

# Design Patterns for Exploring and Prototyping Human-Robot Interactions

Allison Sauppé and Bilge Mutlu

Department of Computer Sciences, University of Wisconsin–Madison  
1210 West Dayton Street, Madison, WI 53706 USA  
asauppe@cs.wisc.edu, bilge@cs.wisc.edu

## ABSTRACT

Robotic products are envisioned to offer rich interactions in a range of environments. While their specific roles will vary across applications, these products will draw on fundamental building blocks of interaction, such as greeting people, narrating information, providing instructions, and asking and answering questions. In this paper, we explore how such building blocks might serve as *interaction design patterns* that enable design exploration and prototyping for human-robot interaction. To construct a pattern library, we observed human interactions across different scenarios and identified seven patterns, such as question-answer pairs. We then designed and implemented *Interaction Blocks*, a visual authoring environment that enabled prototyping of robot interactions using these patterns. Design sessions with designers and developers demonstrated the promise of using a pattern language for designing robot interactions, confirmed the usability of our authoring environment, and provided insights into future research on tools for human-robot interaction design.

## Author Keywords

Design patterns; interaction design; human-robot interaction; authoring environment; design exploration; prototyping; design sessions

## ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: User Interfaces—*input devices and strategies, evaluation/methodology, user-centered design*

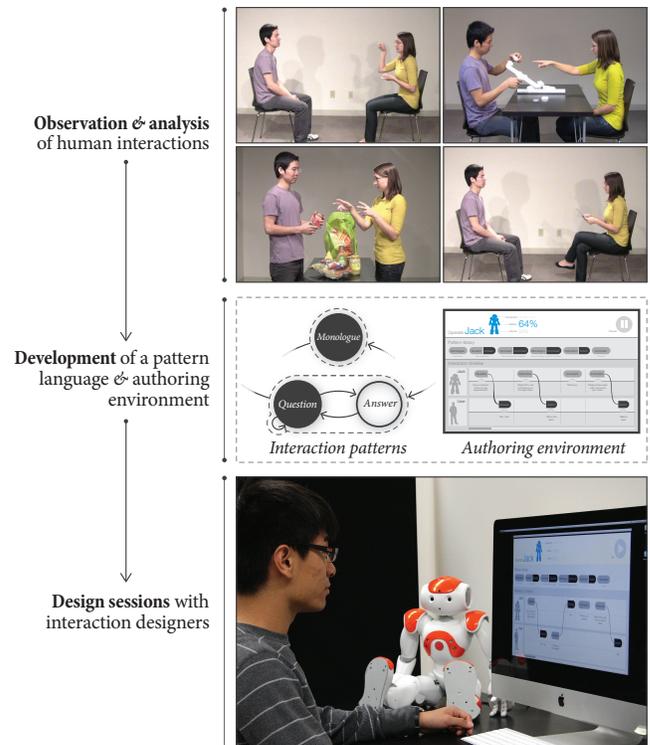
## INTRODUCTION

As robots promise to become commonplace products in a range of settings, serving as receptionists, museum guides, and tutors for children, designers will need materials and tools that will enable them to explore and prototype a range of interactions that robots will offer in these settings. For example, a receptionist robot serving at a doctor’s office may work with a patient to schedule an appointment, instruct a patient on how to use a medication, or inform medical staff about

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. To copy otherwise, or to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI’14, April 26–May 1, 2014, Toronto, Canada.

Copyright 2014 ACM 978-1-4503-1015-4/12/05...\$10.00.



**Figure 1.** In this work, we use observations and analysis of human interactions to develop an interaction design pattern language and authoring environment called *Interaction Blocks* to enable designers to explore and prototype human-robot interactions.

upcoming appointments. While interactions across settings or within a setting may take many forms, they will build on a fundamental set of communicative acts that humans draw on in communication. In this paper, we explore how these communicative acts might serve as interaction design patterns and how we might use a pattern language to enable design exploration and prototyping of human-robot interaction.

Do human interactions follow specific patterns? If so, how can we create a pattern language for human interaction? How can we enable designers to use such a language for exploration and prototyping human-robot interactions? Goffman [11] argued that human interactions follow a specific “order” and characterized a number of patterns in which people interact, such as how greetings unfold and how people leave an interaction. Research across sociology, psychology, and HCI

have sought to identify similar patterns in different forms of interaction, such as interviews [19], conversations [29], and collaborations [5]. Kahn et al. [15] have proposed that the patterns observed in human interactions might serve as designing interactions for robots, describing eight patterns such as “initial interaction” and “in motion together” based on their observations of how children interacted with robots and literature on fundamental constructs in human interactions.

In this paper, we extend this concept toward building a pattern language that might serve as building blocks for human-robot interactions across different forms of interaction, focusing on five scenarios in which robots are expected to engage, including collaboration [8], conversation [12], instruction [13], interviews [9], and storytelling [20]. We present a formative study of human interactions across these five scenarios and describe seven patterns that appear across these scenarios. These patterns informed the design and implementation of *Interaction Blocks*, an authoring environment that enables users to explore and create interactions for social robots. We envision Interaction Blocks having three primary groups of users: (1) interaction designers who have little programming background, (2) designers or programmers who would like to rapidly prototype interactions, and (3) designers or programmers who might not have experience in working with human behavior as a design element. Through a qualitative evaluation with interaction designers and developers, we demonstrate the feasibility of the use of a pattern language for designing human-robot interactions and the usability of our authoring environment. Figure 1 illustrates this process.

The next section provides background on patterns in human interactions and reviews related work on interfaces for authoring robot behaviors. This section is followed by a description of the formative study that informed the development of a pattern language for human-robot interactions and the design of *Interaction Blocks*. Following sections discuss the design and implementation of Interaction Blocks, as well as its evaluation through a series of design sessions. We conclude the paper with a discussion of the findings and their implications for future research into designing human-robot interactions.

