



Distributed Interaction Design

Designing Human-Centered Interactions in a Time of Social Distancing

Insights

- In times of crisis, innovation is essential and human-centered innovation is paramount.
- We have formulated a process to enable the creation and evaluation of interaction designs remotely, by end users for end users.

At the time of this writing, some countries are still enforcing stay-at-home orders and the world population is continuing to practice self-isolation to slow the spread of the coronavirus. In times of crisis, innovation is essential and human-centered innovation is paramount. Despite decades of improved design and usability practices, creating systems with interactions that are highly guessable, learnable, memorable, enjoyable, and accessible is still a persistent challenge. This challenge is exacerbated by the number of emerging intelligent technology platforms and environments such as

wearable devices, drones and robots, interactive surfaces and fabrics, voice-controlled intelligent assistants, and virtual and augmented reality environments. In the field of human-computer interaction (HCI), the practice of human-centered design tackles some of these challenges, but it has not been widely adapted to remote use by physically distant practitioners and users. We have formulated a process to enable the creation and evaluation of interaction designs remotely, by end users, for end users. In this article, we describe the work we have done to take user-centered design approaches out of the lab and online,

and how our work led us to formulate a process we call Distributed Interaction Design (DXD).

Following early work in HCI [1], we think of a single interaction between a user and a technology artifact as comprising three components: 1) a human input, 2) a system computing function, and 3) the system's feedback or output. To best determine what inputs should trigger what functions and outputs in a system, several methodologies, such as participatory design [2], incorporate end users into the design process. From 2005 to 2009, Wobbrock et al. [3,4] developed a related method, the *end-user elicitation study*, to make interactive systems more guessable, learnable, and usable. By incorporating end users of varying abilities, needs, backgrounds, and values directly in the design process, interactive systems could be more usable and inclusive.

The end-user elicitation study works by prompting users with the output of a computing function and asking them to propose the action that would trigger that function to bring about that output. In essence, it asks users to work *backward* from the system's response to the user's action, thereby eliciting the actions that users feel would be most likely to result in the responses they are shown. Over many participants in an elicitation study, patterns of similar proposals start to emerge that can be implemented in an interactive system. End-user elicitation has become popular, with more than 300 published studies by researchers utilizing this method to design a wide range of interactions: gestures for interactive tabletops, gestures for blind users of touchscreens, virtual and augmented reality interactions, smart TV controls, in-vehicle interactions, drone navigations, interactions for Internet of Things

devices, and human-robot interactions, to name a few.

DISTRIBUTED INTERACTION DESIGN

In this article, we are revisiting this methodology more than a decade after its inception and updating it to fit the current state of our world. We created an online research and design platform called the CROWDDESIGN engine [5] (<http://crowddesignengine.com>). With the scaling that this platform makes possible, the end-user elicitation methodology can be conducted completely online, reaching a global pool of participants, remedying the lack of access and user representation that is typical of lab-based studies and allowing researchers to involve users in the design process of future technologies at a global scale [6]. The engine also includes a tool to analyze study results efficiently by utilizing the wisdom of the crowds and machine learning [7], drastically reducing the time it takes to conduct end-user elicitation studies and evaluate their results. In addition, we formulated a method to validate user-generated interactions in a distributed fashion, called *end-user identification studies* [6]. From this work, we extrapolated a six-step process for designing user-centered technologies that we call Distributed Interaction Design (DXD).

Stepping out of the lab. Over time, we identified several areas to pursue opportunities for advancements. To begin with, elicitation studies are traditionally run in laboratory settings with around 20 participants on average. Given social-distancing rules, lab studies are impossible. But even if in-lab studies were possible, the limited number of participants would lead to results that did not necessarily represent a wide range of users. Also, we have shown that interaction designs generated by large numbers of

participants are preferable to those generated by one or a few professional designers [8]. To address the limited number of participants, we built a tool called *Crowdlicit* [6] that reconceptualizes the elicitation process and its best practices to run completely in a distributed manner. Researchers conducting distributed elicitation studies can run their studies either synchronously or asynchronously with any participant in the world who has access to a Web browser. *Crowdlicit* automates the collecting, organizing, and storing of user proposals in an easy-to-analyze manner.

Harnessing the power of the crowd.

Beyond the usual challenges of participant recruitment, study execution, and data capture, elicitation-study data analysis requires a determination of whether two elicited proposals are sufficiently similar to be grouped together as if they were the same interaction. In most elicitation studies, all elicited proposals need to be compared to each other using subjective human judgment, which requires great time and effort. In our platform, we created a tool called *Crowdsensus* to analyze elicitation studies—either by importing the data directly from *Crowdlicit* or manually uploading it—by harnessing the power of online crowd workers and machine-learning algorithms to analyze the results of elicitation studies four times faster than manual human analyses [7].

