

Design Your Robot to Listen Attentively: A Novel Backchanneling Authoring Interface for Non-Experts

by

Samuel Richard Martel Barrett

A thesis submitted to
The Faculty of Graduate and Postdoctoral Studies of
The University of Manitoba
in partial fulfillment of the requirements
of the degree of

Master of Science

Department of Computer Science
The University of Manitoba
Winnipeg, Manitoba, Canada

© Copyright 2026 by Samuel R. M. Barrett

Design Your Robot to Listen Attentively:

A Novel Backchanneling Authoring Interface for Non-Experts

Abstract

We present a novel interface for non-experts (e.g., learners, educators, out-of-field researchers) to author custom human-robot dialogs that include backchanneling—utterances or gestures made while listening. We engaged 36 participants in three sets of workshops, starting from elicitation workshops to investigate how participants perceive, describe, and structure their understanding of robot backchanneling (i.e., their mental models), to inform the design of novel interface prototypes. We conducted paper prototype sessions with participants completing representative tasks to iterate on our design. Finally, we built a high-fidelity prototype (using Blockly and a small robot) and conducted evaluation workshops to assess if the design was successful in being approachable by non-experts: it was a resounding success as all participants were able to quickly and easily engage with our interface and dialog authoring tasks to create a rich variety of nuanced, custom behaviors, demonstrating that the mental models guiding our prototype effectively supported user understanding and rich content creation.

Acknowledgments

First, I extend my sincere appreciation and gratitude to Dr. Jim Young. Your tireless mentorship and paper edits encouraged me to believe in myself and to elevate myself to a next level I didn't know I was capable of. When I thought everything would go wrong, I appreciated your reassurances and obscure Monty Python references. You were right, everything is fine.

I am thankful to my committee, Dr. Andrea Bunt and Dr. Hamid Mansoor, for their feedback; you especially taught me the value of understanding perspective, framing, and a good sales pitch.

To James Berzuk, I am grateful for your steadfastness, for showing me the ropes, helping me execute my studies, and most of all for your camaraderie, including while accompanying me across the world twice: to my first conference at HRI'25, and to Japan. Thank you.

Danika, James, Raquel, and Zachary, I appreciate our group for our many jokes, banter, shared commiseration and support for each other. Every day, I look forward to lunch time, or any time when we no longer feel like working and pull up chairs to chat.

The entire HCI Lab is an amazing group of kind and passionate people. To each of you, I hope I helped make your graduate school experience as fun and rewarding as you have made mine.

To Aidan, Matthew, and Owen, when I want to laugh, feel better, or simply be annoying, I know exactly who to go to. You are my support pillars and my inspiration—I'm told I set a high bar, well, you're the ones pushing it up. You each mean the world to me.

To Maman and Dad, you instilled in me the curiosity, perseverance, and belief in myself that brought me here. Even more, you taught me compassion, kindness, and respect, and to always strive to be a good person. Thank you for cheering me on.

Finally, Sam, your endless belief in me pushed me to grow and succeed. From celebrating my euphoric successes to supporting my stressful lows, your reasoning voice and listening ear are steady and invaluable. I would not have done this without you.

Table of Contents

Abstract	i
Acknowledgments	ii
List of Figures	vii
List of Tables	x
Publications	xi
1 Introduction	1
1.1 Backchanneling	3
1.2 Research Questions	4
1.3 Methodology	5
1.4 Contributions	7
1.5 Overview	8
2 Related Work and Background	10
2.1 User-Centered Design	10
2.1.1 <i>User Models</i>	11
2.1.2 <i>Participatory Design</i>	12
2.2 End-user Creation in Human-Computer Interaction	14
2.2.1 <i>GlucoMaker: Designing Your Own Glucose Monitor</i>	16
2.3 End-user Authoring for Human-Robot Interaction	16
2.3.1 <i>Block Programming Interfaces for Robots</i>	18
2.3.2 <i>Authoring Robot Dialog</i>	20
2.3.3 <i>Granularity of Control</i>	21
2.4 Backchanneling	23

2.4.1	<i>Backchanneling Robots</i>	24
2.5	Chapter Summary	25
3	Elicitation Study: Understanding Non-Expert Mental Models	26
3.1	Tasks and Procedures	27
3.1.1	<i>Analysis Plan</i>	30
3.2	Elicitation Study Results: Non-expert Mental Models of Robot Backchanneling.....	31
3.2.1	<i>Knowledge of Backchanneling</i>	31
3.2.2	<i>Non-experts' Considerations when Authoring Backchanneling</i>	34
3.2.3	<i>Fuzzy Boundaries: Where Non-experts Transcend Backchanneling</i>	36
3.3	Elicitation Study Discussion.....	39
3.3.1	<i>Reflection on Methodology</i>	40
3.4	Design Recommendations for Dialog Authoring Tools.....	41
3.5	Chapter Summary	42
4	Formative Prototyping Sessions: Paper Prototypes.....	44
4.1	Designing an Initial Prototype Interface	44
4.1.1	<i>Informal Interface Brainstorm Workshop</i>	45
4.1.2	<i>Interface Design</i>	47
4.1.3	<i>Pre-study Interface Iteration</i>	49
4.2	Formative Prototype Evaluations with Non-experts.....	51
4.2.1	<i>Tasks and Procedure</i>	52
4.3	Results: Prototype Iteration.....	54
4.3.1	<i>First Workshop: Style Discoverability and Customization</i>	54
4.3.2	<i>Second Workshop: More Style Management Issues</i>	57
4.3.3	<i>Final Workshop: Simplification</i>	58
4.4	Chapter Summary	59
5	High-Fidelity Content Creation: A Tool for Designing Robot Backchanneling	60

5.1	Interface Design.....	60
5.2	Prototype Evaluation Workshop	64
5.2.1	<i>Tasks and Procedure</i>	64
5.2.2	<i>Analysis Plan</i>	65
5.3	Results.....	66
5.3.1	<i>Low Friction and Sustained Engagement</i>	66
5.3.2	<i>From High-Level Structure to Tuning</i>	68
5.3.3	<i>Exploration of Rich and Nuanced Behaviors</i>	69
5.4	Discussion.....	72
5.4.1	<i>Limitations</i>	74
5.5	Chapter Summary	76
6	Conclusion	77
6.1	Contributions	77
6.2	Recommendations	80
6.3	Limitations.....	81
6.4	Conclusion.....	82
	Appendices.....	83
	Appendix A: Elicitation Workshop Study	83
A.1	<i>Recruitment Poster</i>	83
A.2	<i>Recruitment Email Template</i>	84
A.3	<i>Informed Consent Form</i>	85
A.4	<i>Demographic Questionnaire</i>	90
A.5	<i>Elicitation Workshop Protocol Script</i>	91
A.6	<i>Authoring Worksheets</i>	103
	Appendix B: Formative Paper Prototype Evaluation	107
B.1	<i>Informed Consent Form</i>	107

<i>B.2</i>	<i>Formative Paper Prototyping Protocol Script</i>	<i>112</i>
<i>B.3</i>	<i>Task Instruction Script</i>	<i>118</i>
Appendix C: Evaluation Workshop.....		120
<i>C.1</i>	<i>Informed Consent Form.....</i>	<i>120</i>
<i>C.2</i>	<i>Evaluation Workshop Protocol Script.....</i>	<i>125</i>
References		130

List of Figures

Figure 1: Our multi-phase user-centered design method employing three sets of workshops in order: <i>elicitation</i> —to understand mental models and inform design, <i>iteration</i> —to prototype and iterate to improve our design, and <i>prototype evaluation</i> —to assess whether the interaction and prototype design aligned with participants’ expectations and reasoning, enabling free and rich authoring of robot backchanneling behaviors.	5
Figure 2: Elicitation workshop space including paper design activities (mock session shown)....	8
Figure 3: Our design workshop space utilized by participants to create backchanneling robots, including our interface and a VStone Sota robot.	9
Figure 4: GlucoMaker enables users to design their own custom glucose monitoring device using a digital interface, which is then brought to life and worn by the user (taken from [53]).	17
Figure 5: The PODER framework leverages block programming in the conceptual and implementation stages (tangibly and digitally) to support non-expert end users in ideating and creating robot behaviors (from [20]).	19
Figure 6: Zhang et al.'s end-user programming interfaces featured two levels of granularity for behavior authoring: a high-granularity block programming interface (not shown), and a low-granularity interface featuring simple slider controls for basic parameters, and a set of available "personas" to choose from (from [106]).	22

Figure 7: Design workshop space showing a mock session in progress (we did not have permission for photos of participants).	27
Figure 8: Worksheets for authoring dialog included pre-filled samples and blank copies with black text and bubbles only.	30
Figure 9: Prototype robot dialogs and characters, annotated with backchanneling and other dialog behaviors (P25, P6, P1, P7).	34
Figure 10: We tasked participants with authoring robot dialogs with backchanneling on paper to elicit their mental models for designing robot backchanneling behavior. From left to right: P24, P11, P1.	35
Figure 11: Fine-tuning of individual backchanneling behaviors was a theme among sketches from our informal interface brainstorm workshop (R2, R4, R1).	46
Figure 12: Interface sketches for authoring backchanneling emphasized the use of personality descriptors, such as "happy", "extroverted" and "affirming" (R2, R3, R4).	46
Figure 13: Paper prototype of our initial interface design, complete with core block features (toolbox, left, and workspace, middle) and all additional elements including tooltips, keyboard and selection hints, dropdown menus, navigation and callouts.	48
Figure 14: Improved interface design featuring styles as draggable modifier blocks that are attached to core dialog elements (listening, speaking).	51

Figure 15: Users can now create a new style from a blank state, or use an existing style as a starting point, as desired. Options in the blank state appear unchecked, and reveal default values when enabled (e.g., middle slider values).	55
Figure 16: We clarified how attempting to edit a preset style (top) preserves the style, and users are invited to create a new style with the preset as a starting point (bottom).	56
Figure 17: Our final paper interface prototype, deployed in the final workshop.	57
Figure 18: Our interface, based on Blockly, with available blocks in the tool palette (left) and user creations in the workspace (right). The "Play" button on the upper right lets the user test their designs on a connected, physical robot.	62
Figure 19: A diagram showing the blocks available to the user for authoring human-robot dialog with backchanneling.	62
Figure 21: Users can optionally edit backchanneling styles to fine-tune parameters for individual behaviors. Changes are saved only to the modified block; other blocks of the same style remain unaffected.	63
Figure 22: Users authored a range of interactions for their robots and applied backchanneling to them. Pictured is a snapshot of a portion of P10's workspace and their joke-telling robot.	67
Figure 23: P5 used backchanneling styles contextually while authoring dialog, including an <i>excited</i> style that repeatedly says "wow".	68

List of Tables

Table 1: Participation of individuals across all three studies, broken down by each study and each session (e.g., Elicitation study, S4).....	32
Table 2: Our categorization of participant-created backchanneling behaviors from the card sort.	37

Publications

Some ideas in this thesis have appeared previously in the following publication:

Samuel R. M. Barrett and James E. Young, 2025. Toward Designing a Toolkit for Intuitive Design of Backchanneling Behaviour in Social Robots. In *Proceedings of the 2025 ACM/IEEE International Conference on Human-Robot Interaction (HRI '25)*, March 04, 2025. IEEE Press, Melbourne, Australia, 1236–1239.

