutorial in Context

To support users following instructions, we develop HelpViz *Viewer* that serves a HelpViz visual tutorial directly on the mobile device where the user seeks help. In Fig. 2a, the Viewer loads a generated HelpViz tutorial for turning off all notifications. Each step has two screenshot thumbnails: the close-up view is a cropped image around the target element, and the overview of the UI screen highlights the cropped area, similar to prior work [3, 4]. If any step requires scrolling to reach the target element, the thumbnail will be substituted with a short scrolling animation on that screen when the step is either scrolled to the top (including auto-scroll by HelpViz) or being tapped on. When multiple alternatives are available for a step, a user can swipe horizontally to switch between the alternatives.

To follow the instructions, a user switches to the target app and performs the actions described in the tutorial (Fig. 2b). Once the user switches back to HelpViz for further help, HelpViz Viewer automatically scrolls to the step that the user is working on and highlights that step (Fig. 2c). The highlight color automatically fades out within one second to avoid distracting the user.

To synchronize the current step in the tutorial with the actual progress of the user, HelpViz registers a background accessibility service to compare the UI elements on the app screen with the elements of each step in the tutorial. The accessibility service captures the current view hierarchy of the target app every time the screen updates by listening to system *AccessibilityEvents* and records text-related properties (i.e., *Text*, *ContentDescription*, and *HintText*) of each UI element. These texts are matched against the texts stored for each step in tutorial metadata using Jaccard index [16] (also known as Intersection over Union) as a similarity metric. Specifically, it is computed as the number of common tokens in both sets of text (i.e., their *intersection*) divided by the number of all unique tokens in the two sets (i.e., their *union*). To illustrate the matching

Figure 2: HelpViz Viewer serves an automatically-generated visual tutorial in the context of user interaction. In this example, the tutorial illustrates the task of “turning off all notifications.” The original text instruction is: “Open your device’s *Settings* app. Tap *Apps & notifications*. Click *Notifications*. Tap *On lock screen*. Click *Don’t show notifications at all*.” (Formatting added by the authors for clarity.) Note that in (b), elements with text matched to Step 3 are highlighted to illustrate contextual matching. The matching allows HelpViz to scroll to and highlight that step, illustrated in (c). The blue boxes in (b) are *not* present during user interaction.

result, Fig. 2b highlights the elements with text matched to Step 3 of the tutorial.

In rare cases, when multiple dynamically generated elements are present on a screen, the accessibility service may fail to match the current screen with an extracted screen from the pipeline if text in most of these elements are different from what we captured during emulation. As a fail-safe, HelpViz remains at the last viewed step in this case.

4 PIPELINE PERFORMANCE EXPERIMENTS

The HelpViz pipeline robustly handles instructions crawled from the web. We converted the PixelHelp dataset⁴, which contained 187 multi-step text instructions for using Android Pixel phones, to visual tutorials using the pipeline. The execution of each instruction completed in less than 15 minutes and all the executions could be performed in parallel. In our tests, we ran 50 emulators simultaneously to speed up tutorial generation. Including attempts from beam search and look-ahead, our pipeline executed 1997 actions in total and produced and gathered necessary assets (see Table 1 for a summary).

In this experiment, we investigate whether beam search and look-ahead, the two techniques we employed in our pipeline for robust batch execution (Section 3.2), improves the performance of our automatic instruction execution (see Table 2). To do so, we mainly look at two metrics: (1) *steps executed*, defined as the average number of steps successfully executed of each instruction, and (2) *completion rate*, defined as the proportion of successfully executed steps to the total number of steps in each instruction. For example, an instruction with 5 total steps has a completion rate of 60% if 3 of its steps are successfully executed. We compare four configurations of the pipeline: (1) beam search only (*BS*), (2) look-ahead only (*LH*), (3) beam search with look-ahead (*BS+LH*), and finally (4) the baseline condition that uses neither beam search nor look-ahead (*Baseline*).