## RELATED WORK

In the paragraphs below, we review related work on patterns in human interactions, the use of patterns for designing human-robot interactions, and authoring tools and environments developed for programming robot behaviors.

### Patterns in Human Interactions

Human interactions follow an invisible structure, a shared *interaction order* [11], that signals to its participants how they should act and interact with others. For example, the openings of encounters follow a particular “routine” that involves a greeting or a summons by a participant and an in-kind response either in the form of a greeting or an answer [27]. Similarly, conversations might involve the pattern “question-answer pairs,” where one participant is posing a question followed by an answer from a different participant [6]. These routines may be combined to produce a more elaborate interaction, as shown in the example below.

*Participant 1:* <upon seeing a friend at work> Oh, good morning! Are you attending the meeting this afternoon?

*Participant 2:* Yes, I was planning on going.

In this brief encounter, Participant 1 opens the interaction with a greeting and subsequently poses a question, which Participant 2 answers. The participants of this encounter employ both the greeting and question-answer pair routines to seek and provide information and coordinate activities. Research on human communication has shown that the use of such routines or “interaction patterns” facilitates effective communication, specifically processes such as fluency [29] and grounding [5]. On the other hand, breakdowns occur in communication when these patterns are violated [27].

### Patterns in Human-Robot Interactions

Research on human-robot interactions has also considered how robots might display patterned behaviors, such as gaze patterns that facilitate conversational turn-taking and joint attention [13], to enable more effective human-robot interaction. Researchers have also proposed the use of such patterned behaviors as building blocks for constructing human-robot interactions [15, 22], such as an “initial interaction” pattern proposed by Kahn et al. [15]—an interaction pattern analogous to greetings in human interactions. However, this body of work does not yet offer a pattern language that might serve as building blocks for interactions across a wide range of scenarios. Furthermore, this research has not developed tools or environments to support design exploration or prototyping for constructing human-robot interactions.

### Authoring Environments for Robot Behaviors

A complementary body of work has developed several authoring tools and environments to help expert developers and novice users better control and interact with robots [14] and to evaluate human-robot interactions [16]. Environments for developers include the Robot Operating System (ROS), which offers an architecture for abstracting and reusing specific functionalities across different robot platforms [24]. A particular ROS module, the Robot Behavior Toolkit, offers the ability to specify robot social behaviors based on a repository of “rules” or patterns [13]. However, the use of these environments requires a substantial amount of development expertise as well as effort to build suitable pattern repositories.

Authoring environments for novice users include a number of commercial and research tools for programming robot behaviors [10]. For example, Interaction Composer (IC) involves a graphical interface that enables users to coordinate multiple facets of a robot’s behavior, such as dialogue and gestures, by choosing from a set of “behavior blocks” to compose the interaction [10]. RoboStudio offers an expert authoring environment to build graphical interfaces that enable novice users to customize the behaviors of the robot to their needs or preferences [7]. Finally, Lohse et al. [18] has developed a framework for iterative design and evaluation of robot behaviors [18]. While these tools and environments involve easy-to-use interfaces that are accessible to both expert and novice users, they require the designer to have knowledge of a pattern language and to map this knowledge to specific robot behaviors that can be authored using the tools.



**Figure 2.** Examples of the experimental setup for all five scenarios. From left to right, participants are sharing in a storytelling experience, engaging in conversation, conducting an interview, learning how to configure a set of pipes, and collaborating on how to sort foodstuffs.

## FORMATIVE STUDY

To build a pattern language that enables design exploration and prototyping for human-robot interactions, we conducted a formative study of human interactions that involved observations of human interactions and identifying and modeling patterns in which these interactions unfolded. While a variety of behaviors and scenarios have been studied by linguists and psychologists [6, 11, 27], the results have not been constructed into design patterns that can be translated and implemented on a robot. To achieve natural, humanlike robot behaviors, we chose to ground our development of interaction design patterns in observations and detailed analyses of human interactions. We collected data from eight dyads performing in five social interaction scenarios that are representative of the types of interactions robots are envisioned to encounter in their future roles in society: *conversation*, *collaboration*, *instruction*, *interview*, and *storytelling*. We then followed an iterative modeling process to construct models that represented each of the scenarios. These models revealed common interaction design patterns that appeared across multiple scenarios. The paragraphs below detail and discuss the process of discovering and formalizing these patterns.

## Data Collection

### Participants and Data Corpus

A total of 16 native-English-speaking participants from the University of Wisconsin–Madison took part in this study. Participants studied a diverse range of fields, and their ages ranged 19–62 ( $M = 25.44$ ,  $SD = 10.39$ ). Participants were assigned into dyads to jointly interact in the social interaction scenarios. We randomly assigned participants into dyads and conversational roles such that they were fully stratified by gender. The instructor and storyteller roles were never held by the same participant, allowing the instructor to learn the task to be instructed while the storyteller learned the story. The scenarios were executed in the same order for each dyad.

We used a single video camera equipped with a wide-angle lens to capture the entire upper body of the participants for all scenarios. The final data corpus consisted of five scenarios for each dyad, which amounted to 3 hours and 31 minutes of audio and video data. The average lengths for the collaboration, conversation, instruction, interview, and storytelling scenarios were 3:47, 4:47, 5:02, 7:32, and 3:59, respectively.

### Procedure

After arrival in our laboratory, a researcher provided the participants with an overview of the interaction scenarios. Following informed consent, the participants were then led into

a usability laboratory, which was initially configured for the first scenario in which they would be participating. The room was reconfigured for each scenario, as shown in Figure 2. Following the completion of all five scenarios, participants completed a demographic questionnaire. The researchers then debriefed both participants. The five scenarios and post-experiment questionnaire took approximately one hour. Participants were compensated \$10 each for their time.

### Interaction Scenarios

We designed five interaction scenarios that showed the characteristics of the scenarios in which robots are expected to interact with people and that followed previous modeling research, which will be discussed in more detail below. Each scenario was intended for small group interaction. For data collection, we engaged two participants in each interaction as the minimum number required to realize the scenario.