1 Introduction

Human-computer interaction (HCI) has a long and rich history of developing end-user creation tools to enable end users to easily author content and be creative with computing technologies. These tools can empower new demographics of users, to participate in design and creation with computing systems, such as those with minimal experience who may have otherwise been excluded due to technical complexities. In this way, end-user creation tools are critical in empowering end users of technologies to exercise agency and control in how technologies are deployed and used, including customizing systems to function according to their needs and preferences. In this work, we expand on this thread of end-user creation tools by exploring a novel factor in end-user design of human-robot interactions (backchanneling during dialog) and making the authoring of these behaviors approachable to users without robotics experience.

End-user creation tools often aim to simplify authoring of systems and behaviors by providing simplified abstractions that are approachable for end users to design and manipulate. For instance, the Teddy interface allowed users to design 3-dimensional models of stuffed animals by drawing a 2-dimensional representation of the desired output, which the system converted into 3D models [49]. Such creation interfaces allow the user a simplified means of interacting with complex tasks and concepts (such as authoring digital 3D models), enable creative expression in authored outputs, and can even support rapid prototyping and iteration [38]. Further, end-user creation tools can support users in personalizing custom systems that meet their needs and preferences, for instance, the GlucoMaker toolkit for users to design their own personalized glucose monitoring devices, which are worn on the body [53]. Evidently, these creation interfaces need not be constrained to digital inputs and outputs; examples include

toolkits for authoring physical hardware interfaces [39] and for creating personalized glucose monitoring devices [53]. In this work, we investigate digital interfaces for end-user design of rich robot dialog behaviors with a focus on backchanneling.

Software tools that enable people to author and customize personalized robot behaviors without programming (e.g., drag and drop interfaces) have become prevalent in research [20, 23, 26, 42, 99] and commercial packages.¹ These easy-to-learn and use interfaces conceal technical complexities and enable non-experts, including out-of-field researchers, hobbyists, and STEM students to engage in the creative process of robot behavior design [20, 39, 88, 102]. Even for in-field experts (i.e., roboticists), such tools can enable rapid prototyping, supporting exploration and creativity [38]. We present a novel interface design and prototype that enables authors to create and enrich human-robot dialog by customizing the robot’s use of brief reactions (e.g., saying “uh-huh”, nodding, etc., i.e., backchanneling [76]) to convey attentiveness and encourage interaction.

Common tools for authoring rudimentary human-robot dialog focus almost exclusively on text, requiring authors to provide written text to define both robot utterances and keywords or phrases to listen for, as with Misty Studio (Misty robot), VStone Magic (Sota robot), Choreographe (Softbank Nao), EV3 Classroom (LEGO Mindstorms), Sphero Edu (BOLT), among many others, and the focus on text is a theme with robot dialog systems more broadly [10]. However, natural dialog between people is much richer than what is captured in text alone, and is often enriched with other behaviors such as backchanneling [76]. Despite this, many commercial robots’ software packages do not engage with backchanneling, and no tool to date

¹ Such commercial products include: VStone Sota (2015); Misty Robotics Misty II (2019); SoftBank Nao v5 (2014).

enables non-expert creation or modification of backchanneling in dialogs they create. In this work, we present the first end-user tool for designing and customizing the backchanneling behavior of a social robot.

1.1 Backchanneling

Backchanneling is the use of a supplementary channel of communication during dialog, where an interlocutor communicates with another using gestures or short utterances that aim to avoid interrupting speech [76]. For example, listeners indicate attentiveness and encourage interaction using utterances (e.g., “hmm”, “sure”) and gestures (e.g., nodding, facial expressions). Backchanneling serves many important functions in dialog, including to convey attention, feedback, interest or boredom, coordinate turn-taking and rhythm, and build rapport [58, 85, 90, 92]. Further, backchanneling behavior varies greatly by individual and context [15, 27, 103], highlighting the need for flexibility and variability in how robots incorporate backchanneling into their dialog behaviors, and potential for highly specific customizability and personalization of robot backchanneling.

Given the importance of backchanneling for dialog, researchers have investigated the effects of robot backchanneling on human-robot interactions. User perceptions of a robot, such as its perceived personality [14], change based on differences in backchanneling behavior [28, 54]: appropriate use of backchanneling can keep children more engaged in storytelling [76], increase participation of out-group team members [95], encourage students and improve learning performance [78], or encourage sustained interaction in complex tasks such as music co-creation [67]. The range of these interaction effects underline the need for robots to have a range of backchanneling styles and techniques tailored to intended use cases, created by intentioned

designers, rather than simply relying on automated generation methods (e.g., [3, 54]) that may not be aware of the context, personality, or interaction design intent. Our novel tool for authoring custom robot backchanneling expands on these common behavior authoring tool designs to enable fine-tuned control of robot backchanneling behavior for dialog creators.

1.2 Research Questions

To guide our aim to support non-expert end users to author custom robot backchanneling behaviors, we addressed the following research questions:

1. How readily and comfortably do non-experts engage with backchanneling conceptually?
 - a. Which elements of backchanneling can non-expert potential users readily articulate and demonstrate a clear understanding of?
 - b. Which elements of backchanneling present some difficulty for non-expert users to engage with?
2. How might non-experts describe backchanneling elements of robot dialog?
 - a. What vocabulary, metaphors, or comparisons do they use?
 - b. How do they tend to specify elements of backchanneling when tasked with describing a hypothetical instance of human-robot dialog?
 - c. How granular do they tend to describe robot backchanneling? E.g., on a per-strategy basis or an overall robot personality?
3. How do we incorporate non-expert user mental models of robot backchanneling (as elicited in RQ1,2) into the design of a novel interface for authoring robot backchanneling behavior?
 - a. What layout and interaction design can satisfy our requirement of enabling non-experts to author their desired behaviors?

- b. How do users engage with a prototype interface and its features to author backchanneling behaviors?
- c. What challenges arise while using the interface to author behaviors?

1.3 Methodology

To address our research questions described above, we designed and employed a three-phased participatory approach (Figure 2) to understand how non-experts may engage with robot backchanneling, how we can present backchanneling creation controls in an intuitive manner, and how our interface design supports non-experts to author robot backchanneling. To start designing such a tool, in our first phase we sought to understand potential users’ perspectives and ideas about backchanneling robots, what they may want from a backchanneling robot, and how they describe robot backchanneling behaviors (RQ1,2). Following, in phase two, we needed to ensure that our interface design, based on our findings from the first phase, was effective in enabling users to author robot backchanneling (RQ2,3). In the third and final phase, we sought to evaluate how our resulting interface design supported user engagement and authoring of custom robot backchanneling behavior, the range of creative designs made possible, and the limitations of our interface design and approach (RQ3). In the remainder of this section, we expand on our methodology for each of the three phases.




Figure 1: Our multi-phase user-centered design method employing three sets of workshops in order: *elicitation*—to understand mental models and inform design, *iteration*—to prototype and iterate to improve our design, and *prototype evaluation*—to assess whether the interaction and prototype design aligned with participants’ expectations and reasoning, enabling free and rich authoring of robot backchanneling behaviors.

In our first phase, we addressed RQ1 and RQ2 by asking potential users about backchanneling and robots that backchannel, to build an understanding of their mental models. To do this, we conducted design workshops with groups of non-expert participants who were interested in designing social robot behaviors to understand how people conceptualize backchanneling and how they may approach designing backchanneling for robots. We asked participants to complete individual and group activities related to backchanneling, including watching video demonstrations to spark ideas, brainstorming potential robot behaviors, sharing their perceptions and ideas about backchanneling robots, and designing robot dialog with backchanneling on paper. Our findings from this phase informed our prototype design, iteration, and evaluation in subsequent phases.

For our second phase, we aimed to address RQ2c and RQ3a by creating a prototype interface for authoring robot backchanneling, guided by our findings from the first phase, and performing iterative user testing to validate and refine our interface design. First, informed by our findings from the design workshops, we designed initial interfaces for authoring robot backchanneling in dialog. We then conducted iterative formative evaluations of paper prototypes of these interfaces in a small user study with participants. As our goal was to refine our interface design and identify and correct potential design issues leading to user challenges, we updated our interface between each session of the study. Once satisfied that our interface prototype represented a reasonable, effective approach to enabling users to design backchanneling behaviors, we developed a software version of our final interface design before our evaluation in the third and final phase.

Our final phase addressed RQ3—we investigated how our prototype interface design supports non-expert end users to author robot backchanneling, the range of their creations, and uncover limitations with our interface and our design approach from the first two phases. To achieve this, we conducted a prototype evaluation workshop where we invited participants from previous phases to interact with our interface on a computer to program a robot’s dialog behavior, with an emphasis on backchanneling. This participatory approach allowed us to reflect on how well our resulting interface design mapped to the mental models we investigated in the first phase, and how this alignment supported non-experts in creative authoring of robot behaviors. We asked users to complete structured (guided) and unstructured (open-ended creative) tasks to explore the interface and create novel robot backchanneling behavior designs.