We used 50 emulators in parallel to execute the 187 instructions for each configuration during the experiments. Thus, there were 4 batches of execution to process all the instructions for the Baseline or LH conditions, and 12 batches for the BS or BS+LH conditions because they both used beam search with a beam size of 3, resulting in 187×3 sequences to execute. This study evidences why it was necessary to support parallel execution. On average, an execution of one instruction took 8 minutes. Even for the Baseline condition,

⁴PixelHelp dataset [25]: <https://github.com/google-research-datasets/seq2act>

Table 1: Summary of Actions Executed and Assets Collected by our pipeline on processing the PixelHelp Dataset.

	# Actions Executed	# Words	# Images	Video (seconds)	# UI Elements
Minimum	4	11	4	70.1	64
Maximum	28	48	21	573.8	471
Mean	10.7	18.0	8.2	252.9	173.8
Total	1997	3368	1528	47,296	32,494

Table 2: Effect of Beam Search (BS) and Look-Ahead (LH)

	Baseline	BS	LH	BS + LH
Steps Executed	2.82	3.17	2.94	3.35
Completion Rate	67.1%	76.6%	69.1%	80.9%

sequential execution would result in an execution time of over 24 hours. The additional execution required for implementing beam search and the potential need to support different platform or app version further amplifies the need for parallel processing.

Beam Search Only. Of the 187 generated tutorials, 32 tutorials (17.1%) benefited from the extra beams, where the second or third beams executed more correct steps than the first beam. The first beams in these cases (i.e., the baseline condition) averaged only 1.22 steps executed, while the beam selected by beam search average 3.25 steps executed, successfully executing two more steps on average. With beam search, the average completion rate for all instructions improved from 67.1% of the baseline to 76.6%.

Look-Ahead Only. We found that 10 tutorials (5.3%) were improved by look-ahead (i.e., containing at least one look-ahead step in any of its beams), increasing from an average of 2.4 steps executed of the baseline to 4.6 steps executed. Overall, there was also an increase in completion rate versus baseline (69.1% vs. 67.1%).

Beam Search with Look-Ahead. We found that enabling beam search along with look-ahead improved the effectiveness of look-ahead, with look-ahead being applied to 41 tutorials (21.9%), compared to only 10 with look-ahead alone. The average number of executed steps in these tutorials increase from 1.59 of the baseline to 3.46. Meanwhile, the addition of look-ahead also made beam search more effective, with the number of tutorials benefiting from beam search increasing from 32 for beam search only to 39 (20.9%). Overall, the completion rate rose from 67.1% of the baseline to 80.9% of beam search with look-ahead. The number of successful steps executed increased from 2.82 to 3.35, an 18.8% increase. In total, 63 tutorials (33.7%) benefited from beam search, look-ahead, or both. The existing failure cases mostly occurred due to a substantial mismatch between the instruction text and the app interface caused by app updates. In addition, because the emulator was automatically log into some apps, some actions that require authentication failed.

Summary. Our results demonstrate the synergy of adopting both beam search and look-ahead, showing substantial improvement over the baseline technique. While we leverage and adapt

general techniques, our pipeline is novel for fully-automatic tutorial creation from text instructions. With this infrastructure, we can transform more text instructions to visual tutorials at scale.

5 USER EVALUATION

To investigate how mobile users receive visual tutorials generated by HelpViz in comparison with the original text instructions, we conducted a user study with 10 participants. We used the HelpViz tutorials generated from the PixelHelp dataset (see Section 4) for the experiments.

5.1 Participants

We recruited 10 participants from an IT company (2 female, 8 male). Each participant received a gift code worth \$15 US dollars for their completion of the study. All the participants had 5-year experience using a smartphone at minimum. Four participants rated their knowledge of using an Android phone as moderate and that they occasionally seek help; the remaining six participants rated themselves as experts and seldom seek help. While it would be preferable if the participants had little or no prior experience with Android, under the challenges of the COVID-19 pandemic when the research was conducted, each participant needed to use their own device in a remote setting. This limited us to recruiting participants who owned an Android phone, which meant that they were likely to have some experience with Android. Although these users may not rely on tutorials for day to day tasks, they still use tutorials when encountering unfamiliar or complex tasks.

