PONI: A Personalized Onboarding Interface for Getting Inspiration and Learning About AR/VR Creation

Narges Ashtari
nashtari@sfu.ca
School of Computing Science, Simon Fraser University
Burnaby, BC, Canada

Parsa Alamzadeh
palamzad@sfu.ca
School of Computing Science, Simon Fraser University
Burnaby, BC, Canada

Gayatri Ganapathy
gganapath@sfu.ca
School of Computing Science, Simon Fraser University
Burnaby, BC, Canada

Parmit K Chilana
pchilana@sfu.ca
School of Computing Science, Simon Fraser University
Burnaby, BC, Canada

ABSTRACT

New creators of augmented reality (AR) and virtual reality (VR) applications often face a steep learning curve during the onboarding stage of creation and struggle in identifying suitable learning materials that are appropriate for their skillsets. To support the initial learning needs of new AR/VR creators from different backgrounds, we designed and implemented a novel personalized onboarding interface (PONI) that allows users to locate relevant projects based on their programming and 3D modeling skills, development goals, and any constraints, such as time or budget. Our usability evaluation (n=16) showed that most creators found PONI to be intuitive, useful, and saw its potential to be used as a knowledge hub for inspiration and self-directed exploratory learning. We discuss ways in which the personalization could be further enhanced and how the potential of PONI could be explored to improve onboarding in contexts beyond AR/VR development.

KEYWORDS

personalization, software learnability, AR/VR authoring, end-user development

ACM Reference Format:
Narges Ashtari, Parsa Alamzadeh, Gayatri Ganapathy, and Parmit K Chilana. 2022. PONI: A Personalized Onboarding Interface for Getting Inspiration and Learning About AR/VR Creation. In *Nordic Human-Computer Interaction Conference (NordiCHI '22), October 8–12, 2022, Aarhus, Denmark.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3546155.3546642

1 INTRODUCTION

As consumer-level augmented reality (AR) and virtual reality (VR) devices are getting cheaper and easier to access around the world, there has been growing interest in creating new types of AR/VR applications. This growth has led to a proliferation of different AR/VR authoring tools and development environments. For example, commercial frameworks such as *Unity*, *Unreal*, and *A-Frame* allow developers to create industry-level games and other types of creative AR/VR experiences. There is also another class of emerging tools with simpler user interfaces that aim to lower the barrier-to-entry for AR/VR development (e.g., Vizor.io [3], CenarioVR [1], and Cospaces.io [2]). Research in HCI is also pushing the boundary of AR/VR prototyping by exploring tools that eliminate the need for programming or 3D modeling to make AR/VR creation easier to access (e.g., [57, 60, 65–67, 84, 96]).

Despite the availability of several options for authoring AR/VR applications, getting started with AR/VR development still presents a steep learning curve for newcomers to the field [4, 8, 83]. Newcomers often need to first explore existing AR/VR projects and understand the possibilities for design and programming. We refer to this as the *onboarding stage* of the AR/VR creation process as it consists of preparatory activities that newcomers do before actually tinkering with any of the authoring frameworks or development environments. This onboarding stage can be particularly problematic for the growing community of AR/VR creators who come from a range of different domains and may have limited or no professional training in software development, design or engineering [4]. These creators can include artists who are exploring AR to showcase their creations in art installations [88], teachers who are tinkering with VR in their classes to convey complex ideas like geometry [9, 41], and architects who are trying to create virtual physicalizations of their designs [9, 53]. These diverse non-professional AR/VR creators are often not familiar with relevant terminology and concepts and find it difficult to understand the full landscape of AR/VR techniques [4].

To locate examples and learning materials, most new AR/VR creators currently start their informal onboarding process by initiating a web search and peruse through resources such as MOOCs, YouTube videos, and online forums [4]. However, since these creators are not familiar with the AR/VR nomenclature and do not fully understand the interplay between different hardware and software components, they have difficulty in formulating their queries and expressing their desired goals (e.g, should they be choosing *marker-based* or *marker-less* AR?). Another key problem for these creators is assessing the suitability and reliability of the retrieved materials relative to their own skills in programming, 3D modeling, or other technologies [4, 47, 77]. For example, a newcomer may not

NordiCHI '22, October 8–12, 2022, Aarhus, Denmark

ACM ISBN 978-1-4503-9699-8/22/10…$15.00
https://doi.org/10.1145/3546155.3546642

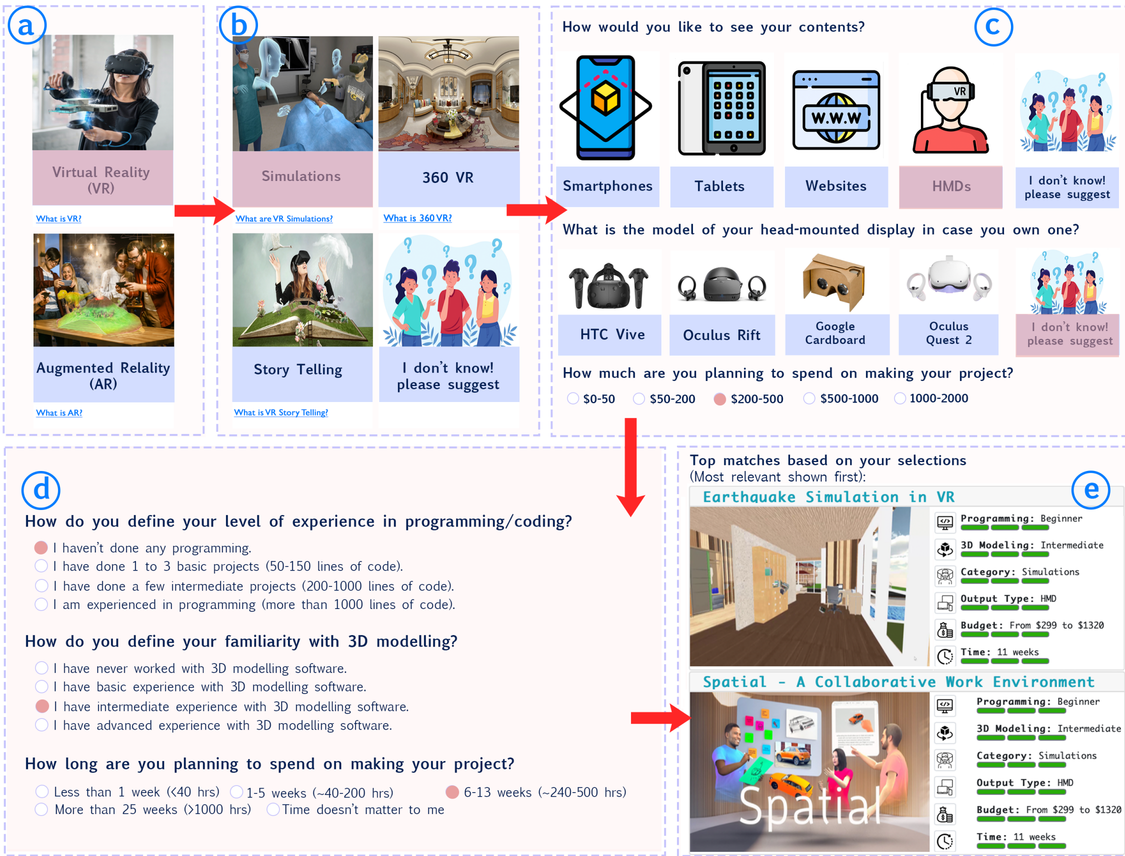

Figure 1: An example of a user journey in PONI that simulates the scenario described in Figure 6: a) The user first determines the type of immersive experience they want (AR or VR); b) Based on the type of experience, the user selects the category of experience (e.g., in VR, they can select simulations, 360° VR, story telling, or leave it open by indicating I don't know; c) The user then specifies their intended way for experiencing the output (e.g., for VR, the options include smartphones, tablets, Web/desktop, and HMDs) and any budget constraints or targets; d) The user specifies their background in terms of technical skills, such as programming and 3D modeling; e) Based on the answers, PONI generates a user profile and directs the user to the suggestion module which shows a ranked list of projects matching the user profile.

realize that the tutorial that they are looking at requires advanced knowledge of 3D geometry or skills in adapting a particular API. Furthermore, creators may be working within the constraints of a timeline or a specific budget and the examples or tutorials that they find online may not be possible to recreate within these constraints. As a result, creators can get entangled in inefficient trial-and-error processes as they look for relevant examples and guidance [47, 76].

Given the difficulties that new AR/VR creators face in understanding the landscape of different design possibilities and determining the suitability of tutorials for their own needs, we wondered how we could use a *personalized* approach to support these creators' onboarding process. In particular, our research question is: *How can we design a* ***personalized onboarding*** *tool for helping new AR/VR creators retrieve learning materials that are appropriate for the creators' level of technical skills, desired goals, and a given set of constraints?*

In this paper, we present the design and evaluation of a "**P**ersonalized **ON**boarding **I**nterface" (*PONI*), that facilitates early stage AR/VR creation for newcomers. PONI is an interactive tool that uses examples and simplified descriptions to incrementally introduce the nomenclature and stages of the AR/VR creation process. It uses a rule-based approach [33, 79] to generate a user profile that captures the user's technical skills, development or design goals, and any constraints, such as time or budget. Based on the user profile, PONI retrieves a ranked list of projects tailored to user needs and characteristics from a database of curated AR/VR examples. To accommodate user navigation and support recognition over recall [70], PONI provides visual cues to show the extent to which each user input matches the retrieved project's characteristics. Users can further customize the suggestions by defining the importance of the factors impacting the ordering of the results or by applying filters on the suggested projects.