Through three workshop phases (Figure 2) with a total of 36 participants (see Table 1 for a breakdown of participant involvement across each study), our work provides insight into how people conceptualize robot backchanneling, illustrates how these mental models can inform interface design, and presents a novel interface design and prototype that enables non-experts to author rich and nuanced backchanneling behaviors. This work demonstrates the feasibility of empowering non-expert users to author subtle elements of robot dialog and motivates future directions for research on the design of accessible robot dialog authoring tools. In the following section (1.4), we briefly highlight our results, with more detail in subsequent chapters pertaining to each of the three phases (chapter 3, 4, and 5).

1.4 Contributions

The contributions of this work are summarized as follows.

1. We demonstrate how our approach to interface design—informed by our investigation of non-expert mental models—enabled non-expert users to design custom robot backchanneling behaviors for dialog.
2. We demonstrate how non-experts, potential users of robots, are readily able and willing to engage robot backchanneling, discuss its implications for human-robot interaction, and design their own customized robot backchanneling behaviors.
3. Our work informs future robot behavior authoring tools to support user agency and control through the creation and customization of rich and nuanced dialog behaviors.

1.5 Overview

In the remainder of this thesis, we discuss related work on social robots, end-user behavior authoring tools, and backchanneling (Chapter 2). We then detail each of the three phases



Figure 2: Elicitation workshop space including paper design activities (mock session shown).

of our work including methodology and results, starting with the elicitation study (Figure 2) which resulted in understanding how readily non-experts engage with and understand concepts surrounding robot backchanneling, and what conceptualizations they used to represent this (Chapter 3). We present our formative paper prototype sessions where participants completed representative tasks to inform our design iteration, resulting in a final interface layout and conceptual design for including backchanneling in a robot dialog authoring interface (Chapter 4). Finally, we implemented a high-fidelity prototype and conducted interactive prototype



Figure 3: Our design workshop space utilized by participants to create backchanneling robots, including our interface and a VStone Sota robot.

evaluation workshops (Figure 3) to examine whether participants understood and could meaningfully engage with the backchanneling system (Chapter 5); the results emphasize the success of our interface design in enabling non-expert users to author and modify a broad range of robot backchanneling behaviors. Finally, we conclude by discussing how this work advances end-user customization and control over social robot behaviors and its implications on future human-robot interaction research and tools, our recommendations to researchers and designers of social robot authoring tools, as well as limitations of our approach and opportunities for future work.

2 Related Work and Background

In this chapter, we describe principles, concepts, and threads of work that informed our approach to designing an end-user authoring tool for robot backchanneling. These related works guide our research questions as follows; first, we note the importance of understanding user mental models for effective interface design and ask how users may engage backchanneling conceptually (RQ1) and how they may describe its features and elements (RQ2). Second, given the prevalence of end-user authoring tools and the importance of backchanneling for dialog, how might end users approach authoring custom backchanneling for a robot? (RQ3).

This chapter proceeds by first outlining the fundamental human–computer interaction (HCI) concepts we followed in our work. Next, we investigate the landscape of end-user creation in HCI, including techniques for involving and enabling end users, and examples of specific tools. We then narrow our focus to end-user authoring of robot behaviors, and further on authoring robot dialog. We conclude this chapter with an overview of backchanneling, including knowledge from human–human interaction, how backchanneling is currently employed in robots and knowledge within the human–robot interaction (HRI) field on how backchanneling impacts HRI.

2.1 User–Centered Design

User-centered design is one of the key tenets of human–computer interaction which considers the user paramount in the design of systems, and is often employed toward designing usable, effective systems that consider user needs and potential challenges [93]. Done properly, user-centered design can alleviate barriers to interaction; for example, gulfs of execution, where the

user struggles to engage the system, or gulfs of evaluation where the user struggles to interpret the resulting action taken by a system [74]. This can be especially critical in systems where the user describes or manipulates complex technical operations, for instance, robot dialog with backchanneling, thus user-centered design underpins our work, informing our methodology to understand user perspectives and involve users throughout the design process.

2.1.1 User Models

Through observation and interaction with a system, the user constructs an internal representation of the system’s function, behavior, or state—a user model—as a framework to understand the system and interact with it successfully [30, 51]. Often called mental models, these models are an important aspect of user-centered design, and understanding the user’s model can facilitate the design of effective, understandable systems that promote successful interactions [29]. As such, human-computer interaction researchers seek to understand how a user’s mental model is formed, how it may evolve as the user interacts with a system, and how a system can leverage an understanding of the mental model in its design to facilitate its use. In this work, we uncover user mental models of robot backchanneling—specifically, the perspectives of non-expert users interested in designing robot behaviors—which informed our interface design for a prototype robot backchanneling authoring tool.

A key challenge of designing for non-experts is the question of how to represent the complex technical and behavioral concepts in ways that are meaningful to people and easy for them to understand and use, despite not having specialized robotics or linguistics background. For example, interfaces that align with people’s existing mental models are more “intuitive” and easy to use [32, 38, 39, 75], and thus human-computer interaction commonly investigates end-user

mental models to understand potential user challenges [7, 21, 68] and to inform interface design [16, 30, 51, 75]. This includes leveraging end-user mental models for authoring interfaces, such as for computer programming, e.g., for school teachers [2, 41, 42]. Similarly, human-robot interaction research continues to explore using end-user mental models to promote successful human-robot interactions [57, 79, 82, 83, 96], and to inform robot behavior authoring tools [1, 48, 65]. Developing an understanding of non-expert mental models of robot backchanneling will be critical to designing these authoring interfaces; unfortunately, there is currently little or no knowledge of non-expert mental models of robot backchanneling. Thus, our work is grounded in eliciting non-expert authors’ mental models—how they talk about and conceptualize robot backchanneling—and using this understanding to inform and guide our interface design and evaluation.

2.1.2 Participatory Design

Participatory design—a design philosophy where a system’s eventual users are actively involved in its design—is a popular approach of user-centered design in HCI, encompassing a broad range of approaches and methodologies [33, 89]. Participatory design as a methodological philosophy seeks to involve users or affected parties throughout the design process, often from the earliest possible stages, with the goal to create systems that account for user needs and preferences by explicitly and directly involving the user. In this way, participatory design combines the systems and interaction expertise of the designer with the lived experience expertise of the target user, with the goal that “more relevant and meaningful technology can be created by giving people who are affected by it some role in its design” [33]. We adopt a participatory approach aligning

with our goal to enable non-experts to engage in creating complex, nuanced dialog behaviors for social robots.

Within the HRI community, participatory design research explores varying degrees of user involvement, for instance, involving children in discussion of ethical concerns about robots [70]. Mott et al. investigated how participatory design methods—specifically a workshop with design activities—could empower children to actively engage in co-design of social robots [70]. In the workshops, children participated in group discussions and a design activity to create prototype human–robot dialog. Our work adopts a similar methodology to involve users through elicitation (discussion) and hands-on design with paper and subsequently with a software prototype. Björling et al. employed a “human–centered and participatory approach” to investigating teens’ perceptions of social robots and how robots could be designed to support their emotional needs [12, 13]. In this case, the teens’ participation involved not designing a robot nor its behavior but rather soliciting their thoughts and ideas from interactions with robots to assess the effectiveness of an early prototype. This contrast in approaches highlights the broad range in user involvement across participatory design methodologies. In our participatory approach, we involve users heavily—from the earliest, pre-interface stage—in expressing their perceptions and ideas about backchanneling robots and engaging creatively to customize their own robots, including discussions, prototypes on paper and designing robot behaviors with a software prototype. Meanwhile, our interface design was informed by our findings from early user participation, but our participants did not directly participate in interface design decisions.

In the broad landscape of participatory approaches, some work has called for accountability in how participatory design research is carried out and reported [33]. Frauenberger et al. argue that,

given the “messy”, “highly contextualised” processes and range of implementations, researchers must critically and transparently reflect on the nature of their participatory design approach, to advance the field and promote rigorous participatory practices, rejecting a positivist lens [33]. They propose a tool for reflection through four lenses: *epistemology*, *values*, *stakeholders* (interested parties), and *outcomes*. One particularly-relevant lens for our work is that of interested parties and the nature of their participation—we did not engage our participants directly in decision-making for interface design, or even whether we should enable end-user design of robot backchanneling, rather, we grounded our design in their perceptions and natural tendencies elicited during the participatory process. This aspect of our work was not participatory—our participatory approach instead involved gathering user perceptions, design tendencies and feedback, and building a system based on how users engaged the concepts.

2.2 End-user Creation in Human–Computer Interaction

Our work follows a theme in human–computer interaction (HCI) of designing tools to enable end users to participate in the development or modification of systems, or, *end-user development* [60]. Among the many benefits of end-user development, our work emphasizes the importance of *tailoring* systems to meet a user’s needs and preferences through customization [63], and some argue that “its most important and far-reaching impact will be to transform cultures by empowering all people to become active contributors in personally meaningful activities” [31]. Research continues to investigate techniques to enable end-user participation in designing or modifying systems by alleviating challenges and barriers [63]. For instance, programming-by-demonstration, where the user models system behavior by providing examples (e.g., programming robot movements by moving one’s own body [34]). Another popular method for

enabling end user participation is visual programming, where programming concepts are simplified into manipulatable objects (e.g., Google Blockly²), enabling the user to focus on learning and applying core concepts without barriers such as code syntax, compilation, obtuse error messages and heavyweight development software [102]. In this section and the following, we outline examples of end-user tools for creative authoring.