Distributed design evaluation. The literature employing elicitation studies published over the past decade shows that most studies conclude by reporting a set of user-generated interactions. However, the elicitation process lacked a formalized method to evaluate or validate these user-generated interactions. We therefore established a method called the *end-user identification study* [6] to evaluate input actions before investing the time and resources to implement these actions into interactive systems. These input actions could be new or existing actions designed by interaction designers, or sets of interactions resulting from end-user elicitation studies or other participatory interaction design methodologies [6]. Identification studies reverse the elicitation process

With the scaling that this platform makes possible, the end-user elicitation methodology can be conducted completely online, reaching a global pool of participants.

by presenting users with a human input action and asking them to propose the system function or system output they expect the input action would trigger. We built the CROWDDESIGN engine in a robust way to run and analyze both elicitation and identification studies in a distributed fashion.

THE SIX STEPS OF DXD

Putting all our work together, we created a six-step iterative process (Figure 1) to designing interactive systems with a global pool of participants. We illustrate the six steps using an example of how the methodology works.

Step 1: Set up the four pillars of a DXD study. From our experience building the CROWDDESIGN engine and running distributed user-centered design studies, we found that there are four foundational elements that need to be established and communicated to remote participants properly to ensure the success of a DXD study:

- *Rules of engagement.* Study instructions preceding the start of an elicitation study should establish the rules of engagement. These rules explain to the participant a) the environment in which they are to imagine the system being designed, b) the form of the system, and c) the system’s sensing capabilities. An example of this would be, “Imagine you are interacting with a TV set in your living room that is able to recognize voice commands.”

- *A list of functions.* Every interactive system has a list of functions triggered by user input—actions. In elicitation studies, these functions are used as prompts. For example, a media player has functions like “play” and “pause.”

- *Prompt modality.* A prompt can be presented to participants in various ways: as a text description; as still images (e.g., before and after pictures of the system state); as audible feedback (e.g., tones and beeps, or natural language output); or as video showing the effects on a system. All these different presentation modalities are available in the Crowdlic tool on the CROWDDESIGN platform.

- *Proposal modality.* The proposals collected from remote participants can take one of many forms: text descriptions of actions, like pressing a button or turning a knob on a physical

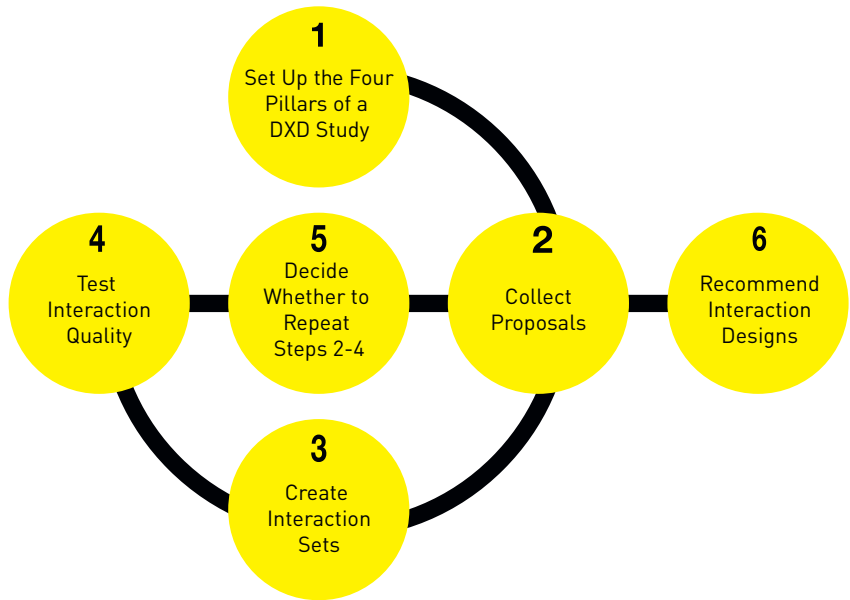


Figure 1. The six-step Distributed Interaction Design (DXD) process.

interactive system; natural language commands; or actual text-based commands for command line interfaces. They could also be still images or sketches. For dynamic proposals, they could be audio or video clips. Proposals can even take the form of annotations on a wireframe or existing user interfaces.