5.2 Apparatus

For this remote study, we asked the participants to run the HelpViz study app on their personal Android phones. Participants used Android phones with a variety of versions, which six of the phones were running Android v10, two on Android v9, one on Android v8, and one on Android v6. Three participants used Google Pixel phones, two used Samsung Galaxy phones, one used a Motorola Moto phone, one used a OnePlus phone, while the others did not disclose their phone models. Of the seven phones that we know the model name, four had a screen size between 5.5 and 5.9 inches, and three were between 6.0 and 6.4 inches. We do not expect the variations in phone models and screen sizes to affect our study, and none of the participants reported issues regarding device compatibility.

We randomly selected 22 tutorials generated from the PixelHelp dataset, generated with beam search and look-ahead in Section 4. Each task involved navigating to a certain page in an Android app.

The apps used in the tasks consisted of four apps, Settings (8 tasks), Clock (2 tasks), Google Photos (4 tasks), and Google Chrome (6 tasks). The tasks required between three and seven steps to complete. The study app logged all user interaction with time stamps during the study. When users are seeking help from a tutorial, they usually do not have a good idea of the steps to take. Therefore, we did not include titles for the tutorials to reduce the likelihood of a participant familiar with the system guessing the steps without looking at the tutorials.

5.3 Experimental Design and Procedure

We employed a within-subjects design for this study, where the only factor was Technique (HelpViz Tutorial or Text-Only). We presented 11 tasks for each technique, with the first one used for warm-up. The presentation order of the two techniques was counterbalanced between participants. The assignment of tasks to each technique was also randomized across participants. We counterbalanced the tasks across techniques such that a task was presented the same number of times in each technique.

Our experiment consisted of three parts. First, we introduced the tasks to the participants and asked them to install the study app on their phones. Then, we asked the participants to complete a warm-up task and 10 tasks for each technique tested. We informed the participants to switch between the tutorial and the target app whenever necessary for completing a task. The participants were allowed to take a break after each task. Finally, we asked them to fill in a post-study survey about their experiences during the study and collected demographic information. Specifically, we asked them to what extent they agreed or disagreed with the following statements for each technique on the five-point Likert items below (i.e., each on a scale of 1 to 5). We referred to their short names (in parentheses) in the Results section. Each participant's subjective ratings for the six items were summed to create an *overall scale* on a range between 6 and 30 for each technique [2].

- (1) "It was easy to understand the instructions." (*Understand*)
- (2) "I could easily find the corresponding action." (*Find action*)
- (3) "I was able to complete the tasks without problems." (*Complete task*)
- (4) "It was faster to use tutorials like this to complete a task." (*Speed*)
- (5) "I was satisfied with tutorials like this." (*Satisfaction*)
- (6) "When finding help online, I would prefer to use tutorials like this." (*Preference*)

We also collected background information, including the number of years participants had used a smartphone, daily smartphone usage in hours, and their knowledge of using an Android phone. Finally, we discussed with the participants on their impressions and experience while using both types of tutorials.

5.3.1 Analysis procedure. In order to test for potential differences in *Completion Time* (i.e., the time elapsed between trial start to reaching the target screen) and the *Number of App Switches* (i.e., the number of times a participant switched between the tutorial and the target app during a trial) for each technique, we performed a Mixed Model analysis of variance on each of the dependent variables, modeling *Technique* as an independent factor and *ParticipantID* as a random effect.

In order to test for potential difference in participants' subjective ratings, we conducted a paired-samples *t*-test on the overall Likert scale to compare the techniques. We also compared the ratings for each technique using the Wilcoxon signed-rank test. Finally, we summarized the participants' subjective feedback.

5.4 Experimental Results

The 10 participants completed a total of 200 trials. Due to the nature of remote studies where participants performed the tasks on their own devices in their own environment, we needed to remove 63 trials. Of these, 49 trials were excluded due to the mismatch of user devices or apps ("android fragmentation"). For example, participants' devices had photos apps from a different vendors, or Samsung's Settings app instead of Pixel's. Inherently, these were different apps instead of different versions of an app. These differences required vastly different tutorials, which could be easily addressed by letting HelpViz generate tutorials corresponding to these devices from instructions specific to these apps, or better screening participants for the remote study. The rest of the cases were due to exceptions in the study that are not related to HelpViz, e.g., apps on participants' devices restarting during the task. All these cases could have been easily avoided in a lab study where a designated device is used, which was not possible during the pandemic. This resulted in a total of 137 trials used for analysis.