An active area of research in HCI is enabling end users to author content and be creative with computing technologies. End-user creation in HCI research includes a rich history of developing novel interfaces and toolkits, for instance, for authoring 3-dimensional models from 2D sketches [49], a physical widget-based toolkit for designing physical interfaces [39], or a tool for rapidly prototyping puppet movement style and personality [104]. The application space is quite broad and remains an active area of exploration with cutting-edge computing technologies, such as enabling users to design video games by engaging with AI chat bots [6], or to create complex non-player character behaviors by painting storyboard-like scenes [80]. There has been particular interest in visual programming, for example, drag-and-drop interfaces that enable people to write computer code by arranging virtual “blocks” [88, 102], or providing graphical interfaces to increase user control over systems such as large language models (LLMs) whose outputs can be difficult to wrangle to achieve the desired outcome [35, 66]. Research efforts extend beyond digital artifacts and into physical, personal systems, such as custom, user-designed wearable glucose monitoring devices [53]. This body of work highlights the continued interest in—and importance of—abstracting complex technological systems into tools accessible to non-expert users.

² Google Blockly: <https://developers.google.com/blockly>.

2.2.1 *GlucoMaker: Designing Your Own Glucose Monitor*

In a recent example, Lakhdhir et al. developed GlucoMaker, a toolkit for people with diabetes to collaboratively (with professionals) design their own custom glucose monitoring device [53]. The GlucoMaker toolkit supports end-user participation in designing a highly personal, body-worn device (a glucose monitor), a technology that may otherwise pose barriers for many end users to design themselves. Our iterative methodology reflects that of Lakhdhir et al.'s, both in the preliminary focus group study to understand desired features and preferences, including physical crafting activities for early prototyping, and in a subsequent study inviting users to digitally customize a model for a glucose monitor in consultation with experts (Figure 4) [53]. By empowering end users of glucose monitors to participate in the design process for their own custom device, GlucoMaker supported end users with varying ability to creatively author a highly personal, tailored device according to their preferences, a thread which we continue in the present work. In our work, we enable customization and tailoring of robot backchanneling behaviors with the goal to promote successful human-robot interaction that meet users' expectations, needs and preferences.

2.3 End-user Authoring for Human-Robot Interaction

In human-robot interaction, end-user authoring tools provide users with agency via customization and tailoring of robot behaviors. For instance, enabling home users to specify cleaning robot routines by sketching its actions on video feeds of their home environment [8, 86], or enabling people with mild cognitive impairment to participate in designing their own robot companion [20]. These authoring tools provide simple ways for a user to exercise control over complex behaviors, such as sketching robot movements [86], or the style of *how* a robot




Figure 4: GlucoMaker enables users to design their own custom glucose monitoring device using a digital interface, which is then brought to life and worn by the user (taken from [53]).

moves and follows the user, which can be easily authored via demonstration by physically moving the robot [105]. One educational robot enabled children to tailor its affective behaviors with a block programming interface [36], while another system enabled the user to select personas for the robot's speech [106]. Our work continues this ongoing thread by investigating how non-expert end-users can participate in designing custom robot dialog and backchanneling, using a novel interface design.

End-user creation in HRI focuses on supporting the user to author behaviors using “intuitive, natural and enjoyable” techniques such as by using speech to program robots [37, 48] or authoring proxemic interaction using their bodies in a spatial environment [65]. This is commonly achieved by translating technical robotics concepts (e.g., spatial sensor data [65]) into languages that users can understand and directly manipulate to describe or “program” a desired behavior. For instance, sketching [86], voice input [37], kinesthetic demonstration [20], use of spatial environments [65] and others, including one interface that used image manipulation on a timeline [64], while another leveraged tangible tabletop puppetry [77]. These tangible, physical approaches can have limitations, however, as it can be difficult to describe complex actions that occur over time, or to exercise control to fine-tune patterns of dynamic behavior, for instance listening behaviors in rich dialog. In this work, we take design inspiration from popular graphical user interface techniques for authoring robot behaviors, including block programming.

2.3.1 Block Programming Interfaces for Robots

On-computer interfaces such as block-based programming interfaces are used extensively by hobbyists and in education to teach programming [45, 88, 101, 102, 108], and have become an industry standard for non-expert programming. Block-based programming interfaces are prevalent in software for dozens of commercially-available robots [25] and in research (e.g., [5, 20, 35, 36, 81, 87, 99, 100, 106]), as they enable a high degree of manipulation and programming-like control over complex robotic systems in an approachable interface paradigm that is easy to learn and use for non-experts, including students, hobbyists, designers and out-of-field



Figure 5: The PODER framework leverages block programming in the conceptual and implementation stages (tangibly and digitally) to support non-expert end users in ideating and creating robot behaviors (from [20]).

researchers. This approach enables inexperienced users to quickly get started programming robots in a simplified, linear fashion and rapidly create working systems without requiring extensive education or programming expertise.

PODER: Ease of Use for Design Participation

Cruz-Sandoval et al. leveraged block programming in their tool for enabling increased participation from people with mild cognitive impairment in HRI research [20]. The PODER framework supports users to prototype behaviors using a mix of analog techniques, including think-aloud, kinesthetic demonstration on a robot, and block programming on paper, and digitally using block programming to define behaviors. PODER is intended for pair programming alongside a researcher to enable flexible, dynamic sharing of tasks based on individual needs and capability (Figure 5). This approach enabled individuals with mild cognitive impairment to increase their confidence and knowledge with respect to technology, to imagine how robots could be helpful to them in new ways and to feel *included* in the design of these systems [20]. While a block programming interface is only one aspect of the tool that contributed to these

outcomes, we follow this popular interface choice for HRI toolkits to leverage the proven learnability and approachability that enables users without technical expertise to engage with what may otherwise be prohibitively complex or abstract concepts [20]. Given our goal to augment robot behavior authoring tools with a novel interface enabling non-expert users to create backchanneling, we adopt a block-based approach for our interface design.

2.3.2 Authoring Robot Dialog

A contemporary approach to providing rich robot dialog is to use large language models or other generative models, including for generating speech [91, 98] and backchanneling [54, 55, 71]. However, there is still a gap in addressing the key limitation of enabling end users to fine-tune the LLM to achieve the desired result. Researchers are developing tools to combat this lack of user control over LLM outputs, for example, structuring LLM prompts with Blockly [35, 62], but this is just one paradigm for behavior design—current state-of-the-art authoring tools (commercial and in research) are still highly structured and provide handcrafted, high-control approaches. Regardless of the underlying implementation of dialog (LLMs or not), there is not yet a way for non-expert end users to create highly personalized, custom backchanneling for a robot.

Current dialog authoring tools often enable user construction of robot speech via text-based input [23, 26, 99], with some tools offering additional customizations, such as selectable robot personas that affect speech [106] or specifying expressive speech using emojis [97]. However, such interfaces typically focus on speech—both content and expressivity—and do not support user customization of backchanneling. We found one system for authoring multi-party AI chatbot interactions that supported enabling or disabling generated backchanneling for each

interlocutor, but the system was geared toward professional designers [47]. No system that we found provides end-user control over robot backchanneling. Our work builds on these dialog tools by exploring how to support end-user authoring of backchanneling.

2.3.2.1 EmojiVoice: Expressing Robot Speech with Simple Heuristics

For authoring intricate, dynamic social behaviors in robots, using simplified representations of behavioral patterns may be a powerful tool to enable participation by non-experts and rapid prototyping for advanced users. Tuttösi et al. developed EmojiVoice for authoring expressive robot speech for human–robot interaction using emoticons [97]. This text-to-speech toolkit targeted to interaction designers and roboticists enables authoring customizable speech by prompting the system with emoticons alongside text. EmojiVoice enabled users to create highly varied, contextually appropriate expressivity in robot text-to-speech output, and improved user perceptions of the robot [97]. In our work, we take some inspiration from this design pattern, in tandem with findings from our elicitation of user perceptions of robot backchanneling, to provide *personality* options as a means of customizing robot backchanneling in human–robot dialog.

2.3.3 Granularity of Control

Many interfaces provide a range of granularity levels in their authoring controls, using various approaches. For instance, block programming approaches can offer fine-grained authoring with a high degree of control [106], while others focus more on offering coarse-grained authoring [17]. In the same vein, some create mixed-initiative systems [46] where the system assists with fine details [37, 47]. Zhang et al. investigated the trade-off between user control and perceived robot agency through coarse- and fine-grained authoring interfaces (Figure 6), finding that users preferred granularity in moderation [106]. Their interfaces included a standard block




Figure 6: Zhang et al.'s end-user programming interfaces featured two levels of granularity for behavior authoring: a high-granularity block programming interface (not shown), and a low-granularity interface featuring simple slider controls for basic parameters, and a set of available "personas" to choose from (from [106]).

programming interface as a high-granularity means of authoring, and a low-granularity interface featuring sliders for basic parameters (e.g., movement speed, talking speed), along with a “vocal persona” option, selected by choosing one of the available persona avatars to illustrate the persona the robot will adopt (Figure 6). Personas include, for example, “surfer dude”, “business woman”, and “robot”, and are also provided in the high-granularity block interface as a dropdown in an overarching encompassing other behaviors [106]. We follow Zhang et al. [106] and Tuttösi et al. [97] in our work to expose nuanced backchanneling using personality and emotion as an abstraction.

In our work, we explore granularity of control for designing robot behaviors all throughout our design process, including at the elicitation stage and in our final interface design. In particular, our work aims to provide varied granularity to make customizing robot backchanneling approachable for non-experts, while providing fine-grained authoring controls for users who want fine-grained control over their creations. We follow this thread of exposing simplified controls for complex robotic behaviors by augmenting interface designs for robot dialog authoring with simple controls to support user authoring of robot backchanneling behavior.

2.4 Backchanneling

In human dialog, interlocutors often perform additional communication alongside primary speech—backchanneling—where an interlocutor who does not currently hold the speaking floor communicates with other interlocutors—often while the current speaker is speaking—using gestures or short utterances that aim to avoid interrupting speech on the primary channel [76]. For example, listeners indicate attentiveness and encourage interaction using utterances (e.g., “hmm”, “sure”) and gestures (e.g., nodding, facial expressions). Backchanneling serves many important functions in dialog, including to convey attention, feedback, interest or boredom, coordinate turn-taking and rhythm, and build rapport [58, 85, 90, 92]. Further, backchanneling behavior varies greatly by individual and context [15, 27, 103], highlighting the need for flexibility and variability in how robots incorporate backchanneling into their dialog behaviors, and potential for highly specific customizability and personalization of robot backchanneling. This section details the state-of-the-art in backchanneling robots and tools for authoring robot backchannel behavior.

