

Negotiating the Creative Space in Human-Robot Collaborative Design

Matthew V. Law¹, JiHyun Jeong¹, Amritansh Kwatra², Malte F. Jung¹, and Guy Hoffman³
 Information Science¹, Computer Science², Mechanical and Aerospace Engineering³
 Cornell University
 Ithaca, NY USA
 [mvl24; jj639; ak2244; mfj28; hoffman]@cornell.edu

ABSTRACT

We describe a physical interactive system for human-robot collaborative design (HRCd) consisting of a tangible user interface (TUI) and a robotic arm that simultaneously manipulates the TUI with the human designer. In an observational study of 12 participants exploring a complex design problem together with the robot, we find that human designers have to negotiate both the physical and the creative space with the machine. They also often ascribe social meaning to the robot's pragmatic behaviors. Based on these findings, we propose four considerations for future HRCd systems: managing the shared workspace, communicating preferences about design goals, respecting different design styles, and taking into account the social meaning of design acts.

Author Keywords

human-robot interaction; collaborative design; tangible user interfaces

CCS Concepts

• **Human-centered computing** → *User studies; Collaborative interaction;*

INTRODUCTION

In this work we study how humans design a complex system together with an autonomous robot. Robots have the potential to be useful collaborators in design activities for two main reasons: humans and AI agents complement each other as designers [14] and designing in a physical space supports collaboration and creative thinking [36, 27].

Motivated by these reasons, we developed a human-robot collaborative design (HRCd) system (Figure 1). In our system, users place and arrange blocks representing design components on a sensor-equipped work surface. A screen gives immediate feedback about the current design configuration

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

DIS '19, June 23–28, 2019, San Diego, CA, USA

© 2019 ACM. ISBN 978-1-4503-5850-7/19/06...\$15.00

DOI: <https://doi.org/10.1145/3322276.3322343>



Figure 1. A human and a robot working together on a design task using the human-robot collaborative design system described in the paper.

and allows recall of previous designs. The user is further assisted by an autonomous robotic arm that tracks the user's action and manipulates the blocks together with the human to achieve better designs.

We conducted an observational study ($n = 12$) in which participants collaborated with the system on a design task. In video observations and interviews, we find three themes. The first relates to how human designers negotiate the physical workspace and take turns with the robot. The second concerns how the human and robot manage the “creative space” of the design process, negotiating roles, goals, and strategies. Additionally, we find evidence that participants attribute social meanings to the robot's pragmatic actions, even though we intentionally did not design any social cues—such as facial expressions, speech, or communicative gestures—into the system.

From these three findings we propose four design considerations for future HRCd systems: managing the shared workspace, communicating personal preferences about design goals, respecting and expressing different design styles, and taking into account the social meaning of design acts.

Why Use Robots for Collaborative Design?

Robots have been suggested as collaborators for humans in physical tasks such as manufacturing or assembly. But humans and robots could make for good collaborators in solving complex design problems as well. This is for two main reasons:

Humans and Machines Have Complementary Design Skills

Collaboration is integral to the practice of designing, as the design process itself can be thought of as a conversation, or dialectic, through which ideas emerge [7]. Multiple designers can consider areas of specialized expertise and different perspectives [4]. Tensions due to diversity in thinking styles can lead to more effective creative solutions [49]. Finally, human creativity is a social phenomenon and new ideas can be generated through situated interactions with others [15].

The particular differences between humans and computers make them especially well-suited to explore the complexities of design together. Design problems are ill-structured, dynamic, and often context-specific [33]. This plays well into human designers' ability to adapt, contextualize, and integrate prior knowledge across domains, whereas a computer's ability to evaluate many designs quickly and have less biases fills crucial gaps often found in human designers [14]. Mixed-initiative design between a human and a computer can foster co-creativity through enhanced re-framing and lateral thinking [31], and co-design between humans and multiple computer agents can increase the number of perspectives considered during design [13]. Finally, humans can account for priorities and preferences that are difficult to formally represent, while computer agents can consider factors that may be difficult for humans to conceptualize, for example, optimizing over high-dimensional spaces.

Physical Embodiment Benefits Design

Engaging physically with a design problem and other designers amplifies the positive effects of designing collaboratively. Physical embodiment allows for co-located, synchronous collaboration, enabling rapid feedback, more nuanced communication, easy joint reference to shared objects, shared local context, access to implicit cues, and easier individual control of interaction dynamics [36]. Embodying the design agent in a robot promises similar benefits to human-computer collaborative design that physical collaboration offers teams of human designers.

Direct physical interaction also affects the way that we think and learn together. Externalizing a design plays a critical role in distributed cognition by enabling a "conversation with materials" that is qualitatively different because it is physical [4]. Physical action is coupled with human cognition, for example, via epistemic gestures that allow us to understand concepts through the movement [27]. The visibility of physical bodies and embodied actions can also facilitate situated learning, which can support the development of shared understanding between collaborators [27]. Finally, embodiment increases the risk of actions. Increased social risk can improve the commitment and focus of designers [27].

Taken together, the complementary design skills of humans and computers and the benefits of physically embodied design

lead us to propose the use of robots for collaborative design. Human-robot collaborative design can contribute not only to the literature in human-computer design, but also extend the current state of the art in human-robot collaboration, which mostly focuses on physical, rather than conceptual, activities.

RELATED WORK

Human-Computer Collaborative Design

Prior work in human-in-the-loop design typically frames the process by either having the human evaluate designs proposed by the agent, e.g. as the fitness function in an interactive genetic algorithm [14, 5, 11], or having the agent find satisficing solutions to a set of goals and constraints defined by the human designer [44]. In some work, agents make suggestions to the human based on user exploration [3, 30].

Human-Robot Physical and Creative Collaboration

Meanwhile, the rich literature on human-robot collaboration is largely concerned with performing physical tasks. These tasks, including assembly [1, 43], fabrication or construction [47, 48], rehabilitation [25], and sorting parts [18], often lack the ill-defined structure and dynamic nature of design tasks. In Wainer et al. [50] a robot coaches a human in solving the Towers of Hanoi puzzle, but this task is more of a mathematical problem than a design problem and the robot does not directly manipulate the puzzle. Kahn et al. [23] found creative effects of interacting with a social robot when designing a Zen rock garden. However, their focus is on the social and conversational capacities of the robot, which does not itself manipulate the garden, but rather prompts the human with encouragement or visual stimuli to facilitate creativity. A handful of projects have explored creative performance with robots [21, 34]. While these deal with ill-structured tasks, we are more explicitly interested in generative design problems than performance or art.

At least one project has looked at using robots in a supportive, but not creative, role for CAD design [37]. In contrast, our work explores human-robot collaboration on an abstract design task where both the robot and the human play a creative role in generating new designs.

Tangible Interfaces with Robots

Tangible user interfaces (TUI) have been investigated as a collaborative design support tool for teams of human designers [26, 51, 28, 41]. Cassell et al. introduced "shared reality" where an embodied conversational agent interacted in a virtual world mapped to a physical interface [10]. In contrast, we propose shared control over physical objects by two actors in the same space. There is also prior work that uses TUIs as a medium of coordination between a human and a robot. However, in these cases, the TUIs are used as a mechanism for the human to program or control the robot [17, 12]. Instead, our work uses a TUI as a shared workspace through which a human and a robot partner can manipulate shared objects to explore and build a design together.

Against this background, we next describe the first contribution of our work: the development of an interactive system for a human and a robot to collaborate on a design task.