5.4.1 Performance. We did not find statistically significant difference in the *Completion Time* ($F_{1,129} = 2.42, p \approx .12$) and the *Number of App Switches* ($F_{1,129} = 1.22, p \approx .27$) for *Technique*, i.e., HelpViz versus Text-Only. The mean *Completion Time* was 33 seconds for HelpViz, and 30 seconds for the Text-Only instructions. The mean *Number of App Switches* was 2.62 times for HelpViz, and 2.34 times for text instructions. To further test whether there was a difference between the two techniques, we also included the *Number of Steps* in each task as another independent factor. Mixed Model ANOVA again showed no significant effect for *Technique*, nor its interaction with the *Number of Steps*, although the *Number of Steps* itself were significant for both *Completion Time* and the *Number of App Switches*, which was expected.

5.4.2 Subjective feedback. We summed subjective scores on each of the Likert items to create an overall scale for each participant, as described in Section 5.3. While there was no significance found on completion time and app switch frequency, participants overall showed preference for HelpViz ($mean = 23.5, sd = 1.96$) over text instructions ($mean = 20.5, sd = 2.72$), $t(9) = 3.10, p \approx .01$. For individual items, a Wilcoxon signed-rank test shows that HelpViz scored significantly higher on the *Preference* item ($Z = 2.55, p \approx .02$, effect size $r = .80$). As shown in Fig. 3, participants rated HelpViz favorably compared to Text-Only help in general. The overall preference assessment received more positive ratings than other subjective feedback items. This is likely because the *Preference* item captures the participants' overall assessment about each technique, which is contributed by all aspects of these techniques. Interestingly, although the actual task *Completion Time* of HelpViz (33 seconds) is on par with that of the Text-Only condition (30 seconds), participants feel HelpViz allow them to finish a task faster, as indicated by the *Speed* preference in Fig. 3. We note that the

Figure 3: The participants' preferences for each technique.

task Completion Time and participants' perception of Speed both showed no statistically significant difference. The faster perceived speed could indicate that the assistance provided by HelpViz made them feel that the tasks took less time.

During the study, several participants expressed their surprise and delight when they found out that the HelpViz tutorial automatically scrolled to their current step (P1, P2, P8). P4 also preferred auto-scroll. None of the participants reported issue with auto-scroll. P2 and P8 praised the quality of the screenshots and auto-generated tutorials. Participants also commented on the benefits of screenshots. 7 out of the 10 participants found the screenshots with UI elements to be helpful. In particular, P2 felt that screenshots make the tutorial easy to follow, and P4, P5, and P8 shared the opinion that screenshots provided expectations for which element to interact with in each step. In addition, 7 out of the 10 participants found the thumbnail to be helpful. P7 favored the screenshots in providing context by showing the neighboring elements and P5 liked seeing an element's general location within a screen. P5 also mentioned that "images stuck in memory". P2 and P5 compared the experience of using tutorials with screenshots with text instructions, and stated that text instructions felt harder.

Participants found the ability to fit all instructions into one page for the text instructions was helpful in completing the tasks (9 out of the 10 participants). P7 preferred text for shorter instructions. P6 preferred text instructions in general, claiming that screenshots might not be so valuable for "power users".

Participants provided insightful feedback and suggestions to improve HelpViz. P1, P3, P8, and P10 all reported that it's less desirable to use the screenshots when they do not match the app UI, and would like to see screenshots created for their versions of apps. P6 and P9 pointed out that in some cases, such as opening a pop-up menu, it would make more sense to combine menu opening with the next step in the tutorial, because a menu might disappear if they switched out of the app.

6 DISCUSSION & OPPORTUNITIES

Overall, we received positive feedback from the user study. Our results indicate that while HelpViz does not directly impact the number of operations or total time needed to complete each task, users showed strong preference for HelpViz when compared to text

instructions. The screenshots in the tutorials made them easier to follow by allowing the users to have an expectation for the target element and its surrounding elements. Although it is widely known that images are useful in tutorials, previous work shed little light on how visual tutorials that are automatically generated from a batch process are useful, which does not involve any human intervention and manual refinements. Our evaluation contributes new insights that automatically synthesized visual tutorials are preferred by participants, which shows the feasibility of the approach.