2.4.1 *Backchanneling Robots*

In HRI, researchers are enhancing social robot dialog capabilities for a variety of interaction benefits. Most implementations are focused on interaction effects, and thus consist of bespoke, researcher-designed backchanneling—without end-user involvement—which is used to study how robot backchanneling impacts a particular kind of interaction. For instance, Park et al. found that making a robot’s backchanneling contingent upon children’s nonverbal cues led to children preferring it over a non-contingent backchanneling robot [76]. Other work has focused on using robot backchanneling to encourage participation in group tasks from individuals who may not otherwise be as vocal. Sebo et al. found that supportive robot backchanneling toward outgroup members who felt excluded increased their verbal participation [95]. Likewise, Cumbal et al. observed a similar effect in group tasks with language learners and native speakers—when a robot performed backchanneling *adaptively*, meaning it backchanneled more often when the least active speaker was speaking, these individuals increased their verbal participation significantly [22]. In all these works, the focus was on either developing accurate backchanneling models or how such models may impact human–robot interactions, and thus no details were provided on how these systems could be easily tweaked or how end users could be involved in system design. Our contribution is a novel angle for backchanneling in HRI: including backchanneling as a component of end-user creation for robot behaviors.

Outside of robot implementations, end-user customizable backchanneling is rare. In one work, Hu et al. [47] developed a toolkit for designing human–AI group conversations with a simple backchanneling feature—select the source and target AI interlocutor, and specify an emotional tone—but it was not possible to customize other aspects of backchanneling, such as the frequency of backchannels and the presence and nature of individual behaviors such as nodding, gaze, and

verbal utterances. We found no tool in the literature or in the commercial sector that enables end users to specify or customize robot backchanneling behavior. In this work, we address this gap and present a novel tool for end-user design of robot backchanneling.

2.5 Chapter Summary

In this chapter, we illustrated a recurring theme of HCI research on end user participation and empowerment in designing and shaping technologies, including robots and other systems. Enabling participation from a broader range of users, rather than just in-field experts, can empower non-experts to customize and tailor system attributes and behavior to their needs and preferences. Many end-user focused authoring tools enable people to program robot behaviors, but we did not find any that supported user design of robot backchanneling. Given that backchanneling is widely investigated, implemented, and impactful for social robots, this gap is a key design element of social robot behavior that remains largely inaccessible for non-experts. Our work draws from the landscape of end-user creation technologies and foundational human-computer interaction principles of user-centered design and participatory design to investigate how we can enable non-expert end users to create custom robot backchanneling. In the remainder of this dissertation, we blend these themes together throughout our methodology in our approach to investigating how we can make authoring robot backchanneling accessible by non-experts.

3 Elicitation Study: Understanding Non-Expert Mental Models

We set out to gain insight into how non-experts might conceptualize backchanneling, which elements they may understand, which may be challenging to understand, and how they engage with ideas about robot backchanneling (RQ1). Further, we sought to investigate how non-expert potential users of robots may describe and create backchanneling for a robot, including what language, metaphors and abstractions they use (RQ2). Once equipped with this knowledge, we could possess an understanding of user mental models to inform our next step of designing an interface for a robot backchanneling authoring tool targeted for use by non-expert users. First, to elicit this knowledge, we needed to ask potential users of robots about these topics and observe how they engage these concepts, including the creative task of designing robot backchanneling. Therefore, we concocted an elicitation study to learn about user mental models about robot backchanneling.

Our elicitation study comprised workshops with groups of participants where we asked them to complete a set of individual and group-based activities related to robot backchanneling, which allowed us to gain insight into their mental models. Through these activities, we explored non-experts' mental models relating to robot backchanneling, including how readily and comfortably non-experts engage with concepts of robot backchanneling, which elements they can readily articulate and which they may struggle to understand or apply. We aimed to investigate what roles non-experts perceive backchanneling can play in dialog, and how they may naturally describe elements of backchanneling including what vocabulary, metaphors, or comparisons they may use. We further investigated how granular people tend to be, e.g., specific behaviors versus abstract personality traits.

3.1 Tasks and Procedures

We conducted two-hour workshops (Figure 7) with groups of participants at a boardroom in our institution or at a satellite location (based on availability), recruited using posters placed around



Figure 7: Design workshop space showing a mock session in progress (we did not have permission for photos of participants).

our city and campus that targeted people “who are interested in designing robots, even if you don’t have any technical skills” (recruitment poster and email communication in appendices A.1 and A.2). Following standard focus group methods for session size [24, 52] we targeted 6-10 participants per session. We opened each workshop with refreshments, an informed consent procedure (Appendix A.3), and administered a demographic questionnaire (Appendix A.4) which included an open-ended question on prior experience with robots. We compensated each participant with \$35 CAD at the beginning of the session. We formed groups of 3 or 4 participants and began with a short ice-breaker activity. Following we administered a series of design tasks informed by participatory design and mental model elicitation methods as described below [21, 40, 59, 89, 109] (Appendix A.5). This research was approved by our institution’s research ethics board.

Unprimed Perceptions of Backchanneling—First, we conducted an unprimed task to explore non-experts’ perceptions of backchanneling without having it explained to them. We produced and showed two short videos of conversations, one naturally between two people (which included backchanneling), and another between a person and a robot which represented typical human-robot interaction, i.e., the robot did not backchannel. We showed the videos back-to-back and

asked participants to play “spot the difference”, brainstorming individually on paper and then in small groups. Following, we invited groups to share their ideas, culminating in a list of behaviors and features of the interactions in the videos, and researcher notes on ensuing discussion, all of which can be analyzed to understand non-experts’ language for describing backchanneling and their perceptions of the roles backchanneling played in dialog.

Primed Backchanneling Brainstorm—Next, we elicited participants’ ideas for potential robot backchanneling behaviors and their engagement with potential reasons for using backchanneling in dialog. We started this task with a short lecture (5 minutes) to formally introduce backchanneling as feedback behaviors that people use in conversation, including examples of behaviors and their functions in human conversations (e.g., acknowledge people, encourage them to speak, listen attentively). We then asked participants to complete a standard card sort task [11, 21, 59, 107, 109] where they brainstormed potential robot backchanneling behaviors on index cards (we gave an example, the utterance “hmm”) and sorted cards into themes representing potential purposes of backchanneling, providing categories based on backchanneling literature: paying attention [56], understanding [44], agreement and disagreement [50, 61], engagement [78], taking turns [18, 73], clarification [76] and building rapport [67]. For flexibility, we followed a hybrid card sort methodology allowing participants to create their own themes [19]. This resulted in a comprehensive list of participant-identified backchanneling behaviors, their categorization of these elements into purposes, and the language used to articulate the concepts.

Authoring Robot Dialog and Backchanneling—To explore how non-experts may approach authoring robot dialog with backchanneling, we asked participants to complete a design task to

individually author robot dialog on paper and subsequently author backchanneling for the dialog. We first asked participants to create conversations between a person and a robot (20 minutes), and then later asked them to consider how their robot might backchannel, how they might indicate this backchanneling for their creations, and to author backchanneling as they saw fit (20 minutes). We provided worksheets (details below), arts and crafts stationery (e.g., colored markers, sticky notes, scissors and glue) and an optional sample scenario to help people get started.

To assist participants in creative authoring, we designed two worksheets on which to describe human-robot dialog and general characteristics of their robot and its behaviors (Figure 8). The first sheet, left, contains dialog bubbles representing a conversation between a person and a robot, on which speech by the two interlocutors can be authored, with spare space to indicate other aspects of dialog. The second worksheet (Figure 8, right) was included to support brainstorming general robot characteristics before or during prototyping dialog, giving participants a workspace to engage these ideas. This “character sheet” is mostly blank to avoid steering design, with only light prompts at the top of the page, such as “how does your robot act?”. We provided blank copies of each worksheet, along with partially filled samples, shown in Figure 8, for the case where participants struggle to start from scratch or simply prefer to work from an existing baseline robot. This task resulted in artifacts of participant-authored robot dialog behaviors including backchanneling, including the concepts used, the vocabulary to articulate behavior, and the layers of granularity on which participants conceptualized backchanneling and authored these behaviors for dialog.




Figure 8: Worksheets for authoring dialog included pre-filled samples and blank copies with black text and bubbles only.

Group Discussion—We closed out the workshop with a brief semi-structured group discussion to summarize participant perceptions of robot backchanneling and authoring backchanneling for robots, and any further commentary, from which we collected observational notes. Guiding questions used as launching points for the group discussion are available in Appendix A.5.

3.1.1 Analysis Plan

We conducted thematic analysis of the resulting materials—including participant-generated paper artifacts, card sorting results, discussion-generated idea lists, transcribed audio recordings, and researcher notes—using a mixed deductive-inductive coding approach [43, 52]. The author and their advisor developed initial codes deductively based on research questions outlined in Chapter 1.2. (specifically, RQ1 and RQ2), after which the author conducted an initial coding pass of the data. As patterns began to emerge, the author iteratively refined codes and added new codes introduced inductively, completing our hybrid approach as emergent patterns solidified.

The author, in regular consultation with their advisor, synthesized key themes on our participants' non-expert mental models.

3.2 Elicitation Study Results: Non-expert Mental Models of Robot Backchanneling

Recruitment resulted in 27 adults (Table 1), 12 women, 15 men, one identifying as non-binary/trans/gender diverse, across four workshop sessions (max 10 participants per session, min 4 due to unexpected participant absence). Most people (17/27) were in the 18-29 age range, with two aged 30-39, three aged 40-49, one aged 50-59, three aged 60-69 and one aged 70-79. Twenty participants indicated little to no prior experience with robots or designing robot behaviors (e.g., 6 participants had used Roomba vacuums or restaurant service robots). The remaining 7 indicated prior experience designing or programming robots, including for hobbies (e.g., Lego Mindstorms, 3 people), undergraduate courses (3 people), and professional experience (teaching a high-school robotics class, 1 person). One participant declined to be audio-recorded (Session 1) resulting in that session not having audio transcriptions. Our thematic analysis resulted in key themes regarding participants' mental models of robot backchanneling and how they approach authoring backchanneling for robots. Note that we refer to whole sessions when results are from group activities, and individuals otherwise.