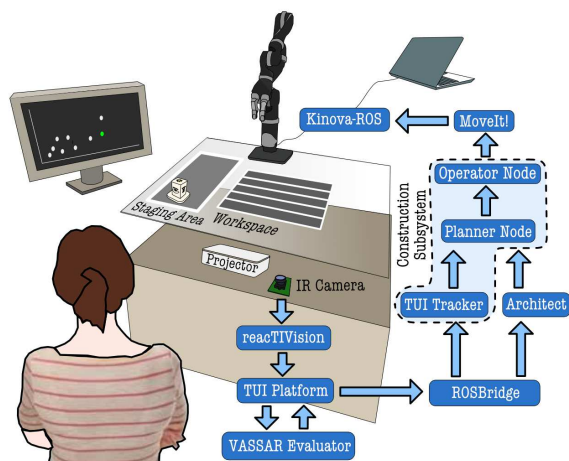


Figure 2. The physical HRC system is composed of a tangible user interface (TUI) and a manipulator arm. The TUI software (bottom center) evaluates designs and visualizes them, and passes the results to the robotic system that selects designs (right) and builds them (top).

THE HUMAN-ROBOT COLLABORATIVE DESIGN SYSTEM

In our setup (Figure 2), the user sits across a tabletop TUI from a robotic arm. Both human and robot can adjust a set of shared blocks which, arranged on the TUI, represent the design they are working on together. As the human adds, removes, and moves blocks around different regions of the table surface, the TUI system tracks the corresponding designs and evaluates them. It visualizes the outcomes on a monitor to the side of the table for the human to evaluate and sends the outcomes to the robot, which reasons independently about the design problem. When the robot wants to modify or try a new design, it rearranges, adds to, or removes the blocks on the table in the same way as the human (Figure 3).

In this work we focus on an important and common class of design problems called *Configuration Design*. Configuration design tasks require selection and arrangement of a set of components to satisfy constraints, requirements, and optional optimality criteria [35]. The configuration design process is thus characterized by iteratively selecting components, defining their relationships to one another in a design, and evaluating and refining according to the prescribed objectives [9]. For example, when designing a race car, design engineers have a choice of frames, engines, fuel components, safety devices, and so on. These can be connected or positioned in different spatial relationships to each other and affect performance outcomes such as cost, fuel efficiency, power, etc. The design outcomes are often tied in complex ways to the selection and integration of the components. Other examples include designing a mobile phone from stock components, or an irrigation system for agricultural purposes. Configuration design problems are extremely common in industry, and map well onto computational systems, including optimization and search algorithms.

The Tangible Workspace

The TUI on which the human and robot explore their designs is based on the reactTIVision ReacTable [24] and our prior

work [30]. Passive blocks, each tagged on the underside with a fiducial marker, are used to represent different design components. Through an infrared camera inside the table, reactTIVision identifies and locates each block on the surface by its fiducial.

The table surface has a “workspace” on which block “components” can be spatially arranged to represent the current design configuration. A separate region, which we call the “staging area,” is reserved to hold blocks that are not currently in use but are available to be added to the configuration by either actor. Any blocks placed elsewhere cannot be seen by the robot, but the entire space is free for the human to use as they see fit.

Design Evaluations and Feedback

Block arrangements are translated into design configurations and evaluated computationally. The outcomes are plotted on a monitor next to the table. A human designer can select a data point on the monitor to see the associated design projected back on the table surface. This allows the human designer to reconstruct a past design. In this work, we use a scatter plot with each point representing one design, and the axes representing the design objectives.

The Robotic Design Collaborator

Our robotic design collaborator is a Kinova Ultra Lightweight 7-DOF manipulator arm, driven by two main logical components: an “Architect” search module that suggests designs to try and a “Construction” subsystem that determines how to create the suggested designs using blocks on the TUI. Communication between the various components of the system is handled using the robot operating system (ROS) [38].

The Architect: Selecting Designs to Build

The “Architect” is a ROS node that listens to all the designs evaluated by the TUI. For each design, it scans random one-step perturbations (e.g. a random component added or removed). The Architect maintains a Pareto frontier of designs across objectives including the human’s designs and its own random perturbations. The frontier is recalculated for every new design on the workspace. On request, it suggests a design from the current Pareto frontier. Suggestions are based on blocks available, prior suggestions, and minimizing the number of blocks that must be moved to realize the change.

The Construction Subsystem: Building Intended Designs

The “Construction” subsystem executes new designs using three ROS nodes: the “Planner,” the “Operator” and the “TUI Tracker.” The Planner obtains a design suggestion from the Architect and translates it into a sequence of block changes. The Operator takes the sequence of block changes and finds the corresponding target blocks and placement locations. It calls a planning library, MoveIt! [45], as well as the Kinova API [40], to plan and execute each associated pick and place action while handling failures that may occur. The TUI Tracker tracks changes in the block arrangement and is responsible for maintaining the mapping from the Kinova arm’s Cartesian space to reactTIVision’s coordinate space.