In addition, auto-scroll and highlighting of the current step reduced the amount of interaction needed while navigating the tutorial and locating the instruction. Although HelpViz does not seem to significantly affect task execution time in our study, it is preferred by users as the screenshot and auto-scroll features received substantial praise. We hypothesize that as the complexity of the tasks increase, HelpViz can show advantages on time performance as well, which deserves further investigation. Below we describe the limitations of our techniques and identify opportunities that can be readily made to HelpViz.

Optimizing Tutorials for Specific Devices & Granularity.

Participants expressed the desire for more accurate tutorials based on the version of the app on their device. Such tutorials can be easily generated by executing them in emulators running different versions of the Android system and third-party apps. As the entire process is automated, this would require minimal human intervention. Furthermore, certain actions can be grouped together to reduce the length of the tutorial. For example, a menu open action followed by a tap on the menu can be consolidated into one step, which can be achieved by checking the class name and text labels for the menu. Lastly, although none of the participants reported any issue with auto-scroll, screen matching accuracy could be improved by adopting more robust heuristics, such as utilizing other features from the Accessibility API [24] and adopting the screen equivalence heuristics proposed in [37].

Supporting Diverse User Groups. To facilitate power users who prefer text instructions over the images [3], HelpViz can present screenshots on demand, or only be displayed for the next step. We suggest that the context-aware capability of HelpViz tutorials will still be valuable in either cases, guiding users to the next step. On the other hand, HelpViz can possibly provide voice guidance from text instructions to support mobile users in need of accessibility support [33].

Generalizing to More Complex Workflows. Because our tutorial generation pipeline is designed to handle raw, unconstrained text instructions, e.g., those directly found on the web, a variety of complex workflows may arise, such as conditional branches and loops. However, these complexities are yet investigated because our experiments focused on PixelHelp, a set of instructions curated from a single website, which contained mostly linear instructions. HelpViz was also constrained by the underlying parsing model, which operated linearly through the instruction text. HelpViz overcomes some of the constraints by supporting branches partially. As typically the branches in a tutorial would transition to different actions for different conditions encountered, the look-ahead mechanism in HelpViz would enable the emulator to skip actions in branches that are not present or invalid. HelpViz does not support loops and would only execute each action in a loop once. These

limitations can be addressed by detecting conditionals and loops in instructions in the future.

There are two directions of future work for HelpViz to successfully handle conditional branches: (1) parsing and traversing all possible branches during emulator execution, and (2) matching current encountered state during runtime with recorded states during execution. To solve (1), a more advanced parsing model would be necessary, as well as an infrastructure to set the appropriate emulator environment according to the different branching options in the instruction. To solve (2), HelpViz needs to keep track of the current state during emulator execution, and record the executed steps along with their corresponding states. During runtime, the recorded states can be matched to the states that a user encounters, and tutorials can be shown accordingly. These conditions propose new challenges for the underlying parsing model, and requires additional handling in our pipeline to accommodate these complexities.

Since online tutorials can be complex and diverse, there might be further unexpected situations where the pipeline fails to yield a valid tutorial. Should this happen, as a fallback, our pipeline combines the visual assets from steps that have been successfully executed with the text of the remaining steps.

HelpViz tutorials could be further expanded to desktop workflows, which calls for parsing models tailored to desktop instructions, emulation infrastructure for desktop software, and expanding user action tracking capabilities to a more dynamic environment, where user focus can easily switch between multiple applications.

7 CONCLUSION

HelpViz automates tutorial generation for Android apps by automatically executing a batch of text-based instructions on a cluster of Android emulators, and extracting steps, screenshots, videos snippets, and other metadata that are needed for creating a visual tutorial. By incorporating a previously trained deep model [25], HelpViz grounds natural language instructions to actions in the emulator—a dynamic environment, and improves the robustness of fully automatic execution by enabling beam search execution and look-ahead. A user study with 10 participants indicates that participants reacted positively to HelpViz, particularly when compared to traditional text-only tutorials. HelpViz demonstrates the feasibility to convert a vast number of text instructions on the web to visual tutorials that are adapted to different versions of mobile apps and platforms and hardware configurations.