To evaluate PONI, we ran an observational usability study with 16 AR/VR newcomers and compared PONI with another non-personalized keyword-based BASELINE interface. We found that almost all of the participants found PONI to be more intuitive, useful, and engaging compared to the BASELINE. In particular, participants indicated that PONI served as a useful centralized hub for learning about AR/VR terminologies and requirements and to get inspiration for potential projects that were actually feasible given one's skillsets. A key advantage of PONI for participants was that they could engage more in systematic self-directed exploratory learning instead of relying on trial-and-error.

The main contributions of this paper are: (1) the design and implementation of PONI, a novel personalized onboarding interface that allows new AR/VR creators to discover learning materials tailored to their skill levels, goals, and constraints; and, (2) insights from an observational usability study that demonstrate the utility of the personalized onboarding approach for newcomers and how it could be used for self-directed exploratory learning. Our findings confirm that one-size-fits-all approaches do not work well for the differing needs of AR/VR creators and that personalization techniques could provide a fruitful starting point for supporting nuances in onboarding. Overall, our work highlights the importance of adopting a user-centred interaction design perspective for designing personalized systems in the context of supporting informal learning.

2 RELATED WORK

This work builds upon prior research on challenges in getting started with AR/VR creation, innovations in software learnability, and personalization approaches used in formal learning.

2.1 Challenges in creating AR/VR applications

New AR/VR creators can face a number of different challenges in getting started as they need to understand the capabilities of various platforms, tools, and devices and determine how they work together to create a cohesive AR/VR experience. Currently available creation frameworks can vary widely in terms of system structure and hardware constraints to support intended use-cases and relevant features [68, 83]. In most cases, experiences created with one framework only run on one kind of device, and repurposing it to another framework or adjusting it to support more devices is either complicated or expensive. For example, marker-based applications are created entirely differently than ones that work with spatial mapping. While a newcomer may start with a quick web search, identifying an appropriate tool-chain requires experience and domain knowledge [4]. This even makes it harder for newcomers to understand the strengths and weaknesses of different frameworks, understand different creation processes, and figure out ways to best combine available resources to satisfy application requirements [68].

One approach for lowering the barriers to entry is automating some of the technical aspects of AR/VR creation, such as generating the initial prototypes [7, 45, 55, 56, 66, 78]. However, recent work indicates that such approaches may not work for AR/VR creators who have little to no software development experience and lack a conceptual model of the overall creation process. For example, Ashtari et al. [4] found that non-professional AR/VR creators, such as hobbyists and domain experts, face difficulty in knowing where to even start, lack access to concrete design guidelines and examples, and struggle in making use of online learning resources. Other research on AR/VR creation [8, 83] has also identified the struggles that creators face with the fragmented landscape of AR/VR authoring tools and how newcomers often fail to select appropriate programming languages, authoring tools, or testing hardware that meets their project-specific needs. Most newcomers try to draw inspiration from existing example AR/VR projects and use them to jump-start their design [4], but struggle in finding learning resources that contain an appropriate amount of high-level (e.g., general rules and strategies of the AR/VR creation process) and low-level details (e.g., software, hardware, and devices used for a particular AR/VR experience) [4, 68]. This makes it difficult to determine the feasibility of a given project that matches an individual's needs and constraints.

Considering the challenges that new AR/VR creators face during onboarding, in this research we attempt to lower the barriers of entry through the design of an interface that 1) personalizes newcomers' initial learning experiences considering their background and constraints, 2) helps them gain domain knowledge through exploration of example projects, and 3) helps them assess the suitability of learning materials.

2.2 Innovations in software help and tutorial systems

Although consumer-level AR/VR creation has only recently started receiving attention in HCI, there is a long history of research on software learnability and supporting help-seeking activities. Since beginners are known to struggle in locating relevant learning materials [47], some research advocates embedding the relevant help in the form of tutorials and Q&A within the target application through overlays and other in-context techniques [17, 27, 61, 62, 89]. Researchers have also explored techniques for improving interaction with video-based tutorials (e.g., [6, 28, 42, 49, 52, 72, 73]), which tend to be more popular way of learning about using a feature-rich application [47]. Other approaches have tried to lower the learning cognitive load by adding gamification elements [58] or augmenting tutorials with input from the user community [12, 54, 74]. Although the general concepts in this software learning and help systems can be applied to specific AR/VR authoring tools, prior work shows that newcomers face challenges in even knowing what tools to select in the first place and their first onboarding need is understanding the overall landscape of AR/VR development.

Individual differences in training and technical expertise also play a significant role when a user starts learning a new programming language or works with a new feature-rich application [22, 24]. Some works have explored ways of detecting software expertise [34] to support users coming from different backgrounds and differing in skill levels, often using low-level operations such as pauses or dwells [73], time of access to the menu [40], mouse motions [32], and usage heatmaps [86]. While these approaches can be effective after a considerable amount of user interaction with the system, they suffer from cold start problem and cannot provide much advantages to the user without exposure to actual user profiles or activities. Another class of tools has explored ways of assessing a tutorial's difficulty by using machine learning techniques to automatically assess a tutorial's difficulty (e.g., [77]) or by using social voting mechanisms to classify difficulty level of instructions [89]. But, these approaches only consider one type of application-specific expertise and, in practice, most software activities span multiple applications. To accomplish this type of activity learners need to equip themselves with a "tool-belt" [85] often differing in characteristics, commands, and output, but we are only starting to see some work in HCI exploring application-independent learning support (e.g, *RePlay* [29]).

In summary, most existing works only provide targeted help for specific tasks that are performed within a single system without acquiring a deep understanding of the user needs, characteristics, and target project. In contrast to the existing approaches in learning and help-seeking, PONI presents a novel design that focuses on personalized onboarding. PONI applies a rule-based [36] method to provide an opportunity for newcomers to declare their own backgrounds and intents. Users can see an overview of the chain of tools used in the creation process of various AR/VR projects activities that is personalized for their needs and level of experience in programming and 3D modeling. Furthermore, PONI personalizes tool suggestions based on user's access to devices and budget constraints.

2.3 Personalization in formal learning

Personalization in education and learning has a long history [10, 43, 44]. The goal of personalization in formal learning is to adopt student-centered practices and design intervening mechanisms to help instructors better individualize learning strategies. They take into account differences in students' skill levels, needs, and interests, and assist learners to *succeed* at a task [63, 75]. By drawing on this method during the delivery of the curriculum, instructors can allow students to move at more individualized paces, assign customized assignments to assess each student's mastery, devise a path that is customized to address each student's needs at the moment, and cluster students strategically [20]. When personalized support is provided [95], the student knows which steps to take and how to proceed independently. When support is non-tailored to students' understanding, students often withdraw from the task as it is beyond or beneath their reach causing frustration or boredom. One large-scale study [87] of a personalized learning classroom intervention program showed that students' achievement (measured with a multiple-choice test and a knowledge assignment) increased with high levels of personalized support compared to limited levels of personalizations in student advising. But, how these interventions could be designed beyond a classroom or formal learning setting is an open research question.

Our work takes inspiration from personalized learning approaches in classroom settings and explores personalization for informal learning practices, such as looking up technical tutorials and examples. Personalization requires information about the user, whether the data are explicitly gathered by asking people to fill out forms (e.g., rule-based personalization [36]) or implicitly through analysis of behavioral data (e.g., data-driven personalization [13, 31] or collaborative filtering (e.g., [80]). The latter techniques require historical data and digital traces of user behavior to tailor the learning materials to user needs. In designing PONI, since we did not have prior access to user profiles or digital traces related to onboarding, we could not use data-driven and collaborative filtering approaches for recommending relevant content. We instead used a rule-based personalization technique to build user profiles that take into account a user's background, skills, and constraints to offer them tailored onboarding content.

3 DESIGN CONSIDERATIONS AND GOALS

In this paper, we explore the design of a personalized onboarding interface that helps newcomers get familiar with the landscape of AR/VR technologies and terminologies, and allows them to retrieve learning materials relevant to their interests, skills, and constraints. Based on the related work discussed above, we considered different aspects of designing and structuring the personalized onboarding process for new AR/VR creators and derived five design goals:

DG1: Locate targeted learning materials given the creator's background and skills. Prior work in learning research shows that people build new knowledge based on what they already know and believe [26, 48, 91] and personal variables (motivations, goals, and self-efficacy) may be predictors of engagement in a development activity. Onboarding approaches for AR/VR creation should take into account learners' prior knowledge (e.g., programming and 3D modelling skills) and should be flexible enough to adapt to learner differences [4, 68].

DG2: Locate targeted learning materials given the desired creation outcomes and constraints. Creating AR/VR experiences requires working knowledge across a chain of software and hardware tools. For example, designing a VR 360° experience for YouTube could consist of an initial 3D prototype in Blender with mock ups and animations added in Unity, and use of the YouTube video player for testing, or optimization for a more immersive experience in an Oculus headset. However, newcomers often fail to find examples and tutorials that cover this entire process [4, 68]. Moreover, newcomers usually have different constraints based on their allocated budget and time commitment [4]. Onboarding for AR/VR creators should allow creators to locate a personalized set of learning materials that are appropriate and feasible for their desired creation outcomes and constraints.

DG3: Get inspiration and browse relevant example projects. Although a plethora of learning resources are available online (MOOCs, YouTube videos, tutorials, forums), they are laden with device-specific or application-specific instructions [4, 68]. However, newcomers often do not even know where to begin as they not understand the landscape of possibilities and are less familiar with the vocabulary used in tutorials [4, 30]. Onboarding for AR/VR creation should facilitate the early stages of exploration by offering examples that creators can use to see design possibilities at a high-level and simplify the descriptions of needed components.