1) *Conversation*: In the conversation scenario, we aimed to observe how participants would handle a set of unstructured exchanges about a given topic. Framing, turn-taking, and recovering from errors are all important aspects of a successful conversation [29]. In particular, participants who have not yet established any level of rapport may favor question-answer pairs, rather than conversational dialogue [29].

In the scenario, the participants discussed their educational experiences and goals. They were instructed to continue this conversation until the researcher re-entered the room when the conversation naturally subsided.

2) *Collaboration*: The collaboration scenario targeted gaining a better understanding of how participants collaborate toward a common goal. Previous research has shown that joint activity requires flexibility and communication in order to effectively coordinate differing opinions from participants [5].

In our collaboration scenario, the participants worked together to sort six grocery bags of foodstuffs onto two tables. The tables were divided into three areas meant to represent common food storage locations in the home: the *pantry*, the *fridge*, and the *countertop*. These areas were further subdivided into areas based on types of food. Participants were asked to place the empty grocery bags on the table to indicate that they were done.

3) *Instruction*: In the instruction scenario, we aimed to observe how an expert conveys knowledge to a non-expert and, when that knowledge is miscommunicated, how the expert might help the non-expert to recover [28].

Our instruction scenario involved first training one of the participants (the *instructor*) in assembling a particular pipe configuration to allow two sinks to drain into a single system. The instructor was given as much time as they needed to learn how to configure the pipes. Upon learning the necessary steps to build the pipes, the instructor walked the second participant (the *student*) through each step of the assembly.

4) *Interview*: The goal of the interview scenario was to capture the process by which one participant questioned a second participant to obtain information. Interviewers may sometimes request sensitive information, requiring substantial rapport between the interviewer and interviewee. The structure and the tone of an interview can help determine the level of rapport between participants [19].

In our interview scenario, one of the participants (the *interviewer*) was told that they would be conducting a job interview for a generic position. The interviewer was given a list of 14 questions and time to review them. After reviewing the questions, the second participant (the *interviewee*) was asked all 14 questions by the interviewer. The researcher reentered the room when all the questions were answered.

5) *Storytelling*: In storytelling, we were interested in observing how one participant relayed a story to a second participant and what discourse patterns (e.g., asking questions) helped participants communicate more effectively [17].

In our storytelling scenario, one participant (the *storyteller*) watched a seven-minute video of a cartoon story. The storyteller was given as much time as they needed to feel comfortable with retelling the story, after which they had three to five minutes to retell the story to the second participant (the *listener*). After the storyteller finished, participants frequently had a conversation concerning the story. The experimenter re-entered the room when the conversation subsided.

## Analysis

For each scenario, we followed an iterative process of data coding and modeling. The coding process involved a researcher iteratively coding all video data and a second researcher annotating 10% of the videos to confirm the reliability of the coding process (81% agreement, Cohen's  $\kappa = .76$ ). In the coding of the data, we annotated the videos for the set of *states*—significant events in the interaction—and the *transitions* between these states, generating a *model* of the interaction scenario. From these models, we extracted what appeared as the core part of the interaction: the states that were essential to characterize the interaction. For example, in an interview, the interviewer asking a question and an interviewee answering it serve as two core states of the interaction. Finally, we examined the models from all five scenarios for states or sequences of states that frequently appeared to identify what might serve as *interaction design patterns*.

Our consideration of significant events in the interaction as states follows prior studies of social interaction [5, 19, 29]. When connected with other states, the resulting model represents the flow of the interaction from one event to the next. The idea of states translates well to social interaction, where one participant usually holds the floor and is engaged in an

event (e.g., sharing a comment) [25]. For the purposes of an interaction with multiple participants, the states of all participants can be mapped to a single model. When a state that is held by only one participant is entered, all other participants implicitly wait. The use of states and transitions also fits well with a common paradigm used in robot programming where “nodes” are used to control functionalities of the robot [24]. The use of states provides a flexible representation for the flow of interaction for multiple participants.

After establishing the core states, we reviewed each video and noted any deviations from this core part of the interaction in both kind and number. An example deviation in the interview scenario is the interviewee asking for clarification regarding the interviewer's question. State models for each scenario were then constructed from the core part of the interaction, any deviations from the core part of the interaction, and notes collected from the video data. The resulting scenario models were then compared against the interactions in the videos for any mistakes or inconsistencies.

## Pattern Discovery

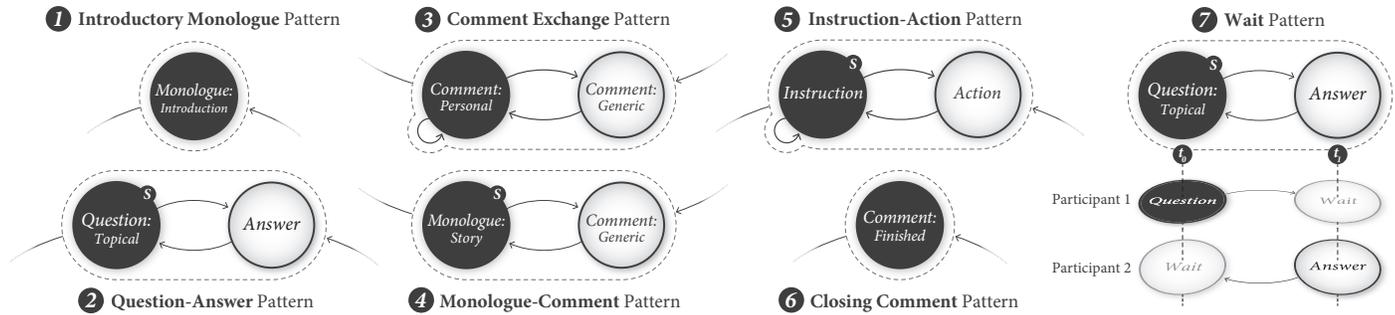
After constructing a model for each scenario, we identified common interaction structures, which served as *design interaction patterns*, or patterns. For example, a question being asked and answered is comprised of two separate states (a question state and an answer state). However, questions are almost always followed by answers, and thus the interaction between both states is codified into a pattern. We identified seven common patterns (Figure 3) across all five scenarios:

1) *Introductory Monologue*: The introduction that begins the interaction is often the most important indicator of how well the remainder of the interaction will play out [26]. A short introduction can be used to introduce other participants to a scenario by giving an overview of the remainder of the interaction. In the interview scenario, some participants started with a short introduction of themselves, describing the interview as part of the “hiring process.” In the storytelling scenario, all participants started with a short introduction that set the stage for the story, such as identifying the primary characters and setting. An example of an introduction used in the interview scenario is shown below:

*Interviewer*: Hi, welcome. ... So, today I'll be asking you a few questions to gauge your compatibility for this job.

Previous work on spontaneous encounters notes that these introductions can take different forms [27]. It is likely that for scenarios where we did not observe an introduction state (conversation, collaboration, and instruction), an introduction state would have occurred in a natural setting. For example, a person may pass by their friend while shopping and summon the friend with a short introduction (“Hello Jane! How are you doing?”) before commencing with a conversation.

2) *Question and Answer*: The idea of a *question-answer pair* is a well-studied component of discourse management [6]. A *question* is a sentence meant to elicit information from other participants. Questions may introduce a topic of interest (e.g.,



**Figure 3.** Models for the seven patterns we discovered. The dark and light-colored states indicate when one participant occupies the dark-colored state, a second participant must occupy the light-colored state. For example, for pattern 5, if one participant asks a question, the second participant gives an answer.

“What has been your most difficult job?”), request elaboration (e.g., “Could you elaborate on that?”), ask for clarification (e.g., “What do you mean by that?”), confirm understanding (e.g., “Does that make sense?”), check background (e.g., “Have you ever watched Looney Tunes?”), and request status (e.g., “How are you?”). In an archetypal question-answer pair, a question is followed by an answer. An *answer* is the response to a question that aims to satisfy the questioning participant’s curiosity [27]. The excerpt below from the storytelling scenario is an example of a question-answer pattern:

*Storyteller:* Do you know who Marvin the Martian is?  
*Listener:* Oh yeah, from Looney Tunes.

3) *Generic Comment and Personal Comment:* A *comment* is a short statement offering the speaker’s opinion. Comments are either generic (e.g., “Wow”) or personal (e.g., “I tried that and didn’t like it”). In our data, participants engaged in exchanging comments frequently move between sharing personal insights—either of their own volition or after being prompted—and offering generic comments. The following excerpt illustrates an exchange of comments from the conversation scenario:

*Participant 1:* Wow.  
*Participant 2:* Yeah...I had never done anything quite like that before.  
*Participant 1:* I had a similar experience once, but it wasn’t nearly that exciting.  
*Participant 2:* Interesting.

In this example, both participants offer both generic and personal comments, highlighting the fluidity between these two types of comments for all participants of the interaction.

4) *Monologue and Generic Comment:* A *monologue* is a longer form of speech during which no response is expected. Monologues may involve the telling of a story (e.g., “...Once Marvin had reached Bugs Bunny, he chose to...”). Although monologues expect no response, listeners may occasionally offer unsolicited commentary, as illustrated by the excerpt below from the storytelling scenario:

*Storyteller:* ...and then, all of a sudden, thousands of these aliens appear on Earth.  
*Listener:* That’s a lot.

5) *Instruction and Action:* An *instruction* is a command offered by one participant to direct the actions of another participant. The proper response to this instruction is often an *action*, although the action might follow the instruction with a delay depending on whether it is an appropriate time to perform that action [4]. Instruction-action pairs are commonly found in teaching scenarios where the teacher is directing the student. Below is an example of an instruction-action pair from the instruction scenario:

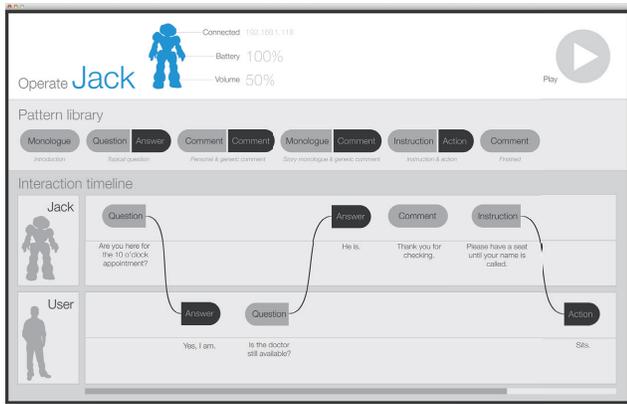
*Instructor:* Now connect the long pipe with the one shaped like an “S”.  
*Student:* <locates both the long and S-shaped pipe, and then connects them>

6) *Finished Comment:* Upon the completion of the goals of the scenario, one or more of the participants will note that the scenario is completed by offering a *finished comment*. For some scenarios (the interview, instruction, and storytelling scenarios), only one of the participants is able to end the scenario (e.g., only the storyteller knows when the story is finished). In the collaborative scenario, either participant is able to end the scenario, as illustrated in the following example:

*Participant:* I think that’s it, so we should be done.

Previous work has shown that conversations frequently have a definite ending initiated by either participant, whether through an unexpected interruption (e.g., a phone call), a forced ending (e.g., a train stop forces participants to go separate ways), or an achievement of the goals of the conversation (e.g., obtaining some piece of information) [21]. However, our instantiation of a conversation scenario lacked a comment confirming the end of the interaction, due to the conversation scenario in our study providing no concrete end goal, except to converse for three to five minutes until the experimenter interrupted the conversation at a natural break.