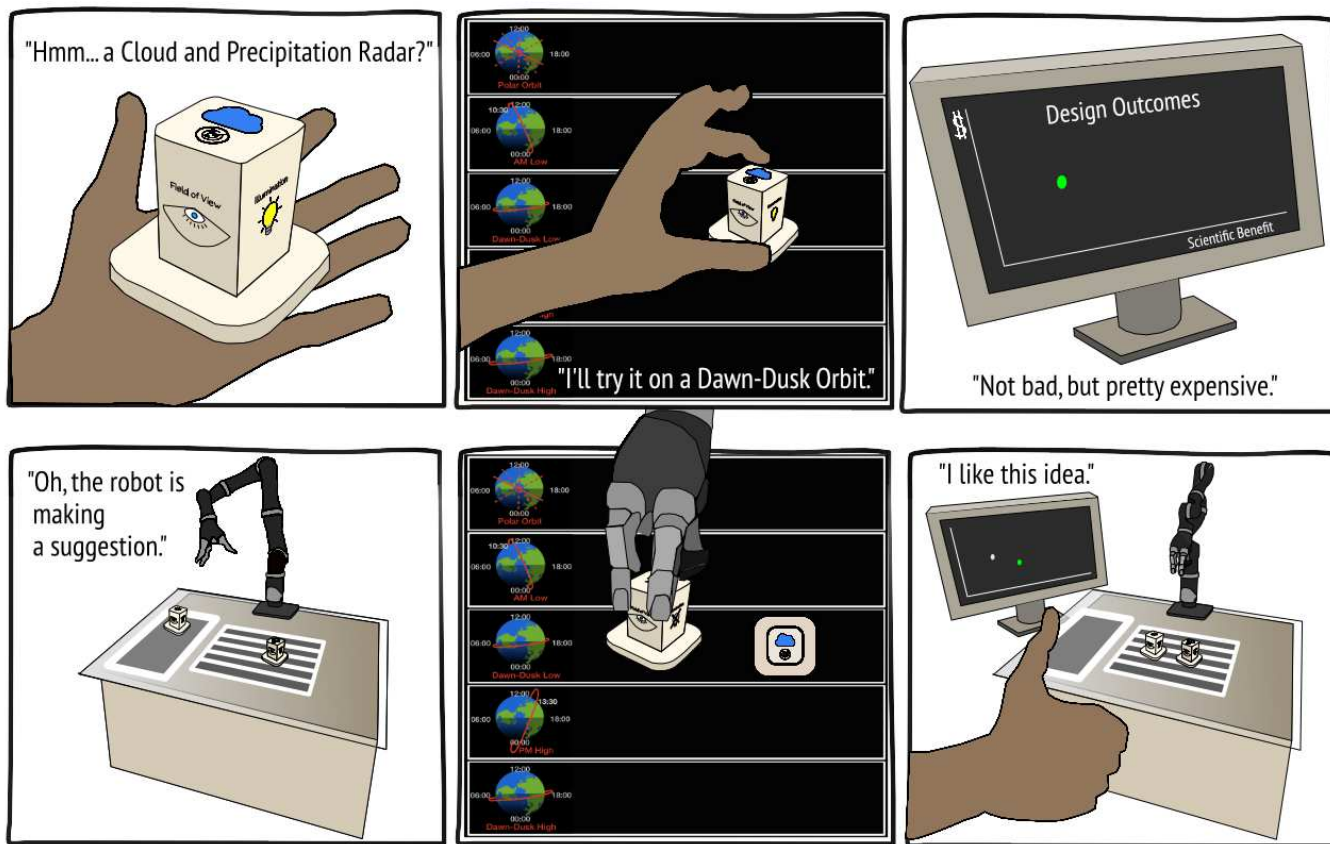


Figure 3. As the human is exploring different designs by placing sensors into orbits on the table and monitoring the outcomes (top row), the robot follows along and performs its own search. The robot suggests designs by moving blocks on the tangible interface in the same way that the human does (bottom row).

The arm then sequentially moves the blocks necessary to realize the new design. It finishes by withdrawing to a retracted position off of the workspace. At this point, the Planner can pause for a tuneable delay before requesting a new design from the Architect. Note that the robotic arm moves slower than the average user: its top linear speed is 20 cm/s [39], and it pauses between individual moves.

STUDY SETUP AND ANALYSIS

The second contribution of this work is an observational study of humans working with our system. We sought to address two major research questions: What are the features of an HRCD interaction; and what are the implications for designing future HRCD systems?

The EOSS Design Problem

The design challenge we use as a case study for our system is the configuration of an Earth-Observing Satellite System (EOSS) [19]. In order to monitor the earth’s climate, designers put together 12 types of sensors into five possible orbits. Designs are evaluated based on how well they satisfy a set of 371 data collection criteria (this is called the design’s “science benefit”), as well as on the system’s overall lifecycle cost.

The orbits pass over the earth at different local times, at different altitudes, and on different paths. The sensors, which

include instruments like an ocean spectrometer and a vegetation laser sensor (LIDAR), measure different climate-related quantities. They can have active or passive illumination, different power requirements, mass, field of view, and other features affecting the overall design. Deploying sensors in different orbits can significantly affect the data they collect and how much they cost. In addition, sensors deployed together can interact in complex positive or negative ways. For example, launching two particular sensors together may cost less than launching them separately, but they may interfere electronically, reducing their usefulness.

We choose this real-work design task for three reasons. Firstly, it is comparable to other complex configuration design problems. Designers choose components and spatially arrange them to build a configuration. Similarly to many other complex configuration design tasks, the choice and arrangement of components and the complex interactions between them have a direct effect on design outcomes. Secondly, there is an established custom rule-based architecture evaluation engine called VASSAR [42] to evaluate designs. In this study, we use VASSAR to simplify the evaluated outcomes and provide two numbers as feedback for designers, a percentage for “science benefit” and a dollar value for “cost.” Lastly, the problem is open-ended enough to allow for diverse design styles and




Figure 4. Each session was video and audio-recorded, with log files recorded for the robot's actions and design evaluations. Video data included a frontal view of the participant, an overhead view of the table, a screencast of the design outcomes plot shown to the participant, and the status output from the robot's "Planner" node.

goals. Participants can be minimalist in their goals: they may value low-cost designs even at the expense of performance. This introduces the possibility of collaborative negotiation of design objectives, which is a characteristic that can generalize to other design problems.

Procedure

The study was conducted in a laboratory room designated for observational studies, with controlled lighting, and no additional visual distractions.

After providing informed consent, participants watched a five-minute video that explained the design task and the tangible interface. We then asked them to spend up to thirty minutes working on the task with the robot encouraging them to think aloud about their design process and considerations. During the first five minutes, participants would design on their own, without the robot intervening, so the agent could collect information about possible designs. At the end of the session, participants took part in a short semi-structured interview that probed their approach to the design task and experience working with the robot. Depending on the participants' engagement, interviews lasted between 4-17 minutes. Participants could withdraw and end their session at any point during the study. Prior to running the study, we refined our system and our study protocol based on feedback we received from six pilot studies.

Participants

We recruited 12 users to design satellite systems with our robotic arm. All participants were students; nine were undergraduate students and three were graduate students. Ages ranged between 19-25 years old. Six of our participants were female. We recruited participants using a university study recruitment system, in addition to advertising through a teaching assistant program and a class. All participants were compensated with \$20 or extra credit for their time.

Data Analysis

With participants' consent, we recorded the studies and interviews using video and audio (Figure 4); log files were also kept of the robot's actions, as well as all the designs and

associated outcomes attempted by the human/robot pair. Participants signed release forms that indicated whether or not they consented to the use of their likeness in publications.

Using thematic analysis [8], we extracted themes from our videos and interview transcripts in two passes. Interviews were transcribed by a third party service and manually corrected for errors by the researchers. On a first pass, the material was divided among three coders who each identified noteworthy events in their portion of the data. From these events, the coders collectively determined five common themes as the focus of the next pass: *Sensemaking*, *Turn-taking*, *Roles*, *Trust*, and *Exploration Styles*.

On a second pass, each coder focused on identifying instances of these themes, compiling a summary for each session and presenting to the rest of the group. From the second pass, we distilled three overarching themes from the five original themes: *Sharing Physical Control of the Design Interface*, *Sharing Creative Control of the Design Process*, and *Attributing Social Meaning to Actions*. For each of these, we compiled emblematic behaviors and counted users that exhibited them. Each instance was verified by a second coder, with disputes settled by the third. Finally, we collectively agreed on strong instances that illustrate each of the findings for publication.

In addition, we extracted data from our log files to characterize each of the sessions (Figure 5). We experienced technical issues with the TUI in two of the sessions and with the robot in two others. However, the participants were still able to explore designs with the robot, and we were able to extract meaningful findings from the interactions despite the issues.

OBSERVATIONS

The 12 participants worked for, on average, 22 minutes and 43 seconds ($sd=7:43$, $min=8:34$, $max=30:18$, $median=25:14$), evaluating on average 138 designs ($sd=54.49$, $min=42$, $max=241$, $median=137$). Designs averaged a science benefit of 0.1479 ($sd=0.0385$, $min=0.0526$, $max=0.1819$, $median=0.1614$) and a cost of \$2.627 billion ($sd=\0.6259 billion, $min=\$1.231$ billion, $max=\$3.605$ billion, $median=\$2.701$ billion).

Figure 5 shows the rate at which participants explored designs alone and with the robot, as well as the design outcomes associated with each user, allowing comparison between user strategies. For example, we can see that Participant 5 explored more slowly than Participant 7, indicated by the lighter shading, but worked longer with the robot. In the scatter plots we can see that, when working with the robot, Participant 5 tended to explore lower-cost designs than Participant 7, and still found some configurations with higher science benefit. Participant 5 said she would work with the robot again, while Participant 7 would not, indicated by an asterisk in the figure.