DG4: Assess relevance of learning materials in relation to the creator's background and desired goals. A key challenge for newcomers often is recognizing relevant projects and tutorials from a list of search results as they lack a mental model of the underlying application [47]. Onboarding techniques should use visualization to provide an intuitive "at a glance" explanation for suggested

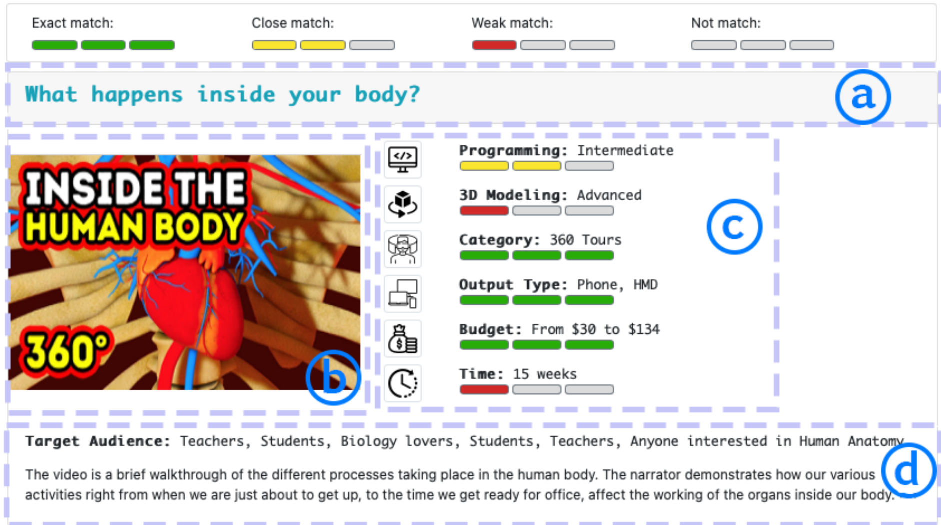

Figure 2: An instance of a project card that constitutes a single result retrieved by PONI. It has three main parts: (a) a header showing the project's title, (b) a preview image from the project, and (c) color-coded bars showing to what extent each project matches the user inputs (e.g., programming, 3D modeling, category, output, budget, time), and (d) a brief project description and the target audience.

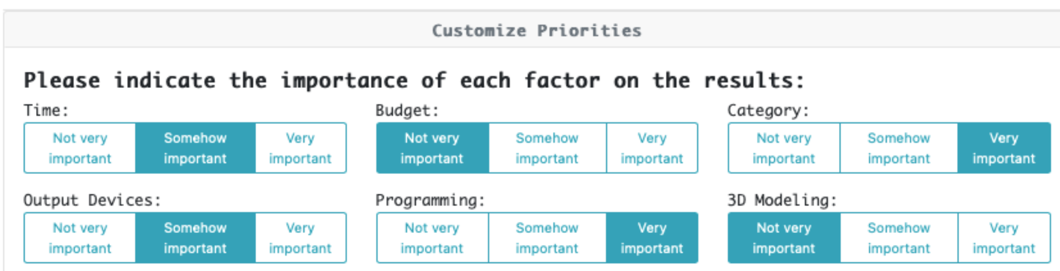

Figure 3: The results page contains customization options for users to define the impact of each user input on showing the ordered list of matching projects. The initial importance (explained in 4.4.3) of each factor is highlighted in blue and can be updated by the user. When the user changes the importance of any factor, the order of the results update accordingly to revised user preferences.

learning materials. For example, highlighting matching metadata when presenting recommended content has been shown to be a useful technique [37]. Moreover, features like adding contextual cues to search results can provide *information scent* [25, 27] that can help users more quickly and easily navigate the results.

DG5: Offer freedom and flexibility in personalizing and exploring relevant learning materials. A known drawback of personalization and profile-based recommender systems is that the underlying algorithms can be opaque to the end user, especially if the system is not open to user inspection or modification [38]. Instead, onboarding should be both *adaptable* (e.g., allow for manual configuration by the user), as well as *adaptive* (e.g., provide proactive personalizations to satisfy the needs of the user) [71]. If the adaptation logic is defined in the form of rules, users can gain control over the system by being able to inspect, understand, and modify the underlying adaptation model. This is particularly important for supporting the wide range of AR/VR creation possibilities and diversity among new creators.

4 PONI: SYSTEM DESIGN AND IMPLEMENTATION

Based on the above design considerations, we followed an iterative design approach consisting of rounds of sketching and wireframing, and elicitation of user feedback [98]. We designed PONI, a novel *personalized onboarding interface* to help AR/VR newcomers with diverse backgrounds, skill levels, and range of development goals to locate relevant learning materials. The design splits across three main modes of interaction: (1) *the input module* where the users specify their background, skills, desired outcomes, and constraints;

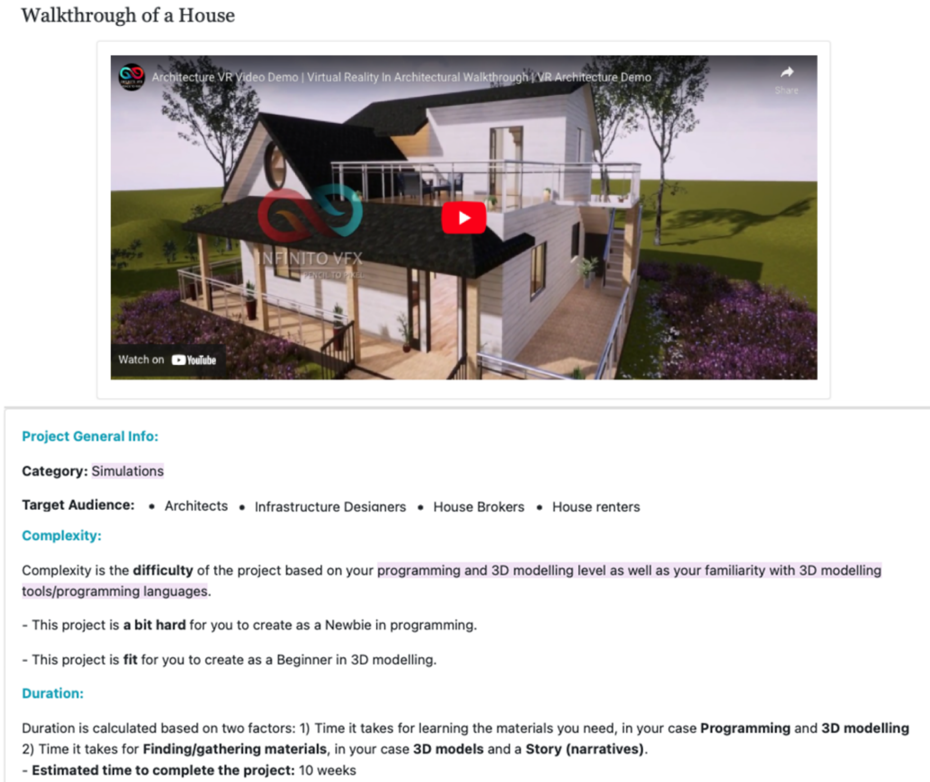

Figure 4: An example simplified project page that provides minimal technical terms and specifications, allowing a novice user to quickly determine the suitability of the project in light of their own background, skills, and personal goals.

(2) *the suggestion module*, which suggests relevant learning materials based on user input and allows users to assess the suitability of each resource; and, (3) the *project description module* allows users to see details of each retrieved result and assess their relevance using metadata such as the required hardware, authoring tools, and programming languages, among others.

4.1 Input Module: Defining user characteristics and desired outcomes

Since PONI is a new design concept, we did not have access to users' information in advance. For providing personalized onboarding experience, we decided to use a rule-based approach [33, 79] that focused on incrementally asking for various user characteristics, preferences, and constraints related to AR/VR development.

4.1.1 Determining the type and category of immersive experience. Fulfilling DG2, the input module enables users to identify their target project by learning about the general concept and vocabulary used in creating immersive experiences. PONI prompts users to choose the type of immersive experience (Figure 1.a) among two available options (*Augmented Reality* or *Virtual Reality*). We arrived at using AR and VR as the primary representative categories of immersive experiences and excluded Mixed Reality (MR) from available options due to the ambiguity and disagreement in the definition of MR in the literature [82]. Next, PONI prompts users to specify the category of experience (Figure 1.b) within AR or VR technologies [50, 82] by adopting the clustering introduced in existing approaches [50, 82, 97]. This step helps users further specify the type of desired immersive experience (e.g., a marker-based vs a location-based AR experience or a simulation vs. 360° VR experience) by browsing examples (more details can be found in Appendix A).

4.1.2 Specifying the target tools and type of outcome. To address DG2, PONI also prompts users to specify their intended ways of experiencing the output. Following DG3, all tools are shown with images to provide learning opportunities and to support recognition over recall (Figure 1.c) [70]. Moreover, users with budget constraints are provided with an option to specify a budget range for further customization and user control. PONI provides an "*I don't know; Please assist*" choice to accommodate the decision-making process and facilitate learning through exploration if users have no particular preferences or are unsure what to choose.

4.1.3 Determining user skill sets and constraints. Considering differences in technical skills and motivations of the creator base of AR and VR experiences (DG1), PONI prompts users to self-define their programming experience, 3D modeling familiarity, and approximate the time they want to spend on their desired AR/VR project (Figure 1.d).

4.2 Suggestion Module

4.2.1 Locating relevant learning materials. In keeping with DG4, PONI's results page (Figure 1.e) lists projects in ranked descending order of match relevance (the matching algorithm is detailed in 4.4.3). Each example project is introduced by a card containing a representative image and information about the factors such as programming level, 3D modeling level, category, output type, and estimated completion time (Figure 2). The suitability of each factor is portrayed through four colours from green being exact, yellow being close, red being weak, and grey being no match to reduce users' cognitive load and support recognition over recall. By clicking on each project card, users gain access to in-depth details of a given project's creation details (see 4.3).

4.2.2 Customizing suggestions. Applying DG5, PONI offers user control by allowing users to manually customize the importance of factors influencing the ordering of suggested projects (Figure 3). To simplify the interpretation of the numeric schema used in the background algorithm (see 4.4.3), PONI presents the importance of each factor using a descriptive schema (e.i., not very important, somewhat important, and very important).