7) *Wait:* One pattern implicit in all scenarios involving two or more participants is the *wait* pattern. The majority of states across our scenarios are intended for a single participant. When a participant transitions into a state intended for a single participant, all other participants enter a wait state, as shown in Figure 3. Data on conversations overwhelmingly supports only a single speaker at a time and other participants listening to the speaker, with multiple speakers being common but brief [25].



**Figure 4.** A screenshot of *Interaction Blocks*, our authoring environment for using a pattern language for synthesizing human-robot interactions.

Our analysis of five common social interactions not only provides a deeper understanding of each scenario, but also reveals the prevalence of a number of patterns across scenarios. These patterns confirmed some previous patterns (e.g., question-answer pairs [6]) as well as discovered new ones (e.g., Generic Comment and Personal Comment).

## DESIGN & IMPLEMENTATION OF INTERACTION BLOCKS

Our analysis of human interactions revealed seven core patterns that appear across the five common social scenarios. To draw on these patterns in design exploration and prototyping for human-robot social interactions, we developed *Interaction Blocks*, an authoring environment that enabled interaction designers to synthesize interactions and prototype them on a NAO humanoid robot, a robot platform commonly used in research and design for human-robot interaction.

### Authoring Environment

Our authoring environment is composed of three sections: the *control panel*, the *pattern library*, and the *interaction timeline*. An example of our tool can be seen in Figure 4.

The *control panel* displays pertinent information to the user. The silhouette of the robot provides a visual indication of connection status, with blue indicating “connected” and grey indicating “disconnected.” If connected, the control panel also informs the user of which robot they are connected to, and the IP address, current battery level, and current volume of the connected robot. The “Play” button on the far right side of the interface is used to upload and execute the user’s synthesized dialogue on the connected robot.

The *pattern library* contains the patterns discovered in our formative study, with each pattern represented in a capsule shape. We chose to represent our patterns in a capsule shape to accommodate dual-colored capsules for those patterns which are comprised of two states. To compose an interaction, patterns can be dragged out of the pattern library and onto the interaction timeline.

The *interaction timeline* is where the user composes the interaction. To accommodate both the user and the robot roles, the timeline is divided in half, each half assigned to one of

the roles. When the user moves a pattern from the pattern library to the interaction timeline, the pattern is automatically added in place to the appropriate role. For patterns with two states, the movement of the pattern into the interaction timeline causes the pattern to divide in half, as shown in Figure 4. A Bézier curve connects the two states, indicating that no patterns can be inserted between them. When a pattern is divided into two states, the user is in control of the first state in the pattern. The second state is automatically added to the opposite role. For example, if a “question” is added to the user’s timeline, an “answer” will be added to the robot’s timeline.

### Implementation on the NAO Robot

To allow users to evaluate their synthesized interaction on a robot, we enabled *Interaction Blocks* to connect and execute the resulting dialogue on a NAO robot. All robot utterances were generated through the NAO’s native text-to-speech application programming interface (API), while participant responses were recognized using the NAO’s native natural language processing API.

In addition to the dialogue, socially appropriate gaze behaviors were automatically incorporated into the robot’s interaction. The NAO employed its native face-tracking capabilities to enable consistent gaze with the participant. We introduced Perlin noise [23]—a technique used in animation and film to simulate randomness that appears natural—to create small head shifts in the robot and create a lifelike appearance. Additionally, we constructed socially appropriate gaze aversion behaviors for the robot, using the timings provided by Andrist et al. [1]. For instance, when the floor was passed to the robot, the robot would gaze away before returning its gaze to the participant, resuming Perlin noise, and continuing with its part of the dialogue, all of which helped the robot display natural social behaviors.

### DESIGN SESSIONS

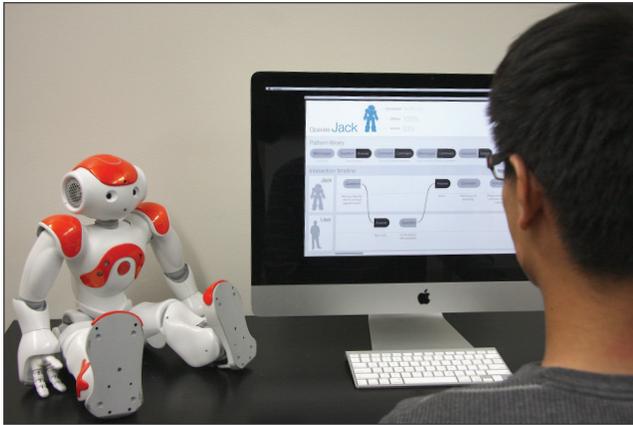
We evaluated our tool in design sessions with local members of the design and development community from two groups: *interaction designers* and *developers*. Using a scenario for an interaction between a robot receptionist and human patient, participants were asked to construct and test an episode of interaction using our authoring environment.

### Design Task

Participants were given a scenario that described an exchange at a dentist’s office between a patient (the user) and the receptionist (the robot). The scenario included a set of micro-interactions, such as “the receptionist greets the patient” or “the patient asks the receptionist when their next scheduled appointment is,” that guided the participant in constructing an interaction episode. Participants took as much time as necessary, taking between 29 and 62 minutes ( $M = 45$  minutes, 30 seconds,  $SD = 10$  minutes, 37 seconds) to complete the design task. The setup of the study is shown in Figure 5.

### Recruitment & Selection

We recruited 10 participants—five interaction designers and five developers—from the local campus community. The interaction designers were selected from among students who



**Figure 5.** A researcher demonstrating the setup of the design sessions with the Interaction Blocks authoring environment and a NAO humanoid robot.

had completed the most advanced level interaction design course offered on campus, which taught principles and methods for user-centered interaction design and provided students with hands-on experience in design exploration and prototyping. Developers were recruited from among upper-level computer science majors with web development experience with the goal of capturing the perspectives of more technical users who might use our tool. Participants were between 19 and 26 years of age ( $M = 22.5$ ,  $SD = 2.68$ ) and were either computer science students or were employed in positions that required a computer-science background.