Video Analysis

From the video analysis of the design sessions and the follow-up interviews, we identify three themes relating to how participants negotiate the collaborative design process with a robot. The first theme relates to negotiating the physical space of the interface, and the second relates to negotiating the creative

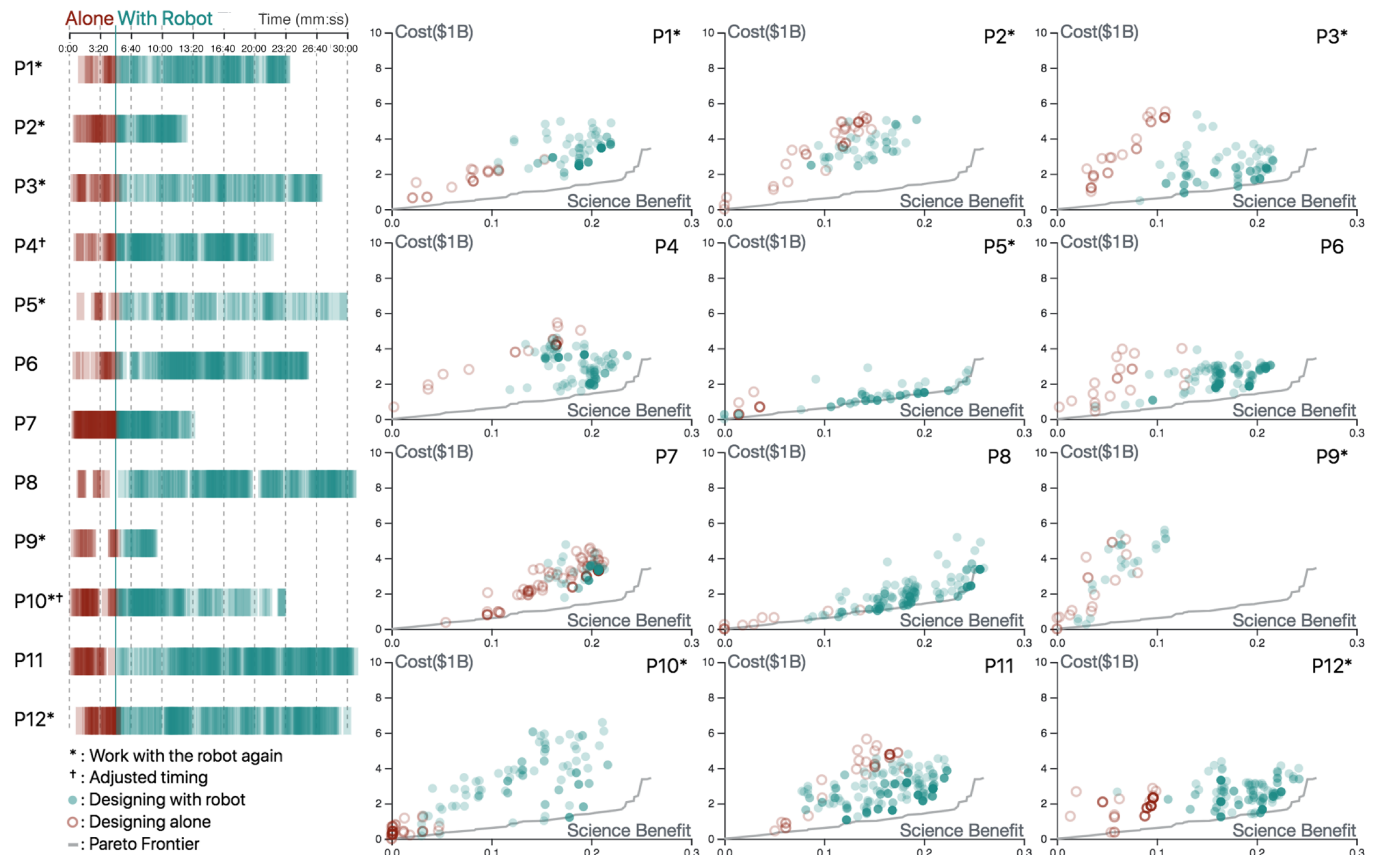


Figure 5. This breakdown of each collaborative design session shows a density plot of the frequency at which participants evaluated designs alone and with the robot (left), and the outcomes for each design evaluated (right). The right plots also show the Pareto frontier of all designs explored during the study as a reference. P4 and P10 have been adjusted to remove time spent repairing technical faults.

space of the design exploration. In a third theme, we observed that people engaged with the robot in a social way, despite the purely pragmatic design of our system.

Negotiating the Physical Space of the Tangible Interface

Observations related to negotiating the physical space can be further broken down into coordinating turn-taking and coordinating the spatial use of the shared table.

Coordinating Turn-Taking

With regards to turn-taking, we observed two issues:

- *Varied Response to the Robot Taking a Turn:* The majority of the time, participants tried to take a strict turn-taking approach with the robot. As one explained, “I didn’t do stuff at the same time [as the robot]. [...] I let it go about its thing” (P6-F-20)¹. When the robot moved to make a change, sometimes to the surprise of participants, most of them responded by pausing for the robot mid-action. Some participants abruptly retracted physically and some adopted deferential poses: stepping back, pulling their hands to their chests, putting them behind their backs, or in their pockets.

¹We refer to participants by *id-gender-age*, e.g. P1-M-21 is Participant 1, male, 21 years old. See also “P1” in Figure 5.

At other times, participants continued experimenting with blocks while the robot moved through the workspace, *de facto* ignoring the robot’s action. Changing the block configuration while the robot moves alters its design suggestions, and can make the robot’s move be less effective. Also, the human might remove a block the robot is reaching for, causing the robot to grasp vainly at the vacated spot and, failing to find the block, retract slowly. Some participants ignored this and kept working, while others who noticed this failure waited for the arm to finish moving. After realizing that she had moved a block the robot wanted to take, one participant (P5-F-19) returned it to its original position and waited, hoping the robot would come back for the block.

- *Determining What Constitutes a Turn:* Given the turn-taking approach that most participants took, understanding what actions delimit a single turn was an important challenge. Our robot defined a turn as the sequence of block movements necessary to realize a single new design. Between block movements in the same sequence the robot would pause momentarily, and once the sequence was complete, it would retract to the home position. Many participants, on the other hand, interpreted the robot’s moving a single block as a complete turn, and a signal to react to that turn. As a result, participants would often add or remove

blocks while the robot was still attempting to construct the design it had intended.

Participants had different conceptions of the duration of their own turns. In the interview, P2-M-19 thought of turns as each agent trying one design at a time. In contrast, P4-M-25 preferred to work for a prolonged period of time and then let the robot manipulate the workspace without him intervening for a while. A third participant thought the robot should intervene when it sensed that she needed help, saying, “the robot came at the right time to help me” (P3-F-25).

Four participants expressed a desire for stronger turn-taking behavior from the robot or control over when it could move. Three of them suggested the ability to ask the robot to stop and wait, either verbally or with some sort of button, e.g. (P4-M-25): “Maybe there is a button that I can control that says take a rest or there’s a button that says do [something].”

Coordinating Shared Space

Sharing the interface involved negotiating the utilization of space as well. Three participants made use of the physical space to communicate with the robot by placing more blocks on the staging area for the robot to see. Similarly, P1-M-21 tried to prompt the robot to move by adding a sensor to the orbit area that he thought it liked. Three participants likewise removed blocks away from where the robot could see after concluding that the blocks weren’t useful in their designs, preventing the robot from using them.