4.2.3 Filtering suggestions. By default, PONI presents all matching projects ordered from most to least relevant. However, addressing DG5, users can also apply hard filters and reduce the size of their search space (e.g., only show projects related to simulations). PONI applies these filters over the system's initial suggestions (see Appendix B).

4.3 Project Description Module

Following DG1 and DG4, each project page is personalized and designed to offer an appropriate level of information based on different user characteristics. In particular, there are two types of project pages:

1) *simplified project pages* (Figure 4) that target non-technical creators and minimize jargon to allow users to quickly see whether or not the project fits their needs.

2) *intermediate project pages* that target creators with intermediate to advanced skills in 3D modelling and programming. This page provides more detailed information about the recommended and optional programming languages, frameworks, 3D modeling tools, output devices, and the hardware requirements for creating a similar project. If a newcomer to programming sought more details of a given project, PONI directs them this page template.

4.4 Implementation

PONI is a platform-agnostic web-based application written in HTML, JavaScript, and Python in the Django framework. The goal was to create a proof-of-concept implementation of the personalized onboarding concept with real-world data that could be evaluated with real users. Once a user submits their initial preferences questionnaire (described in 4.1), PONI matches user preferences (described in 4.4.3) against its back-end curated database of AR/VR projects (explained in 4.4.2). In the current implementation of PONI, the database consists of a manually-curated collection of over 100 AR and VR example video projects and tutorials (this database can be made available to other researchers upon their request).

4.4.1 Gathering example projects and tutorials. Given the popularity of video tutorials in self-directed learning practices [47], we populated PONI's database with English-speaking YouTube videos. We aimed for diversity in the AR/VR technology used, level of complexity in 3D modeling and programming, authoring software, and project topics (see Appendix C). We ensured that all videos showcase at least one project and excluded general explainer videos on AR and VR.

4.4.2 Curating gathered projects. Two of the authors with prior experience in AR/VR creation independently annotated videos gathered from YouTube. To construct a consistent annotation schema, these authors looked at prior work on roles of user expertise interacting with user interfaces and feature-rich software [34] and challenges of non-professional AR/VR creators [4, 68]. As a final schema (see Appendix D), the authors annotated each project based on the type and category of experience within AR/VR, the difficulty level of 3D modeling and programming of each project, estimated budget (including physical equipment and general software licences) and time to complete a project given level of difficulty, recommended and optional creation platforms, 3D modeling software, hardware, and relevant keywords used in the creation of the project. The annotations also included a general project description and possible target audiences of each project. The authors had two rounds of discussions to ensure the consistency of annotations. In the first round, researchers looked for similarities and differences in the curation to reach agreement in annotations. In the second round, after completing the curation on both sides, an inter-rater reliability test was applied to ensure annotation consistency, achieving a Kappa score of *0.81.*

4.4.3 Presenting ranked list of projects. Based on the responses collected through the input module (see 4.1), PONI's internal algorithm starts to match user inputs against its curated database (see 4.4.2). The ranked list of matches consists of individual project cards (see 4.2.1). We defined a scoring function to sort and present a ranked list of matching projects using two criteria: *closeness* of match between user input and the corresponding curated project attributes, and the *importance* of each of the attributes (e.g., programming, 3D modeling, output, budget, time). To visually demonstrate the extent to which each attribute matches user input, PONI renders the closeness factor using the color schema introduced in 4.2.1 (Figure 2).

Closeness of Match: PONI considers the extent to which the user responses match the curated project attributes (e.g., No Match, Weak Match, Close Match, and Exact Match). PONI compares eight attributes of each project's annotations and the user's answers to: type, time, budge, category of experience, head mounted display, output devices, 3D modeling experience, and programming experience questions (see Appendix E for calculation details of the closeness of match).

Importance of Attributes: PONI determines how each of the above factors are important for ordering (e.g., Low, Medium, and High). The importance of each factor can also be customized by the user on the results page (see Appendix E for calculation details of the importance of attributes).

5 USABILITY STUDY

To evaluate the extent to which the personalized onboarding concept introduced in PONI helps AR/VR creators find examples and tutorials, we compared PONI to a non-personalized keyword-based retrieval approach which resembles the "status-quo" of finding online materials. We ran a usability study with 16 users who were new to AR/VR creation and assessed their perceptions of usability (ease of use), utility (ease of locating relevant materials), and engagement (feel of control, confidence, and system demand) using quantitative and qualitative methods.

5.1 Baseline Interface Used for Comparison

To evaluate how the personalized learning materials suggested by PONI are perceived by users, we implemented another retrieval interface (which we will call BASELINE henceforth) that incorporated the same database of AR/VR projects used in PONI but did not personalize the retrieval. This interface (Figure 5) provided a keyword-based query interface and matched user queries against the curated metadata used in PONI. The formatting of the search results displayed in BASELINE was similar to PONI in that users could see each project's metadata through project cards. However, the key difference was that BASELINE did not provide any visual cues indicating to what extent each factor matches user inputs. Since no personalization was provided within this condition, all users saw the default project description page that provided full access to all of the metadata.

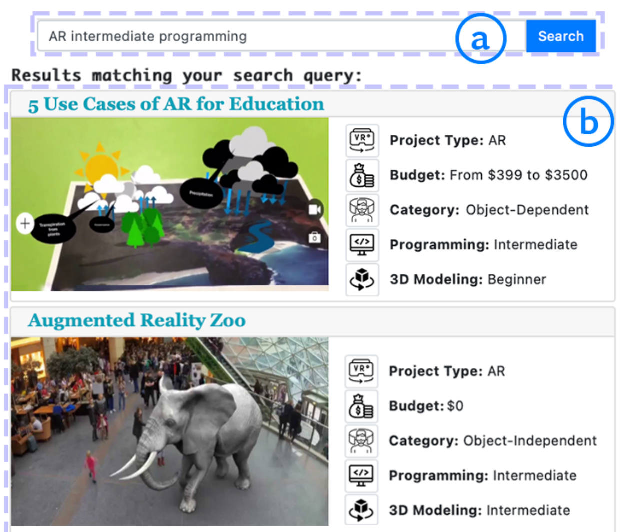

Figure 5: The BASELINE interface consists of a search bar (a) where users issue queries using keywords. The list of results shown only match the user's search query and there is no personalization. Similar to PONI, (b) the projects are listed in separate cards indicating metadata. But, unlike PONI, BASELINE does not provide any visual cues indicating to what extent each factor matches user inputs.

5.2 Participants

We focused on recruiting participants who were new to AR/VR and excluded any experienced AR/VR developers so that we could study their perceptions of the onboarding process and reduce any prior learning effects. To obtain a broad overview of AR/VR creation practices, we recruited a diverse pool of participants (9F/7M) from different backgrounds (CS, Architecture, Arts, UI/UX design, Industrial Design). Our participants were all between the ages of 18-34 and had different levels of education (7 Bachelor's, 6 Master's, and 3 PhDs). All participants reported their initial interest in AR/VR creation. The majority (12/16) of participants were completely new to AR/VR creation while a few (4/16) had tried to create an AR/VR experience in the past but still self-identified as novices. We recruited these participants mainly from university mailing lists and personal connections in the local community.

5.3 Study design and tasks

We used a within-subject design to minimize the impact of high variation among participants. Participants completed four tasks in total under two conditions (two using BASELINE and two using PONI) where each task asked them to select learning materials using one of the interfaces. Participants were asked to locate at least three relevant learning resources for a given AR/VR creation scenario and creator persona (an example is shown in Figure 6. We defined personas and tasks based on documented experiences of newcomers to AR/VR [4, 50].

5.4 Study Procedure and Measures

We conducted the study both in-person and remotely through Zoom, and participants received a $15 Amazon gift card for their time. Participants were asked to log in to our web portal with pre-assigned credentials. This portal allowed them to access both test systems (PONI and BASELINE). To minimize user bias, each system was assigned a pseudonym (e.g., Green for PONI and Blue for the BASELINE condition). Next, participants filled out a pre-test questionnaire (via *SurveyMonkey*) that captured demographics and information about prior experiences in AR/VR creation learning (e.g., familiarity with AR/VR creation software and platforms, general experience in finding learning materials for AR/VR creation). We presented each of the tasks one by one in random order. After completing each task, users filled out a post-task questionnaire (via *SurveyMonkey*) to assess their perception of usability (ease of use), utility (ease of locating relevant materials), engagement (feeling of control, confidence, and system demand), and system reuse. Lastly, we carried out follow-up interviews to further probe into the strengths and weaknesses of each onboarding system design. Sessions were video and audio-recorded for transcription, and the participants were asked to share their screens through Zoom during the usability test. The usability test and follow-up interview took approximately one hour.

5.5 Data Analysis

We used a combination of statistical tests and an inductive analysis approach [19] to analyze the study data. We ran Wilcoxon signed-rank tests with the nominal variable "system" (having two levels: PONI and BASELINE) and ordinal variable "agreement" having five

Task Description
Armstrong school is planning to run a competition with the goal of exposing students to new tech trends. As a starter, the school is running a VR competition. Students are encouraged to be creative and try to build a Simulation experience through which they teach other students any topic they find interesting. The competition is going to happen in 12 weeks. Imagine you are volunteering to encourage students to enter this competition and create a VR simulation experience. A student, Emily, has been assigned to you who has never done any programming and has intermediate 3D modeling experience. To encourage her to continue this path, her parents have allocated a $600 budget for her to buy a head-mounted display device but they don't know what their options are. Your job is to help Emily find example tutorials and projects for developing a VR Simulation application. Please find 3 example projects that you think are most relevant for learning about creating VR simulation projects reflecting Emily's time and budget constraints and her programming and 3D modeling background.