### Session Procedure

Following informed consent, participants were guided into a usability laboratory by a researcher. The researcher demonstrated Interaction Blocks, introducing the participant to the purpose of the tool, the layout of the interface, and the workflow necessary to explore and prototype interaction episodes and to test them with the robot. The participant was given five minutes to explore the interface independent of the researcher and was provided with the opportunity to ask questions at the end of the exploration period. At the end of this period, the researcher provided the participant with the scenario and left the room. The participant was given as much time as necessary to complete the design task. Following the completion of the task, the participant completed the System Usability Scale (SUS). The experimenter then conducted a semi-structured interview with the participant on the design of Interaction Blocks, their experience with using the authoring environment, and how it facilitated their exploration. Participants were compensated \$10 USD for their time.

### Analysis

Following guidelines suggested by Bangor et al. [2], a usability score from 0 to 100 was calculated by rescaling each of the 10 items used in the SUS to range from 0 to 4, summing the scores for the 10 items, and multiplying the result by 2.5.

We used content analysis to analyze the interview data [3]. Each interview was transcribed, and pertinent responses from each participant were included in a spreadsheet. After reviewing the responses across all participants, significant responses

were noted, and affinity diagramming was used to organize the responses into emerging themes, refining as necessary.

### FINDINGS

In this section, we present our findings from two different sources of data: SUS scores and interview data. The SUS scores serve as a metric of the overall usability of the authoring environment in enabling the use of a pattern language for design exploration and prototyping in creating human-robot interactions. Findings from the interview data provide insights into the designers' and developers' experiences in the design sessions and guidelines for future design and development of methods and tools for supporting human-robot interaction design.

### SUS Scores

The SUS scores for the authoring environment ranged from 75 to 95 ( $M = 84$ ,  $SD = 7.38$ ), interaction designers assigning an average score of 82, while the developers assigned an average score of 86. Based on established guidelines for interpreting SUS scores [2], a score of 80 or higher places the interface in the highest quartile when considering a survey of interfaces evaluated using the SUS.

### Interview Data

Below, we highlight themes found from the interview data. Excerpts from the interview are labeled with either ID or DE for the interaction designers and developers respectively, followed by a number.

#### *Patterns' Ease of Use*

Interaction designers and developers stated that the patterns—particularly patterns with multiple states—made it easy to quickly compose and test interactions. For patterns with multiple states, participants felt that the authoring environment's ability to automatically add the second state to the opposite role alleviated the burden of mentally planning the necessary pairings of behaviors and helped prevent creating inappropriate behaviors such as composing a question without an answer. Additionally, ID2 noted that the use of patterns—rather than a generic dialogue box used to accommodate all utterances—forced him to more carefully consider his dialogue prior to adding it. The same participant also inquired about the ability to create custom patterns. The excerpts below illustrate responses related to the ease of use of the pattern language.

*DE1: For most people, programming is not the easiest thing, and this is a very intuitive way to design dialogues.*

*ID2: I think the linkages of showing questions-answer, instruction-action show the clear delineation with that. That would really help in formalizing the structure and kind of removing the ambiguity of real conversation from what you're working on...*

However, while participants commented on the ease of use of the patterns in general, some expressed a desire for additional guidance on the nuances between some of the patterns due to their lack of domain knowledge. For example, DE2 noted

that there was a “monologue” pattern, a “comment” pattern, and a “monologue-comment” pattern. The participant was initially unsure why a “monologue-comment” pattern would be different from the combination of the individual “monologue” and “comment” patterns. DE1 voiced confusion about the presence of four “comment” states across three patterns. While participants generally felt comfortable using the patterns to accomplish the task goals, many participants asked for some form of documentation. Suggestions ranged from a traditional documentation interface with examples or a walk-through (ID4) to tooltips for each pattern (ID1).

#### *Design Exploration*

Participants found the ability to compose interactions on the timeline of the authoring environment and to execute them on the robot to support their design exploration and iterative design of the behaviors and interactions that they were building toward achieving natural dialogue. These comments are best illustrated by excerpts below by ID3.

*ID3: I think that human-human interaction can change all the time, and hearing it more and more and more would make it better in the end.*

*ID3: I mean it seems like this kind of interface is kind of, like, to play, so you can start using questions and answers right. Change the questions and the answers, add a comment in if you want.*

#### *Approachable Interface*

All participants commented on the clean and minimalistic design of the interface for the authoring environment, discussing it as an enabling tool that gave them the confidence that the complex task of synthesizing an interaction would be doable, as highlighted by the excerpts below.

*ID4: It's very clean in design. It seems easy to use for that reason. It doesn't have many functions, so it seems like it's easy to learn.*

*ID3: I think everything's very easy to see, like the flow. It's not easy to get lost. It's the least amount of details for someone who's maybe not majoring computer science.*

*DE1: For most people, programming is not the easiest thing, and this is a very intuitive way to design dialogues.*

#### *Error Prevention*

DE1 and DE3 noted that the use of the pattern language and the visual authoring environment would aid designers who may have little to no development or programming experience in building an interactive application with a robot, as illustrated in the excerpt below. Their comments also suggested that this approach would enable developers to bypass many of the bugs frequently encountered when programming, such as syntax errors and typos, and alleviate the need for them to learn a new application programming interface (API).

*DE3: If it has the functionality to prevent me from making mistakes, then that's good too, right?*

#### *The Need to Support Free-Form Exploration*

Some designers and developers expressed that the addition of a less structured environment might better facilitate their design exploration. Examples included a whiteboard (DE1, DE3), with paper and pencil (ID1), sticky notes (ID4), and a text editor (ID1, ID2). While these participants appreciated the structure and code generation benefits provided by Interaction Blocks, they felt a need for a less-structured initial step to the exploration that allowed them to easily create, edit, and view the behavior and dialogue components that they planned to use in the more structured construction of the robot's exchange with users. Comments from DE2 and ID2 below illustrate the need to support free-form design exploration.