Six participants made use of the tangible interface to spatially organize their thoughts; e.g. sliding blocks to the right to signal their intention of keeping the blocks on the orbit. While the robot never disrupted this new organization of blocks, there was nothing in the robot’s programmed behavior to account for this or to interpret it.

Negotiating the Creative Space of Design Exploration

In addition to the participants’ negotiation of the physical space of the interface, we observed them negotiating creative control in the design process. This included negotiating the roles of the human and the robot, as well as coordinating design goals and strategies to achieve these goals.

Coordinating Roles

Participants interpreted the robot’s role and their own role with respect to it in different ways. We gave no guidance to participants in this regard, and five participants expressed confusion with what role the robot might play before it started, some asking how to control it or turn it on. Two participants said they did not expect the robot to do more than physically assist them with the interface. Interestingly enough, all participants ended up physically assisting the robot by adjusting the blocks it placed or figuring out how to complete actions when the robot missed or dropped a block. Across the roles participants ascribed to the robot, the following stood out:

- *Robot as a subordinate:* Three participants (P4, P6, and P11) adopted a leadership role and took a more dominant stance with respect to the robot. One participant (P6-F-20) often maintained a hands-on-hips pose as the robot was

moving or tapped the the table when she got impatient with the robot. On occasion, she also verbally commanded the robot to “Keep it[the block] there,” when the robot started to move a block she liked. Another participant accepted or rejected the robot’s actions based on whether or not he viewed it helpful to his own approach. At times he granted permission to the robot as it moved, saying, “Ok fine, machine, take this out” (P11-M-25). One participant described what he thought the role of a robotic arm should be during his interview, “If you really have this kind of task I would just let it help me [...] Just like some kind of hard labor, but without literally interacting with you” (P4-M-25).

- *Robot as the leader:* Two participants (P5 and P10) felt that the robot assumed the role of the leader in the interaction and saw themselves as subordinates. Both felt the robot was more capable at performing the task than they were, and trusted that its actions would be effective. One participant interpreted the robot’s removal of blocks from the workspace as a critique of the design, and would propose new changes to the configuration in response to the robot’s suggestion. For instance, in response to the robot’s removing a block she liked, the participant conceded, “I’m going to try to use this instead, it has less mass” (P5-F-19). Another participant found that the robot’s additions to the designs led to better outcomes, and felt that helping the robot successfully construct its designs was more effective than coming up with novel ideas himself. He said, “At this point the robot is doing better than I am, so I’m just trying to help it build” (P10-M-19). Both participants trusted the robot’s actions through the entire study period. Despite regarding some of the robot’s suggestions as sub-optimal, both participants followed the robot’s lead in the exploration.
- *Robot as a colleague:* One participant (P3) adopted a more blended partnership with the robot. Like many others, she helped the robot by adjusting blocks it placed when the detection was obstructed. At two other times, the robot seemed to help her, removing a block that she was reaching to remove. To illustrate, she added a sensor, then considered removing it, when the robot intervened and moved it to a different, more effective orbit, causing her to remark, “that’s better.” Throughout the design session, both she and the robot added new blocks to the design on the table. Both also removed blocks that the other added; on three occasions when the robot removed blocks she had just placed, she verbally agreed, e.g. saying, “Yeah, that was a mistake.” Towards the end of the session, the participant ruminated while placing a sensor she thought the robot might dislike. She said, “It’s answering more questions at a little more cost, so I prefer to leave it there, but maybe he’ll move it,” underscoring the back-and-forth dynamic she had with the robot. The robot left the sensor where it was. The participant reflected afterwards that “it was more than assistance.” “I felt like I was having a colleague, and we were discussing through this process” (P3-F-25).
- *Robot as an adversary:* Three participants (P1, P7, and P9) developed an adversarial relationship with the robot based on their interaction with it. For one participant, the robot

was limited to removing blocks because the participant had laid blocks on areas that the robot could not see. At some point, she declared that all she was doing was undoing the robot's actions, "Actually, this robot is kinda annoying. Now I feel like I'm just trying to move the stuff back that the robot moves from the table," (P7-F-21). Another participant, who confessed that he didn't understand what the robot was doing after a few suggestions, started acting randomly and placed multiple blocks one after another, saying he was trying to "trick the robot." In his interview he speculated, "Maybe [the robot was] trying to confuse me. I don't know if it was trying to beat me. It was just picking up random blocks and putting them down" (P9-M-19). Similarly, one participant described his experience of losing trust in the robot in his interview stating, "I still feel like it did know the answers but it was playing with me" (P1-M-21).

Coordinating Goals

Our design task is a multi-objective problem in which participants need to trade off science benefit and cost. Friction arose around how participants and the robot perceived these trade-offs. Nine participants expressed hesitation over the robot's suggestions based on the trade-offs in the marginal differences for scientific benefit and cost. The robot did not account for these trade-offs amongst non-dominated designs, and participants sensed when this misaligned with their own preferences at a given time.

In one example, a participant (P5-F-19) had been working carefully to keep cost low while adding instruments to increase scientific benefit. The robot removed two blocks that moved the configuration to a more cost-effective but less useful design along the Pareto frontier. The participant, normally deferential to the robot, put both back, and then the robot removed them again; the participant gave in but remarked that the scientific benefit had dropped "significantly" despite the cost savings. Later on, the robot again removed two similar blocks and the participant immediately put them back, saying, "I don't think the difference in cost was that significant," and "I think the percentage [of data collected] is a lot greater [with that block]." After the session, the participant expressed a desire for the robot to explain "why, perhaps, something like its decision was better than mine."

Other participants also reflected on this tension in their post-interviews, suggesting that the robot had its own goals. One participant said, "I wouldn't say it was stupid for wanting its own agenda but it was annoying to me" (P12-F-21).

For another participant, the confusion and frustration around trade-offs eventually led to a breakdown of trust with the robot, as he eventually noted, "I'm not trusting the arm." He elaborated on his interaction during his interview, "I'm not sure if it knew where I was trying to go, like what the goal was" (P1-M-21).

Coordinating Strategies

Finally, coordinating strategies with the robot played a role in how well participants felt they could collaborate with the robot. To begin with, participants had different strategies to approach the design problem, sometimes at different times within the

same session. Based on their reasoning during the session and their self-reported strategies, eight participants tended to be more strategic and reason from theories about the sensors and orbits. Others more naively explored combinations and focused on outcomes. Three participants explicitly mentioned shifting strategies over the course of the session.

Conflicts arose when the robot failed to account for human strategies. P8-F-21 took time to carefully understand each block and orbit before attempting any new designs. She complained that the robot was much faster than her and that she could not remember the designs she tried before because the robot "took a lot of the blocks out." Her frustration resurfaced in her interview: "Before I [can] think very carefully about what this combination means, [the robot] moved it away so I had to rethink." Although she thought "it's helpful," she preferred not to work with the robot on the task because she wanted "plenty of time to study the orbits and the blocks so I can figure [it] out, mathematically or something" (P8-F-21). On the other hand, one participant who adopted a faster brute-force approach to trying new designs, complained that the robot was lagging behind him (P4-M-25).