Figure 6: One example of the scenarios used in the user study.

levels (Strongly Disagree, Disagree, Neutral, Agree, and Strongly Agree) to quantitatively determine the significance of the results. We analyzed the participant's qualitative feedback using affinity diagrams and explored themes around users' perceptions of the personalized onboarding concept. Through discussions with two members of our research team, we categorized our findings and identified key recurring themes related to usability, utility, engagement, and areas of improvements.

5.6 Results

Overall most of our participants preferred PONI (15/16) over the BASELINE condition for locating relevant AR/VR learning materials. We next present users' perception of usability, utility, and engagement as they interacted with the two different conditions in our study.

5.6.1 ***Usability.*** Users found PONI (Mdn = 4) to be easier to use than BASELINE (Mdn = 3.5) and this difference in perceived ease of use was significant ($z = -3.14$, $p < 0.002$). Participants found the question and answer, step-by-step style of PONI to be "light" and "easy to manage" as it provided them with an "intuitive" starting point (even when these creators had limited or no understanding of what AR/VR projects they should look at). In contrast, the open-ended nature of keyword search in BASELINE felt like a "black box" and left them "uncertain" and "clueless" about what needs to be prioritized and what a meaningful query should even look like:

"*In [BASELINE] it got a little bit tricky to know which keyword I need to search for...because, I can go through 10 results or change the keywords 6-7 times, but because I had to check the suitability of them all, it was a bit time-consuming.*" (P5)

5.6.2 ***Utility.*** Overall PONI (Mdn = 4) was ranked higher than BASELINE (Mdn = 3) in terms of ease of locating relevant materials and this difference was significant ($z = -3.39$, $p < 0.001$)). Participants explained that the simplified terms and examples used during input served as "clues" and helped them with understanding the AR/VR nomenclature:

"*It [PONI] is just intuitive;...I appreciated the question and answer method; it held my hand a bit more. Also, the features of filtering and prioritizing kinda gave you some clues about the metrics and terms you need to look for.*" (P1)

In addition, participants indicated that PONI could be a useful "hub" for figuring out what the requirements should be for a given project and what current standards are before they even touch an authoring tool:

"*It [PONI] helped me know what I need to be considering. Because in the beginning, I do not have a set standard in my mind...I did not have to type anything, and it meant that I did not have to figure out the words.*" (P1)

While our quantitative analysis showed that there was significant overall user preference for PONI, some participants who were confident in their programming skills (known as *Hobbyists* in [4]), visual cues and automatic ordering of the results were not enough for locating relevant materials. These participants were ambitious with their learning goals, and found that some of the retrieved results were perhaps "too easy", "limiting" or "not challenging enough".

5.6.3 ***User engagement.*** Participants indicated that PONI (Mdn = 4) provided significantly higher control ($z = -3.13$, $p < 0.002$) compared to the BASELINE condition (Mdn = 3). Users found PONI helpful in making informed decisions as it allowed them to choose their priorities and maintain control over their preferences:

"*...With this one [PONI] I am able to judge better whether those results fit my criteria or if they do not fit how off or different they are from my criteria.*" (P3)

Another aspect of PONI that provided more control to participants was the centralized way of presenting information:

"*...when I was taking my class, I had little knowledge of AR/VR... I had to read Reddit forums or Google Poly and post on them because that was the only way I understood the terms and processes but still did not get a full picture.*" (P9)

While PONI was helpful in anchoring participants to examples that fit their skills and constraints, we observed that some users wanted to also find relevant materials based on topics of interest (e.g., all classroom-related AR examples). The current version of PONI was not able to provide this low-level of control over topic so that users could cluster projects that are thematically similar.

PONI (Mdn = 4) was also ranked significantly higher than the BASELINE (Mdn = 3) condition in helping participants feel confident about their selections ($z = -3.03$, $p < 0.002$). A recurring sentiment expressed by participants was that the simple language used in PONI helped them take the appropriate next step and gave them more confidence about the criteria and keywords they should look for. With BASELINE, all participants reported frustration in formulating and re-formulating search queries from scratch as they struggled to assess the relevance of the results.

Participants ranked interacting with PONI (Mdn = 1) significantly less demanding than the BASELINE (Mdn = 3) condition ($z = -3.49$, $p < 0.001$). They also were more likely to reuse PONI (Mdn = 4) than the BASELINE (Mdn = 3.5) tool in the future ($z = -3.47$, $p < 0.001$). Participants noted that the visual cues provided in PONI made it easier to peruse through search results and explore learning resource alternatives. With BASELINE, users had to spend extra time and effort to determine relevance, usually with little success:

"*I found it [BASELINE] cognitively demanding because I needed to memorize so many different things at the same time. I was trying to make it easier by opening multiple tabs so I did not forget the ones I liked...I felt the need to take a look at all of them, but at the same time, I could not evaluate all of them in time.*" (P7)

5.7 Areas of improvement

Although PONI was perceived to be significantly better than the BASELINE condition for all of our measures, we did synthesize some potential areas of improvement for personalized onboarding based on the user feedback. First, some participants in the interviews reported a slight preference for using a combination of both the keyword-based and a rule-based personalized approach. These participants indicated that they could initially benefit from having PONI's exploratory approach for understanding a project's requirements, keywords, and feasibility, but would like to be able to search by keywords as they gained more experience. Some participants also shared a need for a side-by-side comparison of projects and a more in-depth analysis of how suggested projects map their individual profile. While PONI was helpful for participants in locating and recognizing suitable learning materials, there was no easy way for them to compare the pros and cons of a group of similarly relevant projects. Users also indicated that they would feel more confident about trying out a project if they could see how their skills (e.g., in 3D modeling or programming) map to the different steps of a project's creation process.

6 DISCUSSION

6.1 Key Takeaways

We have contributed the design and evaluation of a novel personalized onboarding interface (PONI) that helps AR/VR newcomers locate relevant learning materials tailored to their programming and 3D modeling skills, development goals, and constraints, such as time or budget. Our initial evaluation indicates that for AR/VR newcomers, the personalized onboarding process offered by PONI can be more intuitive, useful, and engaging compared to exploratory keyword-based search methods. Our work complements ongoing efforts in AR/VR authoring (e.g., [60, 66, 67]) to lower the onboarding barriers for creators from different backgrounds by assisting them in understanding the landscape of AR/VR creation and getting inspiration for projects.

Although our focus in this paper was on AR/VR creation, we believe that our approach of facilitating personalized onboarding can be generalized to other complex design tasks, such as 3D modeling [46, 47], and software development tasks [18, 21, 51, 59, 92] where there is need to understand the landscape and relevant terminologies before beginning the complex authoring process. Our approach complements other innovations in personalized systems for learning and information-seeking by bringing in a human-centered lens and an interaction design approach for tackling the problem. For example, most of the research on personalized information retrieval has focused on the optimization of the underlying algorithms [15, 16, 64], and it is rare to see explicit focus on the interaction design of these systems that captures users' perceptions. Our research shows the importance of observing people using such systems and capturing their perceptions of the effectiveness and utility of personalized results. Insights from this study reveal that there is more in play than just the effectiveness of a retrieval algorithm and factors such as the UI design, user control, transparency, and flexibility all impact users' overall impressions.

We now discuss some limitations of our current work and highlight promising directions for future work in HCI to further expand the design space of personalized onboarding.

6.2 Limitations

Although our implementation of PONI as a proof-of-concept web-based application was useful for assessing users' initial reactions and perceptions of personalized onboarding, more research is needed to fully understand how users would interact with such tools in their actual learning tasks (for example, through a longitudinal field study). Given the scope of this research, the factors that we considered for curating existing projects in our database may not be exhaustive and it is important to keep understanding and addressing the evolving needs of AR/VR creators. However, given our current design and implementation, a natural extension for expanding the curated database would be through the use of methods such as crowdsourcing which have been successful in other learning contexts [94]. Lastly, our current scoring function used for ranking the results only considers three levels for proximity and importance, and future work can explore more granular levels to improve the distinctive power and accuracy of the retrieval in more complex onboarding scenarios.

6.3 Expanding Personalized Approaches

In our current implementation of PONI, we made use of a rule-based approach to retrieve and suggest learning materials based on user input. While this approach has been shown to be accurate in controlled scenarios, it can be blind to user context and may offer low flexibility in supporting spontaneous user interactions [33, 79] (e.g., new items or users that do not fit in any pre-defined clustering). In addition to rule-based approaches, more flexible forms of personalized recommendations can be explored that leverage information from user interaction. For example, collaborative filtering using ratings or other forms of user-generated feedback [11, 39, 81] can determine preference commonalities between groups of users and generate recommendations based on inter-user similarities, which can be particularly useful for supporting the diverse needs of AR/VR creators. Furthermore, content-based filtering [5] also can be used to generate finer-grained recommendations based on the history of a particular user's interactions. Our curated, labeled database of projects can also be used to explore more sophisticated automatic approaches using machine learning and similar techniques [77, 93]. But, some caution has to be used as retrieval and recommender systems that rely on automation and predict user behaviour based on current patterns can suffer from the cold-start problem [33, 79] when they initially lack meaningful data for creating user models.

Based on the insights from our study and prior work on hybrid use of rule-based and adaptive recommender systems [14], there could be some benefits in exploring a combination of both approaches. For example, it can be helpful to learn about user goals, skills, constraints, and context through the initial input and the personalized onboarding systems can adapt as learning progresses. It could be interesting to create experiences that support the delivery of short-term contextual recommendations (e.g., what device is appropriate for my end goal?) and higher-level long-term global recommendations (e.g., task flows and different roadmaps for creating a target project). Such systems can also track and progressively monitor newcomers' progress when applying recommended solutions and completing steps on authoring platforms, predict their needs, and recommend learning materials appropriate to their context. Given the impacts of personalization on how a user experiences and gets exposed to a technology, it can also be interesting to consider the effect of the *filter bubble* [69] of personalization which can isolate people from a diversity of viewpoints or content. An interesting challenge for the future research would be the investigation of interplay between learner-directed exploratory and personalized learning methods.