Example: Suppose a system creator is designing a new robotic arm that performs many functions triggered by midair gestures. The system creator might formulate the following study instructions: “Imagine you are interacting with a robotic arm sitting on your desk. It can sense your body movements and accept them as commands.” One of the functions the arm can perform is gripping an object. The creator then formulates the following prompt to present to participants in an elicitation study: “Perform a mid-air gesture that would make this robotic arm grip an object.” The system creator can represent this prompt in several ways other than text, such as two images, one showing the robotic arm with open fingers—the before state—and one where the fingers are closed together—the after state. Another way would be a video showing the robotic arm closing its fingers, accompanied with the instruction, “Perform a midair gesture that would trigger this movement.” Having viewed and understood the prompt, the participant then would provide a proposal for the gesture,

which the participant could describe in text, sketch as a sequence of images, or, best of all, perform and record as a video.

- **Step 2: Collect proposals.** Human input actions can be captured in many forms, as stated above. In the example of the system creator attempting to design a midair gesture for triggering the grasping function in their new robotic arm, the creator then recruits tens, or maybe hundreds, of participants for an elicitation study. The participants might propose mid-air gestures such as the three shown in Figure 2.

- **Step 3: Create interaction sets.** Once proposals for a prompt are collected, it is time to find the proposal with the highest consensus among participants. All proposals must be compared for their pairwise similarity and put into similarity groups. In our example, the system creator would group the gestures based on their similarity and implement the gesture with the highest consensus in the new system. So, the resulting gesture here would resemble proposals A and C. At the end of this step, the creator will have a list of input action designs—informed by actual end users—that map to the functions of the system.

- **Step 4: Test interaction quality.** Many metrics of the ISO 9241 standard for usability lend themselves to distributed evaluations, such as task, learning, and individualization suitability, and conformity with user




Figure 2. Three gestures by different participants in an elicitation study in response to the prompt, “Perform a midair gesture that would make this robotic arm grip an item.” Participants A and C proposed the same gripping gesture, whereas participant B proposed a different gesture.

expectations. As mentioned earlier, most end-user elicitation studies conclude by reporting a set of user-generated proposals, but without a decisive way to claim whether those proposals are good or not. We have established a method to test the quality of the proposal-prompt relationship called the end-user identification study [6]. An identification study is the reverse of an elicitation study: Participants see a prompt of an input action and guess what the system would do or the feedback it would provide. In our example, when running the identification study, the system creator would recruit new participants and present them with the resulting gestures from the elicitation study, like the hand-closing gesture for grabbing an object (Figure 3). The creator asks participants to propose the function that the robotic arm would perform in response to this gesture.

The creator would then group the proposed functions based on their similarity to find the one with the highest consensus in a similar manner to step 3 above. If the resulting proposed function matches the original function used to elicit the user-generated gesture, this indicates that the input action-system response

relationship is an identifiable one, and that the proposals are a good fit for this prompt. Because of the flexibility of the CROWDDESIGN engine, we can utilize it to run and analyze distributed identification studies with the Crowdlic and Crowdsensus tools. Other usability studies such as learnability and memorability studies can take advantage of Crowdlic and use it as a tool to gather, organize, and store data. Other metrics from the ISO 9241 standard, such as error tolerance and controllability, might be more suited for traditional usability testing, as they work within the context of use with the actual system being evaluated.

Step 5: Decide whether to repeat steps 2 to 4. In cases where the input action-system response relationship is not easily identifiable, a new round of proposal collection and design quality testing can be conducted. The DXD process is iterative for this reason. Researchers can repeat steps 2 to 4 as necessary until they arrive at a set of interactions generated and tested by end users to implement in their system.

Step 6: Recommend interaction designs. Finally, researchers and system creators can build systems informed by real users on a global scale, making future technologies inclusive of different users’

perceptions, values, and physical abilities. The DXD process is meant to inform designers, researchers, and system creators of actual end users’ needs, abilities, and preferences. DXD sits between a wireframing tool and a code editor. After system creators have established the system’s form and functions, they can get insight into how the system’s users would best interact with it before investing the resources to build the system.

REFLECTIONS ON DESIGNING WITH THE WORLD

Distribution trade-offs. Taking elicitation studies out of the lab has revealed many benefits, such as access to more diverse populations, enabling social distancing, cutting down on recruitment time, eliminating the physical spaces needed to run a lab study, and saving time by running dozens of participants within a few hours in an unsupervised manner. On the other hand, an unsupervised DXD study lacks the personal touch of an in-lab study. The result of such a DXD study is a series of discrete responses to a list of prompts augmented by demographic data. This trade-off is one that researchers need to consider when deciding if a DXD study is the correct approach for their design needs. Is it more important to prioritize the quantity of participants and diversity of user input, or the depth of information that comes from observing participants’ body language and taking think-aloud notes in a lab setting?