Many participants attempted to discern the robot's strategy based on its actions. One participant tried particularly hard to understand and rationalize the robot's strategy by examining the block that the robot put down. For instance, he would justify the robot's move by saying "so this one has illumination so it's good to be in here." Or he would comment that the robot was "trying to hint at things" as it took a certain block off the orbit space saying, "so it took off one of the wave ones so I'm thinking maybe it will belong somewhere else" (P1-M-21).

Participants felt that understanding the robot's strategy was critical to figuring out how to coordinate strategies with the robot, and expressed concern about the lack of transparency: "The main problem is I could not read the robot [...] If I could figure out with a systematic pattern then it would help me in my own strategy of trying to solve this" (P11-M-25). Another participant suggested what she thought was missing: "Well I don't know if you can really tell a robot your strategy and have it understand something more complex like that [...] it could have been helpful if we could have collaborated or debated on the strategy [...] and then compromised" (P12-F-21).

Navigating Social Meaning in Actions

We did not program the robot to exhibit explicit social behaviors. Nonetheless, we observed humans addressing the robot and ascribing social meaning to its actions. As one participant mused afterwards, "It kind of felt like I wasn't doing the experiment alone [...] obviously there's no life force in the robot, but it's participating in the study with you" (P9-M-19).

Addressing the Robot

At several times, participants addressed the robot verbally. Two participants thanked the robot for a good suggestion. Five participants asked the robot questions, like "what are you going to do?", "where are you going?", or "why?" Participants also expressed verbal agreement with the robot, for example, "that was a mistake" when the robot removed a block she had placed, or "I would have done that too."

Ascribing Social Meaning to Robot Actions

We also observed instances where the robot's actions unintentionally carried social meaning for the human. During post-session interviews, three participants described feeling dismissed or ignored based on how the robot acted towards them. One participant (P5-F-19) was frustrated that the robot reversed her move without giving adequate explanation on why. Another wished that the robot would wait and be more conscious about her behaviors, "from my impression [the robot is] trying to get the best result there, but it's pretty much ignoring my things that I had done" (P8-F-21). A third participant (P12-F-21) remarked in the interview that the robot was ignoring her designs.

Conversely, one participant came to believe that the robot was actually listening to her (P3-F-25), saying, "I was so surprised because I didn't know that it was listening to what I was talking about." She referred to the robot moving blocks she placed as "helpful" and that she "felt like it was actually listening to me, because I was thinking out loud [...] about how I'm making my decision, and then it kind of came and organized the blocks."

This may have been related to a series of incidents when the robot did seem to defer or respond to her thoughts about the design or process. At the beginning, she started to ask about removing blocks just as the robot started to do so, stopping short to say, "Oh, I think it's doing it for me." Later on, the robot placed a block that raised cost while increasing scientific benefit; she mentioned this out loud and the robot started its next move by removing the block. Finally, two planning failures unintentionally caused the robot to drop a block just when the participant expressed displeasure with the act of removing them. She turned to look at the robot's hand and laughed, "Yeah [...] leave it there, [that] would be good."

DISCUSSION AND DESIGN IMPLICATIONS

Humans and robots designing together is a relatively unexplored scenario in human-robot interaction and in design studies. In this exploratory study, we used the robot in a rudimentary way; its only function was to choose and implement designs based on variations of human-suggested designs and their outcomes. Our findings highlight the importance of negotiating and coordinating both physical and creative control when a human and a robot design together. They also emphasize social aspects of human-robot collaboration.

Based on these observations, we identify four major considerations for designers of future human-robot collaborative design (HRCd) systems.

Manage the Shared Workspace

The shared workspace is the principal interface point between the human and the robot designers. HRCd systems need to include elements to manage the shared workspace beyond pragmatic design acts. This includes managing turns and considering the spatial arrangement of objects.

Taking turns manipulating designs on the workspace is an important mechanism in the designers' shared process; breakdowns in turn-taking undermine the effectiveness of the

human-robot collaboration. Establishing expectations around how and when turns are taken is thus critical towards building a coherent collaborative design process. This may require more explicit ways for the robot to take or yield a turn, as well as mechanisms for the user to have greater control over the turn-taking process.

In pursuit of a generalizable turn-taking framework in HRI, Thomaz and Chao emphasize the importance and difficulty of relinquishing turns gracefully, and point out that turn-taking dynamics are highly domain-specific and contextual [46]. HRCd systems have to address the complexity in negotiating turns across different stages of design as designers adopt different strategies in the design task.

In addition to turn-taking, the spatial arrangement of shared objects should be considered by HRCd systems, as they can hold epistemic and communicative meaning for the designers beyond their pragmatic function. Participants implicitly and explicitly communicated intentions through block placements, and future HRCd systems should detect and interpret these intentions. The robot could also express its own design intentions through spatial cues.

Outside of the immediate workspace, we observed participants using physical gestures and poses (like stepping back or putting hands in their pockets) to negotiate the process. These could provide important cues to the robot and also a medium for it to signal its internal state to the human.

Communicate About Design Goals

Most design tasks involve a trade-off between conflicting objectives. HRCd systems must consider and perhaps explicitly negotiate preferences about design goals. Alternatively, implicit preferences about designs may be detected through observation and experimentation. A useful HRCd system could, for example, use machine learning techniques to reason about the users' design goals at different times and offer the most appropriate assistance. In our example, participants had to consider the trade-off between science benefit and cost. More open-ended design activities can include other ill-defined and vague goals, such as aesthetic preferences.

Hoffman and Breazeal argue for the importance of establishing and maintaining common goals in human-robot joint activities [20]. Shared design goals are more complex to manage than physical goals, as they afford room for interpretation. In addition, while it may not matter how one achieves a physical goal like opening a box, designers may have implicit preferences about how some desired design outcome is realized. We saw how diverging "opinions" regarding trade-offs between design outcomes can lead to confusion over preferred designs, and an eventual breakdown of trust users have in the robot.

Conflict and compromise are important processes that typify the creative freedom of a design task. While it may not be necessary for the robot and human to share the same implicit preferences about common design goals, it is important that they are aware of each others' preferences and how they evolve, in order to negotiate them together.

Respect Different Design Styles

We saw that participants used different strategies and took on a variety of roles. They also dynamically changed these preferences throughout the design process. HRCD systems should have an understanding of the user's collaborative and cognitive styles as they relate to the design task, and adapt to those styles.

Maher et al. [32] describe three styles of collaboration in design teams: mutual collaboration where teams actively work together; exclusive collaboration where they work in parallel on separate parts, coming together to negotiate and exchange advice; and dictatorship, where one member is in charge.

We observed participants adopting roles that fit different styles of collaboration with the robot, for example, taking charge, letting the robot lead, or working together. Engaging in each of these collaboration styles demands different things of those involved; for example, as Kvan [29] points out, mutual collaboration is more socially exacting and necessitates higher commitment to compromise.

Grigoreanu et al. [16] discuss the benefits of considering user strategy and information processing styles, either comprehensive (breadth first) or selective (depth first), when designing a support tool. In physical design tasks, one's exploration styles become more apparent as cognitive strategies manifest in physical actions.

When an agent can physically manipulate a shared interface together with a human designer, it is important not only to match the user's style, but also to communicate its own strategy with the human. We saw that users interpreted the robot's actions with respect to what they thought its strategy was, and expressed frustration when they could not determine the robot's strategy. We also observed conflicts when users lost confidence in the robot's strategy or felt that the robot was not being mindful of theirs. All of these were factors that led to a loss of trust in the robot.

As collaborators, the human and the robot should be aware of each other's strategy and exploration styles to negotiate the corresponding roles and choices in the design process.