6.4 Supporting Long-Term Engagement

6.4.1 Accommodate high-level and low-level user learning needs. AR/VR creation and other digital creative processes [29, 68] often require working across a chain of different tools. Professional developers may be able to work with tools at the high end of the toolchain (e.g., Unity) since they have the relevant training and experience, but newcomers often end up wandering around and working with a large patchwork of tools or get stuck with sub-optimal solutions [4, 47, 68]. In the current design of PONI, participants perceived every suggestion by the system to be click-through to provide a reliable gateway without a need for user validation. This opens up the design space for new support tools to focus on providing a project road map joint with relevant learning materials.

6.4.2 Support learning through embedded communities and automated approaches. While the current design of PONI was perceived to be effective in showing results matching a user's declared interests and backgrounds, there is opportunity to further expand the richness of the retrieved results. For example, as pointed out by participants, having a way to compare the merits of different results would help them make more informed selections. Embedding in additional comments and shared experiences from other creators within each project or tutorial (as has been explored in some recent work [12, 23, 90]) could further enhance the learning experience for new creators.

6.4.3 Facilitate evolving needs and situational interests. Currently PONI only allows users to declare intents once and offers only basic customization and filtering options (see Figure 3). There is little support for users as their individual needs evolve or their situational interests change (e.g., if they change mind about a particular device) over time. A challenge for future work is to investigate flexible ways to provide long-term support for evolving needs of the user. One approach could be providing multiple roadmaps (similar to [35, 52]) for a given project through which users can flexibly change their pathways (e.g., change their method, devices, completion time) based on their evolving interests.

7 CONCLUSIONS

In this paper, we have introduced the design of PONI, a novel personalized onboarding interface that assists new AR/VR creators in locating learning materials that are tailored to their programming and 3D modeling skills, development goals, and any constraints, such as time or budget. Users found the step-by-step question-and-answer style, color-coded suggestions, and personalized results to be intuitive, useful, and saw PONI's potential as a knowledge hub for inspiration and self-directed exploratory learning. Our findings provide an initial lens into the potential benefits of personalization for onboarding and early stages of learning about AR/VR creation and could be extended to other informal learning in technical domains.

ACKNOWLEDGMENTS

We thank the Natural Sciences and Engineering Research Council of Canada (NSERC) for funding this research. The authors would also like to thank Andrea Bunt, Patrick Dubois, Joanna McGrenere, and Laton Vermette.

REFERENCES

[1] 2022. Cenario VR: Develop VR Content With Ease And Speed. https://www.cenariovr.com/.
[2] 2022. CoSpaces Edu: Kid-friendly 3D creation and coding. https://cospaces.io/edu/.
[3] 2022. Vizor: Create 360 tours, stories, sites and WebVR experiences. https://site.vizor.io/.
[4] Narges Ashtari, Andrea Bunt, Joanna McGrenere, Michael Nebeling, and Parmit K. Chilana. 2020. *Creating Augmented and Virtual Reality Applications: Current Practices, Challenges, and Opportunities.* Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376722
[5] Marko Balabanović and Yoav Shoham. 1997. Fab:Content-Based, Collaborative Recommendation. *Commun. ACM* 40, 3 (1997), 66–72. https://doi.org/10.1145/245108.245124
[6] Nikola Banovic, Tovi Grossman, Justin Matejka, and George Fitzmaurice. 2012. *Waken: Reverse Engineering Usage Information and Interface Structure from Software Videos.* Association for Computing Machinery, New York, NY, USA, 83–92. https://doi.org/10.1145/2380116.2380129
[7] Mark Billinghurst. 2020. Rapid Prototyping for AR/VR Experiences. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI EA '20)*. Association for Computing Machinery, New York, NY, USA, 1–4. https://doi.org/10.1145/3334480.3375065
[8] Mark Billinghurst. 2021. Grand Challenges for Augmented Reality. *Frontiers in Virtual Reality* 2 (03 2021), 578080. https://doi.org/10.3389/frvir.2021.578080
[9] Mark Billinghurst, Adrian Clark, and Gun Lee. 2015. A Survey of Augmented Reality. *Trends in Human-Computer Interaction* 8, 2–3 (mar 2015), 73–272. https://doi.org/10.1561/1100000049
[10] Benjamin S Bloom. 1984. The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational researcher* 13, 6 (1984), 4–16.
[11] John S. Breese, David Heckerman, and Carl Kadie. 1998. Empirical Analysis of Predictive Algorithms for Collaborative Filtering. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence* (Madison, Wisconsin) *(UAI'98)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 43–52.
[12] Andrea Bunt, Patrick Dubois, Ben Lafreniere, Michael A. Terry, and David T. Cormack. 2014. TaggedComments: Promoting and Integrating User Comments in Online Application Tutorials. 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, 4037–4046. https://doi.org/10.1145/2556288.2557118

[13] Robin Burke. 2002. *User Modeling and User-Adapted Interaction* 12, 4 (2002), 331–370. https://doi.org/10.1023/a:1021240730564

[14] Robin Burke. 2007. Hybrid Web Recommender Systems. In *The Adaptive Web*. Springer Berlin Heidelberg, 377–408. https://doi.org/10.1007/978-3-540-72079-9_12

[15] Mark J. Carman, Mark Baillie, and Fabio Crestani. 2008. Tag Data and Personalized Information Retrieval. In *Proceedings of the 2008 ACM Workshop on Search in Social Media* (Napa Valley, California, USA) *(SSM '08)*. Association for Computing Machinery, New York, NY, USA, 27–34. https://doi.org/10.1145/1458583.1458591

[16] Pablo Castells, Miriam Fernández, David Vallet, Phivos Mylonas, and Yannis Avrithis. 2005. Self-Tuning Personalized Information Retrieval in an Ontology-Based Framework. In *Proceedings of the 2005 OTM Confederated International Conference on On the Move to Meaningful Internet Systems* (Agia Napa, Cyprus) *(OTM'05)*. Springer-Verlag, Berlin, Heidelberg, 977–986. https://doi.org/10.1007/11575863_119

[17] Parmit K. Chilana, Amy J Ko, and Jacob O Wobbrock. 2012. LemonAid: selection-based crowdsourced contextual help for web applications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1549–1558.

[18] Parmit K. Chilana, Rishabh Singh, and Philip J. Guo. 2016. Understanding Conversational Programmers: A Perspective from the Software Industry. 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, 1462–1472. https://doi.org/10.1145/2858036.2858323

[19] Juliet M. Corbin and Anselm Strauss. 1990. Grounded theory research: Procedures, canons, and evaluative criteria. *Qualitative Sociology* 13, 1 (1990), 3–21. https://doi.org/10.1007/bf00988593

[20] Diana I. Cordova and Mark R. Lepper. 1996. Intrinsic motivation and the process of learning: Beneficial effects of contextualization, personalization, and choice. *Journal of Educational Psychology* 88, 4 (1996), 715–730. https://doi.org/10.1037/0022-0663.88.4.715

[21] Kathryn Cunningham, Barbara J. Ericson, Rahul Agrawal Bejarano, and Mark Guzdial. 2021. Avoiding the Turing Tarpit: Learning Conversational Programming by Starting from Code's Purpose. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 61, 15 pages. https://doi.org/10.1145/3411764.3445571

[22] Andrew Dillon and Charles Watson. 1996. User Analysis in HCI—the Historical Lessons from Individual Differences Research. *International Journal of Human-Computer Studies* 45, 6 (dec 1996), 619–637. https://doi.org/10.1006/ijhc.1996.0071

[23] Patrick Dubois, Volodymyr Dziubak, and Andrea Bunt. 2017. Tell Me More! Soliciting Reader Contributions to Software Tutorials. In *Proceedings of the 43rd Graphics Interface Conference* (Edmonton, Alberta, Canada) *(GI '17)*. Canadian Human-Computer Communications Society, Waterloo, CAN, 16–23.

[24] Dennis E. Egan. 1988. Individual Differences In Human-Computer Interaction. In *Handbook of Human-Computer Interaction*. Elsevier, 543–568. https://doi.org/10.1016/b978-0-444-70536-5.50029-4

[25] Michael Ekstrand, Wei Li, Tovi Grossman, Justin Matejka, and George Fitzmaurice. 2011. Searching for software learning resources using application context. In *Proceedings of the 24th annual ACM symposium on User interface software and technology - UIST '11*. ACM Press. https://doi.org/10.1145/2047196.2047220

[26] Eileen Faulkenberry and Thomas Faulkenberry. 2006. Constructivism in Mathematics Education: A Historical and Personal Perspective. *The Texas Science Teacher* 35 (01 2006), 17–21.

[27] Adam Fourney, Ben Lafreniere, Parmit K. 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

[28] Adam Fourney, Ben Lafreniere, Richard Mann, and Michael A. Terry. 2012. "Then click ok!": extracting references to interface elements in online documentation. In *Proceedings of the 30th ACM SIGCHI Conference on Human Factors in Computing Systems*. ACM, 35–38. https://doi.org/10.1145/2207676.2207682

[29] C. Ailie Fraser, Tricia J. Ngoon, Mira Dontcheva, and Scott Klemmer. 2019. RePlay: 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

[30] George W. Furnas, Thomas K. Landauer, Louis M. Gomez, and Susan Dumais. 1987. The vocabulary problem in human-system communication. *Commun. ACM* 30, 11 (Nov. 1987), 964–971. https://doi.org/10.1145/32206.32212

[31] Fabio Gasparetti, Carla Limongelli, and Filippo Sciarrone. 2015. Exploiting wikipedia for discovering prerequisite relationships among learning objects. In *2015 International Conference on Information Technology Based Higher Education and Training (ITHET)*. 1–6. https://doi.org/10.1109/ITHET.2015.7218038

[32] Arin Ghazarian and S. Majid Noorhosseini. 2010. Automatic detection of users' skill levels using high-frequency user interface events. *User Modeling and User-Adapted Interaction* 20, 2 (Feb. 2010), 109–146. https://doi.org/10.1007/s11257-010-9073-5

[33] Nathaniel Good, J. Ben Schafer, Joseph A. Konstan, Al Borchers, Badrul Sarwar, Jon Herlocker, and John Riedl. 1999. Combining Collaborative Filtering with Personal Agents for Better Recommendations. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence and the Eleventh Innovative Applications of Artificial Intelligence Conference Innovative Applications of Artificial Intelligence* (Orlando, Florida, USA) *(AAAI '99/IAAI '99)*. American Association for Artificial Intelligence, USA, 439–446.