ACKNOWLEDGMENTS

We would like to thank anonymous reviewers for their insightful feedback for improving the paper. We also want to thank our user study participants.

REFERENCES

- [1] Olivier Bau and Wendy E. Mackay. 2008. OctoPocus: A Dynamic Guide for Learning Gesture-Based Command Sets. In *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology* (Monterey, CA, USA) (UIST '08). Association for Computing Machinery, New York, NY, USA, 37–46. <https://doi.org/10.1145/1449715.1449724>
- [2] Harry N Boone and Deborah A Boone. 2012. Analyzing likert data. *Journal of extension* 50, 2 (2012), 1–5.
- [3] Pei-Yu Chi, Sally Ahn, Amanda Ren, Mira Dontcheva, Wilmot Li, and Björn Hartmann. 2012. MixT: Automatic Generation of Step-by-Step Mixed Media Tutorials. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology* (Cambridge, Massachusetts, USA) (UIST '12). Association for Computing Machinery, New York, NY, USA, 93–102. <https://doi.org/10.1145/2380116.2380130>
- [4] Andy Cockburn, Amy Karlson, and Benjamin B. Bederson. 2009. A Review of Overview+detail, Zooming, and Focus+context Interfaces. *ACM Comput. Surv.* 41, 1, Article 2 (Jan. 2009), 31 pages. <https://doi.org/10.1145/1456650.1456652>
- [5] Biplab Deka, Zifeng Huang, Chad Franzen, Joshua Hibschan, Daniel Afergan, Yang Li, Jeffrey Nichols, and Ranjitha Kumar. 2017. Rico: A Mobile App Dataset for Building Data-Driven Design Applications. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology* (Québec City, QC, Canada) (UIST '17). Association for Computing Machinery, New York, NY, USA, 845–854. <https://doi.org/10.1145/3126594.3126651>
- [6] Biplab Deka, Zifeng Huang, and Ranjitha Kumar. 2016. ERICA: Interaction Mining Mobile Apps. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (Tokyo, Japan) (UIST '16). Association for Computing Machinery, New York, NY, USA, 767–776. <https://doi.org/10.1145/2984511.2984581>
- [7] Daniel Eckhoff, Christian Sandor, Denis Kalkoten, Ulrich Eck, Christian Lins, and Andreas Hein. 2018. TutAR: Semi-Automatic Generation of Augmented Reality Tutorials for Medical Education. In *2018 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. 430–431. <https://doi.org/10.1109/ISMAR-Adjunct.2018.00131>
- [8] Jerry Alan Fails and Dan R. Olsen. 2003. Interactive Machine Learning. In *Proceedings of the 8th International Conference on Intelligent User Interfaces* (Miami, Florida, USA) (IUI '03). Association for Computing Machinery, New York, NY, USA, 39–45. <https://doi.org/10.1145/604045.604056>
- [9] James Fogarty, Desney Tan, Ashish Kapoor, and Simon Winder. 2008. CueFlick: Interactive Concept Learning in Image Search. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy) (CHI '08). Association for Computing Machinery, New York, NY, USA, 29–38. <https://doi.org/10.1145/1357054.1357061>