Consider the Social Meaning of Actions

Even our extremely non-social robot design caused social interpretation and responses from users. For example, reversing a human's idea multiple times without explanation could feel dismissive to a user.

Jung's work on affective grounding [22] underlines the importance of establishing common ground about the emotional meaning of actions. This brings to mind the increased social risk that is carried by engaging in both design collaboration, as described by Kvan [29] and physical interaction, as described by Klemmer et al. [27]. This increased social risk can be a boon when it causes partners to act more thoughtfully, but ignoring it can lead to breakdowns.

In light of this, it is important for a robot to understand and consider the social implications of both agents' actions taken

with respect to each others' ideas and design decisions, and establish shared expectations to avoid social harm.

This is particularly salient because the three design considerations listed above involve additional social behavior on the robot's part, such as negotiating turns or expressing design preferences. Addressing these for HRCD systems requires consideration of how to do so in socially appropriate ways.

LIMITATIONS AND FUTURE WORK

The purpose of this work is to begin to characterize the relatively uncharted design space of human-robot collaborative design. As such, the interactive behaviors of the robot we studied were deliberately basic. Future work should measure how the considerations we observed affect collaborative design outcomes, as well as propose and evaluate approaches to address them.

In addition, configuration design is only one, arguably routine, archetype of design tasks and one "ingredient" in a broader design process [9]. Design tasks that involve more ambiguity and room for reformulation could give rise to richer interactions that further complicate the coordination considerations we discuss. Further, all our participants were novices at the design task, while prior work suggests that novices and experts approach design differently [2] and may interact differently with computer partners [6]. Future work will explore other types of design problems in the context of HRCD and with users spanning a broader range of expertise.

CONCLUSION

We present an observational study of humans working collaboratively on a design task with a robot, using a tangible interface. Humans and AI systems have distinct abilities, making them well-suited to undertake design tasks together. The added benefits of physically embodied collaboration in design suggests that robots may be a useful instantiation of an AI design partner. While a rich literature addresses human-robot collaboration in physical tasks like assembly, the idea of robots as creative partners in design tasks is less explored.

Based on our study, we note that users actively negotiate not only the physical design space—through turn-taking and spatial arrangements of design components—but also the creative process—through roles, goals, and strategies. Moreover, participants in our study often interpreted the robot's pragmatic actions in social ways and addressed the robot as a social agent. Based on these observations, we identify four considerations for future human-robot collaborative design systems. These systems should pay attention to the management of the shared space, communicate about design goal preferences, respect diverse design styles, communicate the robot's own strategies, and do all of the above in light of the implicit social meanings of design acts.

ACKNOWLEDGMENTS

The authors would like to thank Nikhil Dhawan for his significant contributions to the project. This work was supported primarily by the Civil, Mechanical and Manufacturing Innovation Program of the National Science Foundation under NSF Award No. 1635253.

REFERENCES

- [1] Henny Admoni, Thomas Weng, Bradley Hayes, and Brian Scassellati. 2016. Robot Nonverbal Behavior Improves Task Performance in Difficult Collaborations. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction*. IEEE Press, 51–58.
- [2] Saeema Ahmed, Ken M Wallace, and Lucienne T Blessing. 2003. Understanding the Differences Between How Novice and Experienced Designers Approach Design Tasks. *Research in engineering design* 14, 1 (2003), 1–11.
- [3] Saleema Amershi, James Fogarty, and Daniel Weld. 2012. Regroup: Interactive Machine Learning for On-Demand Group Creation in Social Networks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 21–30.
- [4] Ernesto Arias, Hal Eden, Gerhard Fischer, Andrew Gorman, and Eric Scharff. 2000. Transcending the Individual Human Mind—Creating Shared Understanding Through Collaborative Design. *ACM Transactions on Computer-Human Interaction (TOCHI)* 7, 1 (2000), 84–113.
- [5] Meghna Babbar-Sebens and Barbara S Minsker. 2012. Interactive Genetic Algorithm with Mixed Initiative Interaction for Multi-Criteria Ground Water Monitoring Design. *Applied Soft Computing Journal* 12, 1 (2012), 182–195.
- [6] Bobby Beaton, Steve Harrison, and Deborah Tatar. 2010. Digital drumming: a study of co-located, highly coordinated, dyadic collaboration. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1417–1426.
- [7] Stephen J Beckett. 2017. The Logic of the Design Problem: A Dialectical Approach. *Design Issues* 33, 4 (2017), 5–16.
- [8] Virginia Braun and Victoria Clarke. 2006. Using Thematic Analysis in Psychology. *Qualitative Research in Psychology* 3, 2 (2006), 77–101.
- [9] David C Brown. 1998. Defining Configuring. *AI EDAM* 12, 4 (1998), 301–305.
- [10] Justine Cassell, M Ananny, A Basu, Timothy Bickmore, P Chong, D Mellis, Kimiko Ryokai, J Smith, H Vilhjálmsson, and H Yan. 2000. Shared Reality: Physical Collaboration with a Virtual Peer. In *CHI'00 Extended Abstracts on Human Factors in Computing Systems*. ACM, 259–260.
- [11] Sung-Bae Cho. 2002. Towards Creative Evolutionary Systems With Interactive Genetic Algorithm. *Applied Intelligence* 16, 2 (2002), 129–138.
- [12] David Johan Christensen, Rune Fogh, and Henrik Hautop Lund. 2014. Playte, a Tangible Interface for Engaging Human-Robot Interaction. In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*. 56–62.
- [13] Ernest A Edmonds, Linda Candy, Rachel Jones, and Bassel Soufi. 1994. Support for Collaborative Design: Agents and Emergence. *Commun. ACM* 37, 7 (1994), 41–47.
- [14] Paul Egan and Jonathan Cagan. 2016. Human and Computational Approaches for Design Problem-Solving. In *Experimental Design Research*. Springer, 187–205.
- [15] Gerhard Fischer. 2004. Social Creativity: Turning Barriers into Opportunities for Collaborative Design. In *Proceedings of the Eighth Conference on Participatory Design: Artful Integration: Interweaving Media, Materials and Practices—Volume 1*. ACM, 152–161.
- [16] Valentina I Grigoreanu, Margaret M Burnett, and George G Robertson. 2010. A Strategy-Centric Approach to the Design of End-User Debugging Tools. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 713–722.
- [17] Cheng Guo and Ehud Sharlin. 2008. Exploring the Use of Tangible User Interfaces for Human-robot Interaction: A Comparative Study. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2008), 121–130.
- [18] Pamela J Hinds, Teresa L Roberts, and Hank Jones. 2004. Whose Job Is It Anyway? A Study of Human-Robot Interaction in a Collaborative Task. *Human-Computer Interaction* 24 (2004), 151–181.
- [19] Nozomi Hitomi, Hyunseung Bang, and Daniel Selva. 2017. Extracting and Applying Knowledge with Adaptive Knowledge-Driven Optimization to Architect an Earth Observing Satellite System. *AIAA Information Systems-AIAA Infotech@ Aerospace* (2017), 0794.
- [20] Guy Hoffman and Cynthia Breazeal. 2004. Collaboration in Human-Robot Teams. In *AIAA 1st Intelligent Systems Technical Conference*. 6434.
- [21] Guy Hoffman and Gil Weinberg. 2011. Interactive Improvisation with a Robotic Marimba Player. *Autonomous Robots* 31, 2-3 (2011), 133–153.
- [22] Malte F Jung. 2017. Affective Grounding in Human-Robot Interaction. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 263–273.
- [23] Peter H Kahn Jr, Takayuki Kanda, Hiroshi Ishiguro, Brian T Gill, Solace Shen, Jolina H Ruckert, and Heather E Gary. 2016. Human Creativity Can be Facilitated Through Interacting with a Social Robot. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction*. IEEE Press, 173–180.
- [24] Martin Kaltenbrunner and Ross Bencina. 2007. reacTIVision: a Computer-Vision Framework for Table-Based Tangible Interaction. In *Proceedings of the 1st International Conference on Tangible and Embedded Interaction*. ACM, 69–74.