Data analysis. We found that the wisdom of the crowd yielded better study-analysis results than the status quo of one or two individual researchers. In our approach, an online crowd comes to consensus on the

In our approach, an online crowd comes to consensus on the similarity of design proposals, drowning out any bias (conscious or unconscious) that only a few individuals might introduce.




Figure 3. The gesture with the highest consensus from participants in the elicitation study of Figure 2.

similarity of design proposals, drowning out any bias (conscious or unconscious) that only a few individuals might introduce. We also showed that the crowd using our platform is four times faster than an individual researcher analyzing an elicitation study. These benefits do come at a price, however, as our crowd-based approach to analysis can become expensive as the number of proposals increases. We plan to investigate a remedy to this drawback in the future by utilizing more advanced machine-learning techniques.

CONCLUSION

We hope readers will adopt or adapt our DXD process to design future technologies with users all over the world, pushing innovation in these difficult times. We believe our process unlocks the possibility for exceptional work in evaluating and improving current interactive systems and exploring ways to interact with future ones.

ACKNOWLEDGMENTS

This work was supported in part by Microsoft Research, the Mani Charitable Foundation, the University of Washington, and the National Science Foundation under grant IIS-1702751.

ENDNOTES

1. Morris, M.R., Wobbrock, J.O., and Wilson, A.D. Understanding users' preferences for surface gestures. *Proc. of Graphics Interface 2010*, 261–268.
2. Wobbrock, J.O., Aung, H.H., Rothrock, B., and Myers, B.A. Maximizing the guessability of symbolic input. *CHI '05 Extended Abstracts on Human Factors in Computing Systems*. ACM, New York,


2005,1869–1872; <https://doi.org/10.1145/1056808.1057043>

3. Wobbrock, J.O., Morris, M.R., and Wilson, A.D. User-defined gestures for surface computing. *Proc. of the 27th International Conference on Human Factors in Computing Systems*. ACM, New York, 2009, 1083–1092; <https://doi.org/10.1145/1518701.1518866>
4. Wobbrock, J.O., Morris, M.R., and Wilson, A.D. User-defined gestures for surface computing. *Proc. of the 27th International Conference on Human Factors in Computing Systems*. ACM, New York, 2009, 1083–1092; <https://doi.org/10.1145/1518701.1518866>
5. Ali, A.X. The CROWDDESIGN Engine; <http://crowddesignengine.com>
6. Ali, A.X., Morris, M.R., and Wobbrock, J.O. Crowdlicit: A system for conducting distributed end-user elicitation and identification studies. *Proc. of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, New York, 2019, 255; <https://doi.org/10.1145/3290605.3300485>
7. Ali, A.X., Morris, M.R., and Wobbrock, J.O. Crowdsourcing similarity judgments for agreement analysis in end-user elicitation studies. *Proc. of the 31st Annual ACM Symposium on User Interface Software and Technology*. ACM, New York, 2018, 177–188; <https://doi.org/10.1145/3242587.3242621>
8. Schuler, D. and Namioka, A. *Participatory Design: Principles and Practices*. CRC Press, 1993.

1 Abdullah X. Ali is a Ph.D. candidate at the University of Washington's Information School. His research focuses on formulating methods and building systems that leverage input from crowds of end users and uses machine-learning algorithms to design technological innovations, which led to the development of the Distributed Interaction Design (DXD) process.

2 Meredith Ringel Morris is a principal researcher at Microsoft Research. Her research primarily focuses on collaborative and social technologies. Her work on interaction techniques for large touchscreen displays led to her interest in gesture elicitation. She received her Ph.D. in computer science from Stanford University.


3 Jacob O. Wobbrock is a professor in the Information School and an adjunct professor in computer science and engineering at the University of Washington. He is a member of the CHI Academy. His research focuses on input and interaction, notably the EdgeWrite alphabet, the first gesture set created via a user-defined gesture methodology.



2018 JOURNAL IMPACT FACTOR: 6.131

ACM Computing Surveys (CSUR)

ACM Computing Surveys (CSUR) publishes comprehensive, readable tutorials and survey papers that give guided tours through the literature and explain topics to those who seek to learn the basics of areas outside their specialties. These carefully planned and presented introductions are also an excellent way for professionals to develop perspectives on, and identify trends in, complex technologies.



For further information and to submit your manuscript, visit csur.acm.org