- [10] Adam Fourney, Ben Lafreniere, Parmit Chilana, and Michael Terry. 2014. InterTwine: Creating Interapplication Information Scent to Support Coordinated Use of Software. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (Honolulu, Hawaii, USA) (UIST '14). Association for Computing Machinery, New York, NY, USA, 429–438. <https://doi.org/10.1145/2642918.2647420>
- [11] C. Ailie Fraser, Tricia J. Ngoon, Mira Dontcheva, and Scott Klemmer. 2019. Re-Play: Contextually Presenting Learning Videos Across Software Applications. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3290605.3300527>
- [12] Mitchell Gordon and Philip J. Guo. 2015. Codepourri: Creating visual coding tutorials using a volunteer crowd of learners. In *2015 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. 13–21. <https://doi.org/10.1109/VLHCC.2015.7357193>
- [13] Floraine Grabler, Maneesh Agrawala, Wilmot Li, Mira Dontcheva, and Takeo Igarashi. 2009. Generating Photo Manipulation Tutorials by Demonstration. *ACM Trans. Graph.* 28, 3, Article 66 (July 2009), 9 pages. <https://doi.org/10.1145/1531326.1531372>
- [14] Tovi Grossman and George Fitzmaurice. 2010. ToolClips: An Investigation of Contextual Video Assistance for Functionality Understanding. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1515–1524. <https://doi.org/10.1145/1753326.1753552>
- [15] Andrew Head, Jason Jiang, James Smith, Marti A. Hearst, and Björn Hartmann. 2020. Composing Flexibly-Organized Step-by-Step Tutorials from Linked Source Code, Snippets, and Outputs. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3313831.3376798>
- [16] Paul Jaccard. 1912. The distribution of the flora in the alpine zone. 1. *New phytologist* 11, 2 (1912), 37–50.
- [17] Caitlin Kelleher and Randy Pausch. 2005. Stencils-Based Tutorials: Design and Evaluation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 541–550. <https://doi.org/10.1145/1054972.1055047>
- [18] Md Adnan Alam Khan, Volodymyr Dziubak, and Andrea Bunt. 2015. Exploring Personalized Command Recommendations Based on Information Found in Web Documentation. In *Proceedings of the 20th International Conference on Intelligent User Interfaces* (Atlanta, Georgia, USA) (IUI '15). Association for Computing Machinery, New York, NY, USA, 225–235. <https://doi.org/10.1145/2678025.2701387>
- [19] Juho Kim, Philip J. Guo, Carrie J. Cai, Shang-Wen (Daniel) Li, Krzysztof Z. Gajos, and Robert C. Miller. 2014. Data-Driven Interaction Techniques for Improving Navigation of Educational Videos. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (Honolulu, Hawaii, USA) (UIST '14). Association for Computing Machinery, New York, NY, USA, 563–572. <https://doi.org/10.1145/2642918.2647389>
- [20] Juho Kim, Phu Tran Nguyen, Sarah Weir, Philip J. Guo, Robert C. Miller, and Krzysztof Z. Gajos. 2014. Crowdsourcing Step-by-Step Information Extraction