[34] Tovi Grossman and George Fitzmaurice. 2015. An Investigation of Metrics for the In Situ Detection of Software Expertise. *Human-Computer Interaction* 30, 1 (jan 2015), 64–102. https://doi.org/10.1080/07370024.2014.881668

[35] Tovi Grossman, Justin Matejka, and George Fitzmaurice. 2010. Chronicle: Capture, Exploration, and Playback of Document Workflow Histories. In *Proceedings of the 23nd Annual ACM Symposium on User Interface Software and Technology* (New York, New York, USA) *(UIST '10)*. Association for Computing Machinery, New York, NY, USA, 143–152. https://doi.org/10.1145/1866029.1866054

[36] Frederick Hayes-Roth. 1985. Rule-Based Systems. *Commun. ACM* 28, 9 (september 1985), 921–932. https://doi.org/10.1145/4284.4286

[37] Marti A. Hearst. 2009. *Search User Interfaces*. Cambridge University Press. https://doi.org/10.1017/cbo9781139644082

[38] Jonathan L. Herlocker, Joseph A. Konstan, and John Riedl. 2000. Explaining Collaborative Filtering Recommendations. In *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work* (Philadelphia, Pennsylvania, USA) *(CSCW '00)*. Association for Computing Machinery, New York, NY, USA, 241–250. https://doi.org/10.1145/358916.358995

[39] Will Hill, Larry Stead, Mark Rosenstein, and George Furnas. 1995. Recommending and Evaluating Choices in a Virtual Community of Use. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI '95)*. ACM Press/Addison-Wesley Publishing Co., USA, 194–201. https://doi.org/10.1145/223904.223929

[40] Amy Hurst, Scott E. Hudson, and Jennifer Mankoff. 2007. Dynamic Detection of Novice vs. Skilled Use without a Task Model. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '07)*. Association for Computing Machinery, New York, NY, USA, 271–280. https://doi.org/10.1145/1240624.1240669

[41] Hannes Kaufmann and Dieter Schmalstieg. 2002. Mathematics and Geometry Education with Collaborative Augmented Reality. In *ACM SIGGRAPH 2002 Conference Abstracts and Applications* (San Antonio, Texas) *(SIGGRAPH '02)*. Association for Computing Machinery, New York, NY, USA, 37–41. https://doi.org/10.1145/1242073.1242086

[42] Caitlin Kelleher and Randy Pausch. 2005. Stencils-Based Tutorials: Design and Evaluation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Portland, Oregon, USA) *(CHI '05)*. Association for Computing Machinery, New York, NY, USA, 541–550. https://doi.org/10.1145/1054972.1055047

[43] Fred S. Keller. 1968. Goodbye Teacher. *Journal of Applied Behavior Analysis* (1968), 79–89. https://doi.org/10.1901/jaba.1968.1-79

[44] Fred S. Keller. 1974. Ten Years of Personalized Instruction. *Teaching of Psychology* (Oct. 1974), 4–9. https://doi.org/10.1177/009862837400100102

[45] Annie Kelly, R. Benjamin Shapiro, Jonathan de Halleux, and Thomas Ball. 2018. *ARcadia: A Rapid Prototyping Platform for Real-Time Tangible Interfaces*. Association for Computing Machinery, New York, NY, USA, 1–8. https://doi.org/10.1145/3173574.3173983

[46] Kimia Kiani, Parmit K. Chilana, Andrea Bunt, Tovi Grossman, and George Fitzmaurice. 2020. "I Would Just Ask Someone": Learning Feature-Rich Design Software in the Modern Workplace. In *2020 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. 1–10. https://doi.org/10.1109/VL/HCC50065.2020.9127288

[47] Kimia Kiani, George Cui, Andrea Bunt, Joanna McGrenere, and Parmit K. Chilana. 2019. Beyond "One-Size-Fits-All": Understanding the Diversity in How Software Newcomers Discover and Make Use of Help Resources. 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, 1–14. https://doi.org/10.1145/3290605.3300570

[48] Ada S. Kim and Amy J. Ko. 2017. A Pedagogical Analysis of Online Coding Tutorials. In *Proceedings of the 2017 ACM SIGCSE Technical Symposium on Computer Science Education* (Seattle, Washington, USA) *(SIGCSE '17)*. Association for Computing Machinery, New York, NY, USA, 321–326. https://doi.org/10.1145/3017680.3017728

[49] 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

[50] Kangsoo Kim, Mark Billinghurst, Gerd Bruder, Henry Been-Lirn Duh, and Gregory F. Welch. 2018. Revisiting Trends in Augmented Reality Research: A Review of the 2nd Decade of ISMAR (2008–2017). *IEEE Transactions on Visualization and Computer Graphics* 24, 11 (2018), 2947–2962. https://doi.org/10.1109/TVCG.2018.