*DE2: Well honestly I'd probably write out a script in a text editor, and then drag out all the parts and fill in the text.*

*ID2: This more feels like that I know what my design is and I am going to formalize it somehow rather than 'this is how I lay it out and tweak it as I go kind of thing.' I almost feel this is the last stage.*

#### *Need for Branching*

While participants appreciated the simplicity of the current interface, several participants raised the question of accommodating alternative ways in which the interaction or the dialogue between the robot and its user might unfold. Questions focused on two facets of this problem: accommodating different utterances with the same semantic meaning (e.g., “Yes” versus “Yes, please”), and accommodating utterances with different semantic meanings (e.g., “Yes” versus “No”). Participants suggested adding into the authoring environment support for a list of possible responses for accommodating utterances with the same semantic meaning. To handle multiple branches of dialogue that result from utterances with different semantic meanings, participants suggested creating multiple timelines that can be collapsed and expanded (DE5) and displaying multiple timelines at once (ID4).

In addition to these suggestions, participants commented on the possibility of additional visualizations, color coding patterns for easier identification and disambiguation, and adding additional information concerning the robot's status.

## **DISCUSSION**

Results from the design sessions focused on three main findings: the use of design patterns, the design of the authoring environment, and the workflow required to synthesize dialogue. These results offer many future directions for improving and expanding the pattern language and authoring environment to better support human-robot interaction design.

### **Design Patterns**

The use of a pattern language was reviewed favorably by interaction designers and developers alike, with many noting the ease in which they could explore and prototype exchanges for the robot and users through their use. Additionally, participants highlighted the structure that patterns afforded as an advantage, noting that patterns forced them to consider their design choices for dialogue and interaction elements and how they contributed to the overall goals of the interaction.

Although participants found patterns easy to use, what interaction elements some patterns represented was not always clear to them. Participants suggested various forms of providing help within the authoring environment, including tooltips and examples to help acclimate users to the pattern language. Documentation of the patterns would not only help interaction designers better understand the design elements, but would also inform researchers interested in using and extending the pattern library presented here.

### **Authoring Environment**

The authoring environment was cited by all participants for its ease of use and approachable design, which was confirmed by the SUS scores. The requests and suggestions for improvement discussed in the previous section underline the importance of accommodating more advanced functionality, such as branching, in future versions of the authoring environment without sacrificing too much from the simplicity of the current design. Some participants suggested the use of a “superuser” mode or the ability for users to reveal or hide detailed options or functionality for patterns.

### **Workflow**

Participants frequently cited the ability to rapidly modify and evaluate their design ideas as a useful feature that aided them in the iterative development of their final design. This finding confirms prior work that highlights the importance of rapid design exploration and prototyping in the development of human-robot interactions to provide designers with a better understanding of how their designs would perform in real-world settings [18]. Additionally, the use of Interaction Blocks frees the designers’ workflow of debugging, enabling them to concentrate on designing interactions.

While participants expressed enthusiasm for the use of Interaction Blocks as a part of their workflow for designing and prototyping human-robot interactions, they also expressed a desire for a less formal design step prior to using Interaction Blocks. Some participants viewed using the authoring environment as a last step to formalize their work, while others suggested using Interaction Blocks earlier in the design process after brainstorming to create a basic idea of how the designed interaction may unfold. Further studies of the design practices of interaction designers might inform the design of flexible authoring environments that support informal as well as formal exploration at different stages of the design process.

Finally, designers and developers voiced a need for the ability to sketch out dialogue exchanges, see how they fit within the interaction, and easily refine individual dialogue elements. To support this need, the authoring environment might provide a view of the script of the entire dialogue, enabling users to easily gain an overview of their design and make changes at each exchange from a global perspective. Additionally, this sketching environment might allow users to draft the dialogue first and use a drag-and-drop interface to add these dialogue elements to the patterned interaction elements.

### **Interaction Designers vs. Developers**

While there were many topics that interaction designers and developers agreed on, there were also differences in how

these two groups evaluated the use of the pattern language and authoring interface. For instance, interaction designers were more concerned with how new visualizations might enable users to better compose and manage complex interactions, while developers focused on the reduction in errors enabled by the use of the pattern language and visual authoring, as well as the desire for more information on the robot’s status, the ability to batch edit and move large portions of the interaction, and the need for incorporating branching into the development of the dialogue. These trends might reflect how interaction designers with different training, expertise, or backgrounds might have different needs for authoring human-robot interactions; those with a development background might be empowered by the ability to more precisely control the robot, while those with a design focus might want more flexibility in design exploration.

### **Limitations**

The pattern language that we used in building our authoring environment relied on our observations and analysis of interactions in five scenarios. While we carefully chose these scenarios to represent many of the interactions robots will encounter, additional or more complex scenarios might reveal additional patterns. Furthermore, the evaluation of the use of the pattern language and authoring environment involved primarily student designers and developers in relatively short design sessions, due to the limited volume of interaction design practice for robotic technologies and limited number of human-robot interaction designers. Future explorations might seek to engage professional interaction designers with experience in human-robot interaction design in longer-term design sessions to better understand how the approach presented here might support design exploration and prototyping human-robot interactions. Given the difficulty and overhead involved in programming complex robot systems, design tools such as Interaction Blocks might significantly benefit such users.

### **Lessons Learned**

The development of the pattern language and the authoring environment required decisions that might prove valuable for future work in this or related areas. One key challenge was selecting a diverse set of social interaction scenarios for our formative study of human interactions. While there are other roles that robots will likely fulfill, such as coaching, we needed to balance the diversity of scenarios with the workload involved in analyzing a large corpus. Another challenge was discovering an appropriate level of abstraction for our models (Figure 3) that would enable their use as design patterns. We followed an iterative process of reviewing the data from human interactions and sketching, constructing, and refining our models until no further modifications could be made.