- to Enhance Existing How-to Videos. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) (CHI '14). Association for Computing Machinery, New York, NY, USA, 4017–4026. <https://doi.org/10.1145/2556288.2556986>
- [21] Nicholas Kong, Tovi Grossman, Björn Hartmann, Maneesh Agrawala, and George Fitzmaurice. 2012. Delta: A Tool for Representing and Comparing Workflows. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) (CHI '12). Association for Computing Machinery, New York, NY, USA, 1027–1036. <https://doi.org/10.1145/2207676.2208549>
- [22] Mackenzie Leake, Hijung Valentina Shin, Joy O. Kim, and Maneesh Agrawala. 2020. Generating Audio-Visual Slideshows from Text Articles Using Word Concreteness. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–11. <https://doi.org/10.1145/3313831.3376519>
- [23] Guo Li, Tun Lu, Jiang Yang, Xiaomu Zhou, Xianghua Ding, and Ning Gu. 2015. Intelligently Creating Contextual Tutorials for GUI Applications. In *2015 IEEE 12th Intl Conf on Ubiquitous Intelligence and Computing and 2015 IEEE 12th Intl Conf on Autonomic and Trusted Computing and 2015 IEEE 15th Intl Conf on Scalable Computing and Communications and Its Associated Workshops (UIC-ATC-ScalCom)*. 187–196. <https://doi.org/10.1109/UIC-ATC-ScalCom-CBDCom-IoP.2015.50>
- [24] Toby Jia-Jun Li, Amos Azaria, and Brad A. Myers. 2017. SUGILITE: Creating Multimodal Smartphone Automation by Demonstration. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 6038–6049. <https://doi.org/10.1145/3025453.3025483>
- [25] Yang Li, Jiacong He, Xin Zhou, Yuan Zhang, and Jason Baldridge. 2020. Mapping Natural Language Instructions to Mobile UI Action Sequences. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 8198–8210. <https://doi.org/10.18653/v1/2020.acl-main.729>
- [26] Justin Matejka, Wei Li, Tovi Grossman, and George Fitzmaurice. 2009. CommunityCommands: Command Recommendations for Software Applications. In *Proceedings of the 22nd Annual ACM Symposium on User Interface Software and Technology* (Victoria, BC, Canada) (UIST '09). Association for Computing Machinery, New York, NY, USA, 193–202. <https://doi.org/10.1145/1622176.1622214>
- [27] Clara Meister, Tim Vieira, and Ryan Cotterell. 2021. If beam search is the answer, what was the question? arXiv:2010.02650 [cs.CL]
- [28] Alok Mysore and Philip J. Guo. 2017. Torta: Generating Mixed-Media GUI and Command-Line App Tutorials Using Operating-System-Wide Activity Tracing. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology* (UIST '17). Association for Computing Machinery, New York, NY, USA, 703–714. <https://doi.org/10.1145/3126594.3126628>
- [29] Suporn Pongnumkul, Mira Dontcheva, Wilmot Li, Jue Wang, Lubomir Bourdev, Shai Avidan, and Michael F. Cohen. 2011. Pause-and-Play: Automatically Linking Screencast Video Tutorials with Applications. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology* (Santa Barbara, California, USA) (UIST '11). Association for Computing Machinery, New York, NY, USA, 135–144. <https://doi.org/10.1145/2047196.2047213>
- [30] Philip Quinn and Shumin Zhai. 2016. A Cost-Benefit Study of Text Entry Suggestion Interaction. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 83–88. <https://doi.org/10.1145/2858036.2858305>
- [31] Vidya Ramesh, Charlie Hsu, Maneesh Agrawala, and Björn Hartmann. 2011. ShowMeHow: Translating User Interface Instructions between Applications. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology* (Santa Barbara, California, USA) (UIST '11). Association for Computing Machinery, New York, NY, USA, 127–134. <https://doi.org/10.1145/2047196.2047212>
- [32] André Rodrigues, Leonardo Camacho, Hugo Nicolau, Kyle Montague, and Tiago Guerreiro. 2018. Aidme: Interactive Non-Visual Smartphone Tutorials. In *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct* (Barcelona, Spain) (MobileHCI '18). Association for Computing Machinery, New York, NY, USA, 205–212. <https://doi.org/10.1145/3236112.3236141>
- [33] André Rodrigues, André Santos, Kyle Montague, Hugo Nicolau, and Tiago Guerreiro. 2019. Understanding the Authoring and Playthrough of Nonvisual Smartphone Tutorials. In *Human-Computer Interaction – INTERACT 2019*, David Lamas, Fernando Loizides, Lennart Nacke, Helen Petrie, Marco Winckler, and Panayiotis Zaphiris (Eds.). Springer International Publishing, Cham, 42–62.
- [34] Anh Truong, Peggy Chi, David Salesin, Irfan Essa, and Maneesh Agrawala. 2021. Automatic Generation of Two-Level Hierarchical Tutorials from Instructional Makeup Videos. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 108, 16 pages. <https://doi.org/10.1145/3411764.3445721>
- [35] Cheng-Yao Wang, Wei-Chen Chu, Hou-Ren Chen, Chun-Yen Hsu, and Mike Y. Chen. 2014. EverTutor: Automatically Creating Interactive Guided Tutorials on Smartphones by User Demonstration. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) (CHI '14). Association for Computing Machinery, New York, NY, USA, 4027–4036. <https://doi.org/10.1145/2556288.2557407>
- [36] Xu Wang, Benjamin Lafreniere, and Tovi Grossman. 2018. Leveraging Community-Generated Videos and Command Logs to Classify and Recommend Software Workflows. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3173859>
- [37] Xiaoyi Zhang, Anne Spencer Ross, and James Fogarty. 2018. Robust Annotation of Mobile Application Interfaces in Methods for Accessibility Repair and Enhancement. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology* (Berlin, Germany) (UIST '18). Association for Computing Machinery, New York, NY, USA, 609–621. <https://doi.org/10.1145/3242587.3242616>