2868591
[51] Amy J. Ko, Brad A. Myers, and Htet Htet Aung. 2004. Six Learning Barriers in End-User Programming Systems. In *Proceedings of the 2004 IEEE Symposium on Visual Languages - Human Centric Computing (VLHCC '04)*. IEEE Computer Society, USA, 199–206. https://doi.org/10.1109/VLHCC.2004.47
[52] 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
[53] Peter G. Krogh. 2000. "Interactive Rooms—Augmented Reality in an Architectural Perspective". In *Proceedings of DARE 2000 on Designing Augmented Reality Environments* (Elsinore, Denmark) *(DARE '00)*. Association for Computing Machinery, New York, NY, USA, 135–137. https://doi.org/10.1145/354666.354681
[54] Benjamin Lafreniere, Tovi Grossman, and George Fitzmaurice. 2013. Community Enhanced Tutorials: Improving Tutorials with Multiple Demonstrations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) *(CHI '13)*. Association for Computing Machinery, New York, NY, USA, 1779–1788. https://doi.org/10.1145/2470654.2466235
[55] Felix Lauber, Claudius Böttcher, and Andreas Butz. 2014. PapAR: Paper Prototyping for Augmented Reality. In *Adjunct Proceedings of the 6th International Conference on Automotive User Interfaces and Interactive Vehicular Applications* (Seattle, WA, USA) *(AutomotiveUI '14)*. Association for Computing Machinery, New York, NY, USA, 1–6. https://doi.org/10.1145/2667239.2667271
[56] Gun A. Lee, Gerard J. Kim, and Mark Billinghurst. 2005. Immersive Authoring: What You EXperience Is What You Get (WYXIWYG). *Commun. ACM* 48, 7 (jul 2005), 76–81. https://doi.org/10.1145/1070838.1070840
[57] Gun A. Lee, Claudia Nelles, Mark Billinghurst, and Gerard Jounghyun Kim. 2004. Immersive Authoring of Tangible Augmented Reality Applications. In *Proceedings of the 3rd IEEE/ACM International Symposium on Mixed and Augmented Reality (ISMAR '04)*. IEEE Computer Society, USA, 172–181. https://doi.org/10.1109/ISMAR.2004.34
[58] Wei Li, Tovi Grossman, and George Fitzmaurice. 2012. *GamiCAD: A Gamified Tutorial System for First Time Autocad Users.* Association for Computing Machinery, New York, NY, USA, 103–112. https://doi.org/10.1145/2380116.2380131
[59] Henry Lieberman, Fabio Paternò, and Volker Wulf. 2006. *End User Development (Human-Computer Interaction Series).* Springer-Verlag, Berlin, Heidelberg.
[60] Blair MacIntyre, Maribeth Gandy, Steven Dow, and Jay David Bolter. 2005. DART: A Toolkit for Rapid Design Exploration of Augmented Reality Experiences. *ACM Trans. Graph.* 24, 3 (jul 2005), 932. https://doi.org/10.1145/1073204.1073288
[61] Justin Matejka, Tovi Grossman, and George Fitzmaurice. 2011. Ambient Help. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada) *(CHI '11)*. Association for Computing Machinery, New York, NY, USA, 2751–2760. https://doi.org/10.1145/1978942.1979349
[62] Justin Matejka, Tovi Grossman, and George Fitzmaurice. 2011. IP-QAT: In-Product Questions, Answers, & Tips. 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, 175–184. https://doi.org/10.1145/2047196.2047218
[63] Jonathan F. Mattanah, Michael W. Pratt, Philip A. Cowan, and Carolyn P. Cowan. 2005. Authoritative parenting, parental scaffolding of long-division mathematics, and children's academic competence in fourth grade. *Journal of Applied Developmental Psychology* 26, 1 (Jan. 2005), 85–106. https://doi.org/10.1016/j.appdev.2004.10.007
[64] Phivos Mylonas, David Vallet, Pablo Castells, Miriam FernÁndez, and Yannis Avrithis. 2008. Personalized Information Retrieval Based on Context and Ontological Knowledge. *Knowledge Engineering Review* 23, 1 (mar 2008), 73–100. https://doi.org/10.1017/S0269888907001282
[65] Michael Nebeling, Katy Lewis, Yu-Cheng Chang, Lihan Zhu, Michelle Chung, Piaoyang Wang, and Janet Nebeling. 2020. *XRDirector: A Role-Based Collaborative Immersive Authoring System.* Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3313831.3376637
[66] Michael Nebeling and Katy Madier. 2019. 360proto: Making Interactive Virtual Reality & Augmented Reality Prototypes from Paper. 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/10.1145/3290605.3300826
[67] Michael Nebeling, Janet Nebeling, Ao Yu, and Rob Rumble. 2018. *ProtoAR: Rapid Physical-Digital Prototyping of Mobile Augmented Reality Applications.* Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3173574.3173927
[68] Michael Nebeling and Maximilian Speicher. 2018. The Trouble with Augmented Reality/Virtual Reality Authoring Tools. In *2018 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. 333–337. https://doi.org/10.1109/ISMAR-Adjunct.2018.00098
[69] Tien T. Nguyen, Pik-Mai Hui, F. Maxwell Harper, Loren Terveen, and Joseph A. Konstan. 2014. Exploring the Filter Bubble: The Effect of Using Recommender Systems on Content Diversity. In *Proceedings of the 23rd International Conference on World Wide Web* (Seoul, Korea) *(WWW '14)*. Association for Computing Machinery, New York, NY, USA, 677–686. https://doi.org/10.1145/2566486.2568012
[70] Jakob Nielsen. 1994. Enhancing the Explanatory Power of Usability Heuristics. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, Massachusetts, USA) *(CHI '94)*. Association for Computing Machinery, New York, NY, USA, 152–158. https://doi.org/10.1145/191666.191729
[71] Jakob Nielsen and Uffe Lyngbæk. 1989. Two Field Studies of Hypermedia Usability. In *Hypertext: State of the Art. Papers presented at the Hypertext 2 conference, York, 1989*, Ray McAleese and Catherine Green (Eds.). Intellect Ltd., Oxford, 64–72.
[72] Amy Pavel, Floraine Berthouzoz, and Björn Hartmann. 2013. Browsing and Analyzing the Command-Level Structure of Large Collections of Image Manipulation Tutorials.
[73] Suporn Pongnumkul, Mira Dontcheva, Wilmot Li, Jue Wang, Lubomir D. Bourdev, Shai Avidan, and Michael F. Cohen. 2011. Pause-and-play: automatically linking screencast video tutorials with applications. *Proceedings of the 24th annual ACM symposium on User interface software and technology* (2011).
[74] Luca Ponzanelli, Gabriele Bavota, Andrea Mocci, Massimiliano Di Penta, Rocco Oliveto, Mir Hasan, Barbara Russo, Sonia Haiduc, and Michele Lanza. 2016. Too Long; Didn't Watch! Extracting Relevant Fragments from Software Development Video Tutorials. In *Proceedings of the 38th International Conference on Software Engineering* (Austin, Texas) *(ICSE '16)*. Association for Computing Machinery, New York, NY, USA, 261–272. https://doi.org/10.1145/2884781.2884824
[75] Anna-Katharina Praetorius, Gerlinde Lenske, and Andreas Helmke. 2012. Observer ratings of instructional quality: Do they fulfill what they promise? *Learning and Instruction* (2012), 387–400. https://doi.org/10.1016/j.learninstruc.2012.03.002
[76] John Rieman. 1996. A Field Study of Exploratory Learning Strategies. *ACM Transactions on Computer-Human Interaction* 3, 3 (sep 1996), 189–218. https://doi.org/10.1145/234526.234527
[77] Shahed Anzarus Sabab, Adnan Khan, Parmit K. Chilana, Joanna McGrenere, and Andrea Bunt. 2020. An Automated Approach to Assessing an Application Tutorial's Difficulty. In *2020 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. 1–10. https://doi.org/10.1109/VL/HCC50065.2020.9127271
[78] Akira Sano. 2011. An Application for Creating Full-Scale Augmented Reality Content without 3d Modeling Skills. In *Proceedings of the 2011 ACM Symposium on The Role of Design in UbiComp Research & Practice* (Beijing, China) *(RDURP '11)*. Association for Computing Machinery, New York, NY, USA, 19–24. https://doi.org/10.1145/2030031.2030039
[79] Andrew I. Schein, Alexandrin Popescul, Lyle H. Ungar, and David M. Pennock. 2002. Methods and Metrics for Cold-Start Recommendations. In *Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Tampere, Finland) *(SIGIR '02)*. Association for Computing Machinery, New York, NY, USA, 253–260. https://doi.org/10.1145/564376.564421
[80] Avi Segal, Ziv Katzir, Ya'akov Gal, Guy Shani, and Bracha Shapira. 2014. EduRank: A Collaborative Filtering Approach to Personalization in E-learning. In *EDM*.
[81] Upendra Shardanand and Pattie Maes. 1995. Social Information Filtering: Algorithms for Automating "Word of Mouth". In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI '95)*. ACM Press/Addison-Wesley Publishing Co., USA, 210–217. https://doi.org/10.1145/223904.223931
[82] Maximilian Speicher, Brian D. Hall, and Michael Nebeling. 2019. What is Mixed Reality?. 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, 1–15. https://doi.org/10.1145/3290605.3300767
[83] Maximilian Speicher, Brian D. Hall, Ao Yu, Bowen Zhang, Haihua Zhang, Janet Nebeling, and Michael Nebeling. 2018. XD-AR: Challenges and Opportunities in Cross-Device Augmented Reality Application Development. *Proceedings of the ACM on Human-Computer Interaction* 2, EICS, Article 7 (jun 2018), 24 pages. https://doi.org/10.1145/3229089
[84] Maximilian Speicher and Michael Nebeling. 2018. *GestureWiz: A Human-Powered Gesture Design Environment for User Interface Prototypes.* Association for Computing Machinery, New York, NY, USA, 1–11. https://doi.org/10.1145/3173574.3173681
[85] Tamara Sumner and Markus Stolze. 1997. *Evolution, Not Revolution: Participatory Design in the Toolbelt Era.* MIT Press, Cambridge, MA, USA, 1–26.
[86] Dereck Toker, Ben Steichen, Matthew Gingerich, Cristina Conati, and Giuseppe Carenini. 2014. Towards Facilitating User Skill Acquisition: Identifying Untrained Visualization Users through Eye Tracking. In *Proceedings of the 19th International Conference on Intelligent User Interfaces* (Haifa, Israel) *(IUI '14)*. Association for Computing Machinery, New York, NY, USA, 105–114. https://doi.org/10.1145/2557500.2557524
[87] Janneke van de Pol, Monique Volman, Frans Oort, and Jos Beishuizen. 2015. The effects of scaffolding in the classroom: support contingency and student independent working time in relation to student achievement, task effort and appreciation of support. *Instructional Science* 43, 5 (June 2015), 615–641. https://doi.org/10.1007/s11251-015-9351-z
[88] Wim van Eck and Yolande Kolstee. 2012. The augmented painting: Playful interaction with multi-spectral images. In *2012 IEEE International Symposium*

on Mixed and Augmented Reality - Arts, Media, and Humanities (ISMAR-AMH). 65–69. https://doi.org/10.1109/ISMAR-AMH.2012.6483990

[89] Laton Vermette, Shruti Dembla, April Y. Wang, Joanna McGrenere, and Parmit K. Chilana. 2017. Social CheatSheet: An Interactive Community-Curated Information Overlay for Web Applications. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW, Article 102 (dec 2017), 19 pages. https://doi.org/10.1145/3134737

[90] Laton Vermette, Joanna McGrenere, and Parmit K. Chilana. 2020. *Peek-through Customization: Example-Based In-Context Sharing for Learning Management Systems*. Association for Computing Machinery, New York, NY, USA, 1155–1167. https://doi-org.proxy.lib.sfu.ca/10.1145/3357236.3395507

[91] Lev S. Vygotskiĭ and Michael Cole. 1978. Mind in society : the development of higher psychological processes.

[92] April Y. Wang, Ryan Mitts, Philip J. Guo, and Parmit K. Chilana. 2018. *Mismatch of Expectations: How Modern Learning Resources Fail Conversational Programmers*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi-org.proxy.lib.sfu.ca/10.1145/3173574.3174085

[93] Xu Wang, Benjamin Lafreniere, and Tovi Grossman. 2018. *Leveraging Community-Generated Videos and Command Logs to Classify and Recommend Software Workflows*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3173574.3173859

[94] Sarah Weir, Juho Kim, Krzysztof Z. Gajos, and Robert C. Miller. 2015. Learnersourcing Subgoal Labels for How-to Videos. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (Vancouver, BC, Canada) *(CSCW '15)*. Association for Computing Machinery, New York, NY, USA, 405–416. https://doi.org/10.1145/2675133.2675219

[95] James V. Wertsch. 1979. From Social Interaction to Higher Psychological Processes A Clarification and Application of Vygotsky's Theory. *Human Development* 22, 1 (1979), 1–22. https://doi.org/10.1159/000272425

[96] Wolfgang Wohlgemuth and Gunthard Triebfürst. 2000. ARVIKA: Augmented Reality for Development, Production and Service. In *Proceedings of DARE 2000 on Designing Augmented Reality Environments* (Elsinore, Denmark) *(DARE '00)*. Association for Computing Machinery, New York, NY, USA, 151–152. https://doi.org/10.1145/354666.354688

[97] Feng Zhou, Henry Been-Lirn Duh, and Mark Billinghurst. 2008. Trends in augmented reality tracking, interaction and display: A review of ten years of ISMAR. In *2008 7th IEEE/ACM International Symposium on Mixed and Augmented Reality*. 193–202. https://doi.org/10.1109/ISMAR.2008.4637362

[98] John Zimmerman, Jodi Forlizzi, and Shelley Evenson. 2007. Research through Design as a Method for Interaction Design Research in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '07)*. Association for Computing Machinery, New York, NY, USA, 493–502. https://doi.org/10.1145/1240624.1240704