### **CONCLUSION**

In this work, we explored how a design pattern language and visual authoring environment might enable designers to rapidly perform design exploration and prototyping for human-robot interaction. We observed and analyzed interactions from eight dyads engaged in five scenarios and developed seven interaction design patterns. We then built an authoring environment, *Interaction Blocks*, to enable interaction

designers to use these patterns to rapidly construct, evaluate, and refine human-robot interaction. We conducted a qualitative evaluation of the use of the pattern language and authoring environment with a group of ten interaction designers and developers, as they prototyped an exchange between a robot and its user. Our results highlight the potential for the use of design patterns and the workflow that our authoring environment promotes to design, prototype, and evaluate interactions, enabling interaction designers to take advantage of patterns to synthesize complex interactions.

#### ACKNOWLEDGMENTS

This work was supported by National Science Foundation award 1149970 and National Institutes of Health award 1R01HD071089-01A1. We would like to thank Brandi Hefty, Jilana Boston, Ross Luo, Jason Sauppé, Lindsay Jacobs, and Catherine Steffel for their help with our work.

#### REFERENCES

1. Andrist, S., Mutlu, B., and Gleicher, M. Conversational gaze aversion for virtual agents. In *Proc. IVA* (2013), 249–262.
2. Bangor, A., Kortum, P. T., and Miller, J. T. An empirical evaluation of the system usability scale. *Intl. Journal of Human-Computer Interaction* 24, 6 (2008), 574–594.
3. Berelson, B. Content analysis in communication research.
4. Brandstätter, V., Lengfelder, A., and Gollwitzer, P. Implementation intentions and efficient action initiation. *Journal of personality and social psychology* 81, 5 (2001), 946–960.
5. Clark, H. *Using language*. Cambridge U, 1996.
6. Clark, H. H. *Arenas of language use*. U Chicago, 1992.
7. Datta, C., Jayawardena, C., Kuo, I. H., and MacDonald, B. A. Robostudio: A visual programming environment for rapid authoring and customization of complex services on a personal service robot. In *Proc. IROS* (2012), 2352–2357.
8. Fong, T., Thorpe, C., and Baur, C. Collaboration, dialogue, human-robot interaction. 255–266.
9. Fussell, S. R., Kiesler, S., Setlock, L. D., and Yew, V. How people anthropomorphize robots. In *Proc. HRI* (2008), 145–152.
10. Glas, D. F., Satake, S., Kanda, T., and Hagita, N. An interaction design framework for social robots. In *Proc. RSS* (2012).
11. Goffman, E. The interaction order. *American Sociological Review* 48, 1 (1983), 1–17.
12. Hsiao, K.-Y., Mavridis, N., and Roy, D. Coupling perception and simulation: steps towards conversational robotics. In *Proc. IROS* (2003), 928–933.
13. Huang, C.-M., and Mutlu, B. Robot behavior toolkit: generating effective social behaviors for robots. In *Proc. HRI* (2012), 25–32.
14. Humphrey, C. M., Henk, C., Sewell, G., Williams, B. W., and Adams, J. A. Assessing the scalability of a multiple robot interface. In *Proc. HRI* (2007), 239–246.
15. Kahn, P. H., Freier, N. G., Kanda, T., Ishiguro, H., Ruckert, J. H., Severson, R. L., and Kane, S. K. Design patterns for sociality in human-robot interaction. In *Proc. HRI* (2008), 97–104.
16. Kooijmans, T., Kanda, T., Bartneck, C., Ishiguro, H., and Hagita, N. Interaction debugging: an integral approach to analyze human-robot interaction. In *Proc. HRI* (2006), 64–71.
17. Langellier, K. M. Personal narratives: Perspectives on theory and research. *Text and Performance Quarterly* 9, 4 (1989), 243–276.
18. Lohse, M., Siepmann, F., and Wachsmuth, S. A modeling framework for user-driven iterative design of autonomous systems. *International Journal of Social Robotics* (2013), 1–19.
19. Matarazzo, J., and Wiens, A. *Interview: Research on Its Anatomy and Structure*. Aldine De Gruyter, 2007.
20. Mutlu, B., Forlizzi, J., and Hodgins, J. A storytelling robot: Modeling and evaluation of human-like gaze behavior. In *Proc. HUMANOIDS* (2006), 518–523.
21. O’Leary, M. J., and Gallois, C. The last ten turns: Behavior and sequencing in friends’ and strangers’ conversational findings. *Journal of Nonverbal Behavior* 9 (1985), 8–27.
22. Peltason, J., and Wrede, B. Modeling human-robot interaction based on generic interaction patterns. *AAAI Symposium on Dialog with Robots* (2010).
23. Perlin, K. Improving noise. *ACM Transactions on Graphics (TOG)* 21, 3 (2002), 681–682.
24. Quigley, M., Conley, K., Gerkey, B., Faust, J., Foote, T., Leibs, J., Wheeler, R., and Ng, A. Ros: an open-source robot operating system. In *ICRA Workshop on Open Source Software* (2009).
25. Sacks, H., Schegloff, E. A., and Jefferson, G. A simplest systematics for the organization of turn-taking for conversation. *Language* 50, 4 (1974), 696–735.
26. Schegloff, E. *The first five seconds: The order of conversational opening*. PhD thesis, University of California, 1967.
27. Schegloff, E. A. Sequencing in conversational openings. *American Anthropologist* 70, 6 (1968), 1075–1095.
28. Skehan, P. A framework for the implementation of task-based instruction. *Applied Linguistics* 17, 1 (1996), 38–62.
29. Tannen, D. *Talking voices: Repetition, dialogue, and imagery in conversational discourse*, vol. 6. Cambridge U, 1989.