3.2.1 *Knowledge of Backchanneling*

Participants recognize and understand backchanneling—During our unprimed spot the difference task, before we explained backchanneling to participants, all participants demonstrated rich understanding of backchanneling, its purposes and potential behaviors, even if they did not use the term specifically. When comparing the human-human discussion video (natural) to the

human-robot one (no robot backchanneling), groups noted that the humans used *acknowledgment* behaviors (Session 1), *indication of listening* (S2), or *filler words* (S3, S4).

Participants further noted

Table 1: Participation across all three studies, broken down by each study and each session.

Participant	Elicitation				Prototyping			Evaluation	
	S1	S2	S3	S4	S1	S2	S3	S1	S2
1									
2									
3									
4									
5									
6									
7									
8									
9									
10									
11									
12									
13									
14									
15									
16									
17									
18									
19									
20									
21									
22									
23									
24									
25									
26									
27									
28									
29									
30									
31									

32									
33									
34									
35									
36									

how backchanneling supported engagement and turn-taking, with many noting the importance of small interruptions and overlaps often found in human conversations (P6, P9, P10, P15, P24, P27): for instance, *“when the humans were talking, lots of nodding and other non-verbal cues that keep the conversation flowing”* (P7). Participants further identified that the robot lacked backchanneling *“it wasn’t really showing active listening, because it wasn’t, it didn’t have any body language unless it was talking”* (P9). Many noted negative impacts on interaction, for instance, *“can it not just tilt its head or whatever and show that it’s listening somehow? ... it just seemed odd to me that it wouldn’t move while listening.”* (P9). One participant succinctly stating *“It needs more pizzazz”* (P23). After we introduced backchanneling (card sort task), this theme of demonstrating clear and nuanced understanding of backchanneling continued.

Participants described impacts of backchanneling—Throughout the workshop, participants relayed their thoughts on the impacts of backchanneling, especially on taking turns in conversation. In contrast to the card sort task, where *turn-taking* accounted for just 12 of 227 items categorized (5.2%), participants readily explained how backchanneling affects turn-taking. One participant explained backchanneling for taking the speaking floor, *“even just kind of saying ‘well yeah’, you know, just kind of giving yourself... ‘this is my turn to speak’”* (P9). Further, P7 noted how backchanneling helped the dialog flow: *“when the humans were talking, lots of nodding and other non-verbal cues that keep the conversation flowing”*. Multiple participants commented on the impact of small interruptions and overlaps often found in human conversations (P6, P9, P10, P15, P24, P27), for instance, P24 articulated how talking with a robot differs from the turn-taking

dynamics of siblings, *“that’s how like a natural conversation between siblings would go is they would be talking over each other, you don’t really get that when you have robots because they have to stop and listen”* (P24). Another participant in a different session described turn-taking dynamics in the demo videos, *“the person would interject sometimes, and it had to do with those filler words... the other person I guess technically interrupted you to agree, whereas the robot only spoke when it was clear that you were finished”* (P15). Participants articulated awareness of the role backchanneling plays turn-taking for dialog.

3.2.2 Non-experts’ Considerations when Authoring Backchanneling

Participants readily describe desired backchanneling behaviors—In our robot dialog creation task we observed most participants to readily engage with authoring dialogs that include backchanneling, observing no difficulty or confusion with the concept. Only 9 (of 27) did not indicate any robot backchanneling behavior but rather explained higher level behaviors where backchanneling may be implicit. Participants annotated dialogs with descriptions of backchanneling (6 participants), e.g., *“robot interjects with laughter”* (P15) or *“uses filler words (hmm)”* (P16). Participants also described non-backchanneling behaviors, for example, *“*sarcastic tone*”* (P7) or *“moves arms around excitedly in wide, unnecessarily large gestures”* (P1, Figure 9).

Participants typically annotated dialogs with written descriptions, and one (P24) drew robot




Figure 9: Prototype robot dialogs and characters, annotated with backchanneling and other dialog behaviors (P25, P6, P1, P7).

facial expressions (Figure 10). For verbal backchanneling, it was most commonly (15 participants) incorporated into robot speech text, for example, “*hmm... let’s see!*” (P16). Our analysis did not uncover any predominant style or technique that we could use to inform specific interface designs.




Figure 10: We tasked participants with authoring robot dialogs with backchanneling on paper to elicit their mental models for designing robot backchanneling behavior. From left to right: P24, P11, P1.

Appropriate use of realistic behaviors—Although participants attributed backchanneling behaviors to robot realism, there was significant discussion over the downsides of realistic robots, “*I struggle with the idea that, partly, it’d be good to have a robot that looks sort of more human because they could be more expressive, but I think that would creep me out a bit, do you know what I mean? I’m not sure what the balance is to that.*” (P21). Many noted wanting to make robots undeniably robotic to avoid confusion (P2, P9, P11, P21, P23, P25), with two (P2, P27) noting the “uncanny valley” [69], “*mannerisms aren’t too human, avoids uncanny valley*” (P2) and one referring to media “*you should always be aware that a robot is a robot... I don’t know, I’ve watched Bladerunner*” (P11). Some noted specific desires for *appropriate* robot backchanneling to mitigate this (P2, P10, P11, P22).

Participants author backchanneling as personality traits—We observed 26 (of 27) participants to describe high-level general descriptions of robots’ tendencies and personality traits, to complement the dialog, such as “*she is very affectionate, loving and compassionate*” (P26). Many of these were directly related to how the robot would listen (P1, P2, P6, P10, P23), for instance, “*he likes to listen but chatters on a bit*” (P1), “*always listens to you carefully*” (P10), or “*good listener, able to draw out responses with open-ended questions*” (P23).

Authoring using human-like language—some participants drew comparisons to human behavior to explain desired behavior, for instance, directly comparing how a robot behavior may resemble human-like behaviors (P2, P3, P5, P10, P11, P15, P21, P22). For example, participants described robots as “*humanoid-ish*” (P11), “*not a wannabe human*” (P10), that its’ “*mannerisms aren’t too human*” (P2), or even just, “*a bit robotic*” (P22). One participant made direct reference to being as human-like as possible, writing, “*Amy acts like a human. Amy wants to be human*” (P5).

Importance of context of use for authoring—Participants commonly grounded their descriptions in the robot’s role or context of use, such as envisioning a service robot to support continued independence and make their lives easier (P7, P9, P21, P22). Participants noted the importance of this context for authoring appropriate dialog, for example, “*what is the robot’s overall purpose in your life? Because that’ll affect how they’re going to talk to you*” (P15).

3.2.3 Fuzzy Boundaries: Where Non-experts Transcend Backchanneling

Participants considered backchanneling broadly—While our participants demonstrated awareness and understanding of backchanneling, they also considered other aspects of dialog and general robot characteristics. For instance, during the *Primed Analysis of Backchanneling* task, participants listed backchanneling behaviors and categorized them by purpose, resulting in 227

Table 2: Our categorization of participant-created backchanneling behaviors from the card sort.

CATEGORY	SUBCATEGORIZATION	EXAMPLES
backchanneling	fine-grained	“tilt head” (P24-27), “laughter” (P11, P12, P14), “reinforce with ‘mhm’” (P24-27)
	coarse-grained	“body language” (P21-23), “appropriate facial responses (surprise, anger)” (P13, P15, P16), “use of filler words” (P17)
	adjacent	“ask follow up questions” (P21), “interrupt” (P9)
not backchanneling	personality	“tone to indicate meaning (like sarcasm, interest, joking)” (P8) “caring compliments” (P17)
	capability	“try to anticipate when responding”, “make the robot bigger” (P13, P15, P16)
	interaction	“pivot a conversation” (P21-23), “robot call you by name” (P21-23)

combinations of behaviors and purposes (Table 2). We analyzed these items and found 167 backchanneling behaviors across 3 categories, including 85 items (37.4%) representing *fine-grained* behaviors such as “smile” (P24), 44 items (19.4%) representing *coarse-grained* behaviors such as “body language” (7 participants), and the remainder were *adjacent* behaviors including interruptions and asking questions (38 cards, 16.7%).

To understand where our participants’ conceptualizations escaped the boundaries of backchanneling, we analyzed the remaining 60 items that were not related to backchanneling, for instance, “make the robot bigger” (group of P13, P15, P16), and “reload its weapon :)” (P7), resulting in 3 subcategories (Table 2). These included *personality* behaviors, such as tone, expressions, and empathy (24 cards), *capability* behaviors such as the robot’s size, ability to anticipate speech, and memory of past interactions (15 cards), and *interaction* behaviors including ending the conversation or calling the user by name (21 cards).

Our findings indicate that when shown backchanneling in dialog, our participants demonstrated awareness and understanding of the concepts and behaviors, however, when we asked our

participants to brainstorm backchanneling behaviors for robots, they identified myriad backchanneling behaviors but also produced a wide range of behaviors outside of backchanneling and dialog.

Some purposes and impacts stood out—Participants demonstrated understanding of a range of purposes and impacts of backchanneling in dialog. The *Primed Analysis of Backchanneling* task resulted in 227 backchannel actions written on cards, which participants sorted into 9 pre-defined categories and one participant-created category, each representing a purpose of backchanneling (as explained in Chapter 3.1). Participants assigned their backchanneling ideas most frequently to the *engagement* (23.7% of all cards) and *paying attention* (14.5%) categories, suggesting that these purposes may be more clearly recognized by non-experts.