- [25] Kazuhiko Kawamura, Sugato Bagchi, Moenes Iskarous, and Magued Bishay. 1995. Intelligent Robotic Systems in Service of the Disabled. *IEEE Transactions on Rehabilitation Engineering* 3, 1 (1995), 14–21.
- [26] Mi Jeong Kim and Mary Lou Maher. 2008. The Impact of Tangible User Interfaces on Spatial Cognition During Collaborative Design. *Design Studies* 29, 3 (2008), 222–253.
- [27] Scott R Klemmer, Björn Hartmann, and Leila Takayama. 2006. How Bodies Matter: Five Themes for Interaction Design. In *Proceedings of the 6th Conference on Designing Interactive Systems*. ACM, 140–149.
- [28] Scott R Klemmer, Mark W Newman, Ryan Farrell, Mark Bilezikjian, and James A Landay. 2001. The Designers' Outpost. *Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology - UIST '01* May 2001 (2001), 1.
- [29] Thomas Kvan. 2000. Collaborative Design: What is it? *Automation in Construction* 9, 4 (2000), 409–415.
- [30] Matthew V Law, Nikhil Dhawan, Hyunseung Bang, So-Yeon Yoon, Daniel Selva, and Guy Hoffman. 2018. Side-by-side Human-Computer Design using a Tangible User Interface. In *Eighth International Conference On Design Computing and Cognition*.
- [31] Antonios Liapis, Georgios N Yannakakis, Constantine Alexopoulos, and Phil Lopes. 2016. Can Computers Foster Human Users' Creativity? Theory and Praxis of Mixed-Initiative Co-Creativity. *Digital Culture & Education* 8, 2 (2016), 136–153.
- [32] Mary Lou Maher, Anna Cicognani, and Simeon Simoff. 1996a. An Experimental Study of Computer Mediated Collaborative Design. In *Enabling Technologies: Infrastructure for Collaborative Enterprises, 1996. Proceedings of the 5th Workshop on*. IEEE, 268–273.
- [33] Mary Lou Maher, Josiah Poon, and Sylvie Boulanger. 1996b. Formalising Design Exploration as Co-Evolution. In *Advances in Formal Design Methods for CAD*. Springer, 3–30.
- [34] Claire Mikalauskas, Tiffany Wun, Kevin Ta, Joshua Horacsek, and Lora Oehlberg. 2018. Improvising with an Audience-Controlled Robot Performer. In *Proceedings of the 2018 on Designing Interactive Systems Conference 2018*. ACM, 657–666.
- [35] Sanjay Mittal and Felix Frayman. 1989. Towards a Generic Model of Configuraton Tasksex. In *IJCAI*, Vol. 89. 1395–1401.
- [36] Gary M Olson and Judith S Olson. 2009. Human-Computer Interaction Distance Matters. *International Journal of Human Computer Interaction* 0024, July 2014 (2009), 37–41.
- [37] Huaishu Peng, Jimmy Briggs, Cheng-Yao Wang, Kevin Guo, Joseph Kider, Stefanie Mueller, Patrick Baudisch, and François Guimbretière. 2018. RoMA: Interactive Fabrication with Augmented Reality and a Robotic 3D Printer. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 579.
- [38] Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Y Ng. 2009. ROS: an Open-Source Robot Operating System. In *ICRA Workshop on Open Source Software*, Vol. 3. Kobe, Japan, 5.
- [39] Kinova Robotics. 2018a. Kinova Gen2 Ultra Lightweight Robotic Arm User Guide. *Online at <https://www.kinovarobotics.com/en/knowledge-hub/all-kinova-products>* (2018).
- [40] Kinova Robotics. 2018b. Kinova-ROS. *Online at <https://github.com/Kinovarobotics/kinova-ros>* (2018).
- [41] Bertrand Schneider, Patrick Jermann, Guillaume Zufferey, and Pierre Dillenbourg. 2011. Benefits of a Tangible Interface for Collaborative Learning and Interaction. *IEEE Transactions on Learning Technologies* 4, 3 (2011), 222–232.
- [42] Daniel Selva and Edward F Crawley. 2013. VASSAR: Value Assessment of System Architectures Using Rules. In *Aerospace Conference, 2013 IEEE*. IEEE, 1–21.
- [43] Julie Shah, James Wiken, Brian Williams, and Cynthia Breazeal. 2011. Improved Human-Robot Team Performance Using Chaski, a Human-Inspired Plan Execution System. In *Proceedings of the 6th International Conference on Human-Robot Interaction*. ACM, 29–36.
- [44] Gillian Smith, Jim Whitehead, and Michael Mateas. 2010. Tanagra: A Mixed-Initiative Level Design Tool. In *Proceedings of the Fifth International Conference on the Foundations of Digital Games*. ACM, 209–216.
- [45] Ioan A Sucas and Sachin Chitta. 2013. MoveIt! *Online at <http://moveit.ros.org>* (2013).
- [46] Andrea L Thomaz and Crystal Chao. 2011. Turn-Taking Based on Information Flow for Fluent Human-Robot Interaction. *AI Magazine* 32, 4 (2011), 53–63.
- [47] Lauren Vasey, Long Nguyen, Tovi Grossman, Heather Kerrick, Tobias Schwinn, David Benjamin, Maurice Conti, and Achim Menges. 2016. Human and Robot Collaboration Enabling the Fabrication and Assembly of a Filament-Wound Structure. In *ACADIA//2016: Posthuman Frontiers: Data, Designers, and Cognitive Machines, Proceedings of the 36th Annual Conference of the Association for Computer Aided Design in Architecture*. 184–195.
- [48] Vighnesh Vatsal and Guy Hoffman. 2017. Wearing Your Arm on Your Sleeve: Studying Usage Contexts for a Wearable Robotic Forearm. In *Robot and Human Interactive Communication (RO-MAN), 2017 26th IEEE International Symposium on*. IEEE, 974–980.

- [49] Antonio P Volpentesta, Salvatore Ammirato, and Francesco Sofo. 2009. Thinking Style Diversity and Collaborative Design Learning. *IFIP Advances in Information and Communication Technology* 307 (2009), 785–796.
- [50] Joshua Wainer, David J Feil-Seifer, Dylan A Shell, and Maja J Mataric. 2007. Embodiment and Human-Robot Interaction: A Task-Based Perspective. In *Robot and Human interactive Communication, 2007. RO-MAN 2007. The 16th IEEE International Symposium on.* IEEE, 872–877.
- [51] Manuela Waldner, Jörg Hauber, Jürgen Zauner, Michael Haller, and Mark Billinghurst. 2006. Tangible Tiles: Design and Evaluation of a Tangible User Interface in a Collaborative Tabletop Setup. *OZCHI, ACM Press* (2006), 1–8.