We analyzed and categorized all cards for their relevance to backchanneling, and our subsequent analysis revealed that these more commonly identified purposes contained a higher proportion of cards that represented fine- or coarse-grained backchanneling behaviors. For *paying attention*, 27 of 33 cards (81.8%) described such behaviors, and for *engagement* this number was 36 of 54 (66.6%). Other purposes of backchanneling like *turn-taking* (5.2% of all cards) and *clarification* (7%) were less commonly identified by participants for the items they brainstormed. These purposes also had a lower frequency of fine- and coarse-grained backchanneling among the items, with 2 of 12 (16.67%) for *turn-taking* and 3 of 16 (18.75%) for *clarification*, meaning participants assigned items that were not backchanneling behaviors to these categories more frequently. We note that turn-taking may be technical terminology less familiar to end-users than to researchers interested in dialog. Our findings indicate participants thought of

backchanneling primarily to indicate *engagement* and *attention*, while other purposes were less clear.

3.3 Elicitation Study Discussion

Our analysis revealed key insights into non-experts' understanding of robot backchanneling. First, our non-expert participants readily understand and engaged with backchanneling and readily demonstrated awareness of its purposes and impact on dialog. Most participants designed backchanneling for their robots, employing a variety of techniques and approaches to do so. Further, participants articulated a wide range of backchanneling styles, applications, and behaviors using rich and varied language, and readily demonstrated awareness of its purposes and impact on dialog. We further highlight how participants went *beyond* backchanneling *per-se* and focused more on nuanced robotic dialog behaviors, for example, sarcasm and humour, varying tone of voice, smooth turn-taking and attentive listening. Further, participants had a strong desire for *appropriate* backchanneling and had differing views on what kinds of backchanneling—if any—would be suitable for robots. These results highlight just how well participants understand the nuances and potential uses of these techniques and support our approach of enabling non-experts to author nuances of robot dialog, including backchanneling behaviors.

In describing backchanneling behaviors, participants used a mix of coarse-grained approaches (e.g., considering role, context and personality), and fine-grained approaches such as annotating dialog with individual behaviors (Figure 4). From the high level, participants frequently referred to personality traits and human-like emotions when describing robot behaviors and even backchanneling specifically, both during group discussion and in their design artifacts. Some

focused heavily on authoring their robot’s personality to inform their dialog design. From the low level, participants applied various individual behaviors in their authoring, including nodding, verbal utterances, facial expressions and other gestures (e.g., nodding slowly while listening, waving arms around). These results provide a window into how non-experts understand backchanneling, as a mixture of both personality characteristics and specific gestures and dialog behaviors.

Our results highlight the importance of context of use, application, and even participant preferences on how a robot dialog should be designed. Many participants grounded discussions and creations in the robot application and included elements such as preferring a robot to not be too realistic, or how a robot should act “appropriately” based on context. This supports our motivation of enabling creators to have flexibility in designing nuanced human-robot dialogs.

3.3.1 Reflection on Methodology

We analyzed comments from participants during the *Group Discussion* task at the end of the workshop to reflect on our method. Our participants encountered struggles mainly in prototyping their own robots during the *Designing a Backchanneling Robot* task. Three participants noted having difficulty in authoring behaviors while considering realistic limitations, due to uncertainty of possibilities (P22), external media influences (P15), or their programming knowledge (P7). “*I have a hard time differentiating between, how idealistically could you have a robot behave to what's actually practical for a robot,*” (P22). Others expressed struggling with the semi-structured nature of the task and designing their own robot. For some, having a clear understanding of their robot’s purpose was essential to understanding how to author its dialog.

While we carefully opted for a less-structured approach in the *Designing a Backchanneling Robot* task to let participants explore and avoid oversteering them toward authoring certain behaviors, more structure may have enabled them to engage more directly in backchanneling design, without needing to first address the robot’s purpose, character and human-like attributes. For instance, providing more structured ideation through a heavier-weight co-design method [4], or by requiring participants to use the provided sample dialogs and robot character.

3.4 Design Recommendations for Dialog Authoring Tools

We synthesize the results from our elicitation workshop into design recommendations for non-expert authoring tools for human-robot dialogs design, recommendations which we use directly in the following sections in creating our own interface.

1. Design Dialog Authoring Tools to Include Backchanneling. We found people to be acutely aware of backchanneling and its impacts on dialog. Non-experts engaged heavily in discussion and design activities on robot backchanneling and noted the contextual nature of these behaviors. All this highlights that dialog authoring tools should include backchanneling as a core feature.

2. Richer Dialog: Beyond Backchanneling. We observed participants to focus on rich dialog broadly (beyond just speech), including personality, empathy, tone of voice and emotive speech. Thus, while we still advocate for the necessity of backchanneling in authoring tools, where possible also include other ways of enriching dialog, such as specifying the robot’s personality or enabling emotive speech [97].

3. Support Personality-Level Control of Backchanneling. Participants often considered backchanneling patterns and tendencies at a high level, such as robot personality or emotions, not necessarily using specific robot behaviors. These participants considered backchanneling as coarse-grained listening styles and described them using the robot’s personality traits. For instance, employing sarcasm, humour, affection, or being less agreeable. We provide mechanisms to support authoring robot behaviors using the robot’s personality.

4. Support Contextual Customizability. Participants identified how backchanneling will vary based on interlocutors or contexts, highlighting that backchanneling for a given robot is variable and not fixed. We provide the ability to vary the robot’s backchanneling behavior for one type of context: different segments of a dialog.

5. Support Behavior-Level Fine-Grained Control of Backchanneling. Participants readily understood specific backchanneling behaviors and techniques, ranging from utterances to gestural actions and facial expressions. Although many defined backchanneling at the higher, coarse-grained level, they do understand the lower-level mechanics and may want to adjust or customize at this level, e.g., tune the frequency or intensity of nodding. Therefore, authoring interfaces should provide fine-grained control of backchanneling to support authoring specific behaviors.

3.5 Chapter Summary

In the elicitation study, we set out to investigate how non-experts conceptualize backchanneling (RQ1), how they describe its elements and features (RQ2a) and how they may approach designing backchanneling for a robot (RQ2b, RQ2c). We conducted workshops with non-experts, completing elicitation and design activities to learn about how they approach robot

backchanneling conceptually and for authoring robot behaviors. Our findings show that non-experts intuitively understand and engage with backchanneling, including the associated behaviors, its impact on dialog interactions and how backchanneling behaviors can affect one's perceptions of others and robots. Non-experts engage backchanneling when designing robot dialog, including at the fine-grained behavior level and on the broader, coarse-grained layer, considering robot personalities, emotions and other complementary dialog behaviors.

Our findings reveal key implications for dialog authoring tools. We uncovered that potential non-expert users of robot behavior authoring tools are capable and ready to engage with backchanneling in robot behavior design tasks, and we suggest that robot behavior authoring tools should include backchanneling design options on multiple layers of granularity (e.g., personality). We synthesized these results into design recommendations for dialog authoring tools, and proceed by leveraging these recommendations in subsequent chapters to ground our interface design and prototype evaluation in our elicitation study findings (Chapter 4 and Chapter 5).

4 Formative Prototyping Sessions: Paper Prototypes

Equipped with the elicited knowledge from our design workshops, which addressed RQ1 and RQ2, we turned to address our third research question, specifically, what layout and interaction design can satisfy our requirement of enabling non-experts to author their desired backchanneling behaviors? (RQ3a). In this phase, we investigated layouts and interaction designs, refined our interface through iterative user testing with paper prototypes, and synthesize a novel interface design for authoring robot backchanneling behavior. We began by informally consulting with other HRI/HCI researchers in an interface brainstorm session, guided by our findings from the design workshops. We then designed an initial interface for authoring robot dialog with backchanneling, drawing directly from the results of our elicitation study and our informal brainstorm, and constructed paper prototypes for rapid evaluation and iteration. To answer whether our interface enabled non-experts to engage this creative task (RQ3a), and to initially investigate how they may engage such an interface and its features (RQ3b) and identify potential issues early on (RQ3c), we conducted formative prototype evaluation sessions with non-expert users, where we instructed them to complete simple tasks, observed and analyzed their interactions with the tool, and rapidly iterated on our design for the next session.

4.1 Designing an Initial Prototype Interface

Based on the high learnability and ease of use of block-based programming interfaces for non-expert users [9, 88, 101, 102], and their widespread popularity among existing robot authoring tools, we elected to use a block interface as the base of our prototype. Our interface thus leverages the known existing design pattern of block authoring and extends this design to accommodate authoring backchanneling for robots. Adopting this block pattern for our design

solves several problems, namely, how to present basic functionality like a robot speech action (typically a block with a text field for defining what the robot will say), and presenting robot actions in a visual, ordered manner. Our remaining challenge, then, was to identify how to present backchanneling features to users in a way that enabled their participation and creative expressivity in authoring. To explore early prototypes of potential interfaces that expose backchanneling behavior, the author of this thesis met with four other human-robot interaction and human-computer interaction researchers (colleagues) for an informal workshop to brainstorm interface designs, including the general approach and specific layouts.

4.1.1 Informal Interface Brainstorm Workshop

The author met with 4 human-robot interaction researchers for a 1-hour, informal workshop to brainstorm potential interface designs for authoring backchanneling behavior. Our goal was to explore how we might present backchanneling behaviors within a block-based programming interface, including controls for adding and/or adjusting the backchanneling. We began the workshop by reviewing our results from the elicitation study to ground our brainstorming and discussion, including how participants readily understood backchanneling concepts and tended to consider backchanneling through a lens of personality or emotion. We then asked researchers to individually sketch as many interface ideas as they could come up with in a 25-minute brainstorm exercise (with a goal of 10 interface designs each). Finally, we closed the workshop with an extended discussion to share ideas and discuss potential design challenges. Following the workshop, the author analyzed the interface sketches informally, and then discussed with their supervisor to create the initial prototype design.

This workshop resulted in 27 sketches of interface designs for authoring robot backchanneling. One recurring theme was exposing backchanneling behaviors as tunable parameters, for instance, a slider to control how often the robot nods its head or makes a particular utterance (Figure 11), though one researcher mentioned potentially using an LLM to fine-tune parameters on behalf of the user (R4). This basic level of authoring—tuning parameters by hand—exposes simple controls for backchanneling at the fine-grained layer of each behavior.




Figure 11: Fine-tuning of individual backchanneling behaviors was a theme among sketches from our informal interface brainstorm workshop (R2, R4, R1).

Another major theme was personality-based authoring of backchanneling behaviors, featured in several sketches in different forms and capacities (Figure 12). All researchers agreed that personality may prove necessary in enabling interaction by non-experts, given the abstract, intangible nature of backchanneling as a conversational behavior. Some further suggested that




Figure 12: Interface sketches for authoring backchanneling emphasized the use of personality descriptors, such as "happy", "extroverted" and "affirming" (R2, R3, R4).

the user should be able to start by selecting one of many “preset” personalities, and then offer customization or creation of new ones (see Figure 12, left). Another sketch suggested that a robot could have “listening modes” (e.g., “affirming” and “skeptical”), acting as buckets to which the user adds specific behaviors. In general, our takeaway was to support users with preset starting points for authoring behavior with personality, and to empower them with some manner of creating and customizing new personalities.

4.1.2 Interface Design

Following the workshop, the author informally analyzed the interface sketches, then met with their advisor to discuss the generated ideas and recurring themes throughout the sketches and elicitation study results. Through this informal process we arrived at an initial design which the author implemented in the first paper prototype (Figure 13).

We decided to adopt the block-based programming language approach as the foundation of our work, given its field dominance and success at being accessible to non-experts [20, 88, 102]. Following common block-based programming patterns, we created tool palettes on the left for blocks representing dialog elements which can be dragged from the palette into the main window, where a copy is created, and users are able to customize these by clicking on them, for example, entering text for the robot to say. Elements on the main window click together if dragged within close proximity, dictating the serial order of the events in the behavior. Our design aimed to include backchanneling behaviors directly into the creation interface without hindering authors’ creation of the core text content: authors can select *styles* for listening



Figure 13: Paper prototype of our initial interface design, complete with core block features (toolbox, left, and workspace, middle) and all additional elements including tooltips, keyboard and selection hints, dropdown menus, navigation and callouts.

elements to encapsulate backchanneling. Further, we aimed to incorporate both coarse-grained, personality-level definition of robot behaviors as well as the ability for fine-tuned editing and control, and adopted a coarse-to-fine workflow: users first specify a high-level style to encapsulate template backchanneling behavior, including *happy*, *sad*, *excited* or *bored*, but can refine the behavior specifics through fine-grained modifications to these styles using the panel on the right. Here, users can modify parameters such as the frequency, intensity and direction of nodding behavior, utterance content, frequency and volume, and frequency of eye contact and

gaze shifts. This menu also supports turning off each of these types of backchanneling behaviors within a style, and new styles can be created as desired.

We constructed the paper prototype using cardstock with hand-written text, and cut-out to shape. We photocopied base elements to facilitate rapid exploration and to quickly make new elements. We made the tabletop workspace using high-tack paper sheets (we used photo album inserts) to prevent elements from slipping around and to lock them to their positions. Drop-down menus, call outs, etc., were all prepared on paper. We prepared analogs for computer usage where needed: when the user of the paper prototype checks a box or drags a slider, overlay elements react to reflect the new state. When the user wants to type, they tap their finger on a text box and a “keyboard” element appears to indicate the action the user has invoked; we then provide the user with a pencil and a blank sticky note to write on, attached to the text box. Our paper prototype interfaces are fully featured for accurate use and exploration by users.

4.1.3 Pre-study Interface Iteration

Before our formative paper prototype evaluation with non-experts, we identified during informal researcher pre-testing that the initial way we presented styles may be difficult or tedious to engage efficiently. In the initial prototype (Figure 13), the user’s first interaction with styles is likely to be the listening style dropdown provided in the Robot Dialog block, which encapsulates other dialog behaviors added by the user. This may be unexpected as the user is likely to want to author what the robot *says* as a first step, and there is no indication of where this style is applied or what it may do. Further, the presence of an “auto” default may be unclear (will the robot automatically listen?), and *editing* styles occurs in a separate side panel (right) with no intuitive connection to the behaviors authored in the workspace. The concepts are disconnected,


for instance, does modifying the *happy* style in the side panel affect all block instances with *happy* selected in their dropdowns, or just one, or only future instances? We came to understand that this lack of clarity represented a usability issue that made styles difficult to interpret and manage while exploring and authoring.

To make our interface design more intuitive, we decided we needed to centralize all the pieces of our conceptual model for robot dialog: put styles in the toolbox with other dialog elements. Styles became a type of block in themselves, a decorator complementing other dialog elements, by way of attaching to them (Figure 14). This unified approach remedies the fragmented nature of styles in the previous version and integrates our representation of backchanneling behaviors into the core interaction loop (dragging blocks from the toolbox to the workspace), making *styles* idiomatic to, and benefitting from, the core block-based programming paradigm. We maintain support for presets, user-created styles, and editing parameters within a style.


The other notable change is that we redefine styles from “listening styles” to simply *styles*, allowing the user to place a style on any type of block, following our design recommendation for enabling *Richer Dialog: Beyond Backchanneling*, and expecting that users may not draw a clean boundary at backchanneling-specific behaviors when operating through the lens of *styles*. For instance, a style on a Listen block defines backchanneling, on a Say block it defines speech tone or style, and on a Robot Dialog block it defines the style to apply for all nested blocks.

This version of our prototype constituted our starting point for the formative prototype evaluation study with non-experts, detailed in the following section.

1. interface showing a blank state with dialog toolbox open



2. populated workspace with the styles toolbox and preset styles shown



3. create a new style by “dragging” a preset into the Your Styles space in the toolbox.



4. edit the new style's parameters with the style options pane




Figure 14: Improved interface design featuring styles as draggable modifier blocks that are attached to core dialog elements (listening, speaking).

4.2 Formative Prototype Evaluations with Non-experts

Given our target user base and goal to rapidly refine our candidate design, we conducted a series of three formative prototype evaluations, rather than a formal usability study. We designed the evaluations to assess user learnability and alignment between our interface structure and users' conceptual understanding of dialog and backchanneling, as well as whether participants could use the interface to complete simple dialog and backchanneling authoring tasks. We developed

a set of simple tasks relating to participants authoring human-robot dialogs, ranging from simple robot speech, simple listening personality, to complex dialogs with advanced features such as customizing individual backchanneling instances. For each, we specified clearly what we want participants to accomplish but did not explain how to accomplish it using the prototype. Between the three sessions we refined our prototype, improving it for the next one.

4.2.1 Tasks and Procedure

Tasks—For the first task we asked users to complete a pre-designed and highly prescribed dialog, starting from scratch, that included the robot speaking and listening, and the participant both “playing” the dialog (test it on an imaginary robot), and modifying their creation. This was explained to participants verbally in several steps (e.g., make the robot say “hi, how was your day at work?” and make it listen to the user’s response) (see Appendix B.3). The second task was an open creation phase where participants were asked to create a new dialog. We asked participants to imagine a make-believe scenario to help them get started (a robot acting as a cooking assistant) but left all dialog creation and backchanneling up to them and provided no additional instruction.

During sessions the author manually operated the prototype, updating the interface in response to participants’ indicated interactions (e.g., clicking or dragging with their finger). The author had numerous templates and blank widgets prepared, as well as various pens, pencils, and highlighters, to enable on-the-fly interface extension or creation based on participant actions. Participants were also encouraged to use these materials to explain what they are trying to accomplish, though only one participant engaged in this way (P3).

Procedure—We invited individuals to attend a 1-hour workshop on our institution’s campus. Each participant was compensated \$20 CAD. We conducted workshops with groups of people, although only one person interacted with the paper prototype at a time given the requirement for the research operator. We began the session with consent forms (Appendix B.1), followed by refreshments and an introduction by the author who moderated each session. We gave our explanation of backchanneling from the elicitation study (Appendix B.2), and a video tutorial of the interface explaining each of its features (tool palette, blocks and their functionality, workspace and play button) and how to interact with a paper prototype (picking up and arranging blocks, tapping with a finger to “click” elements, writing text with a pencil). Following, the moderator took one participant at a time to a separate room to complete the tasks on the paper prototype (Appendix B.3), with an assistant moderator took observational notes. The remaining participants completed basic dialog authoring activities from the elicitation study, provided with the same stationery and sample dialogs (Appendix A.6). This research was approved by our institution’s research ethics board.

To inform prototype iteration, the author and assisting researcher noted for participant ability to complete representative tasks, paying particular attention to how participants explored the interface and how they engaged with their authored creations during the open-ended task. The author also watched for participant ability to achieve their desired outcomes with the interface tools provided and without complex workarounds or appropriation. All analysis of this observational data was completed by the author, who subsequently iterated the interface design in regular consultation with their advisor.

4.3 Results: Prototype Iteration

We had 7 participants across the three sessions, recruited from posters placed on our institution's campus, 4 men and 3 women, six of whom were aged 18-29 years and one aged 30-39 years. One participant indicated coding experience, and another had previously taken a robotics course in secondary school.

Overall, participants readily understood our interface design and were able to, with minimal explanation assistance, author human-robot dialogs with backchanneling behavior, giving us confidence in our overall interface design. The largest trouble spot was management and use of the dialog styles: we initially provided complex style templating, customization, and management features. However, we observed participants struggling and getting confused regarding managing styles, which appeared to inhibit dialog authoring and increase overheads. Therefore, we greatly simplified the style system by completely removing management of presets (e.g., adding, deleting, naming, etc.), and instead simply enable users to customize individual style instances in place. In the following sections, we detail our prototype iteration after each session, arriving at a final prototype which we later implement as a high-fidelity software prototype in Chapter 5.

4.3.1 *First Workshop: Style Discoverability and Customization*

From the first session, we observed that a major challenge for users was interacting with our styles features as currently presented. First, discoverability—people could not find the styles tab of the toolbox, even when guided towards the tab bar, and the interactivity of the tab bar (that it might reveal another page of the toolbox) was not intuitive, at least on paper. This was at least partially a procedural oversight; our tutorial avoided priming participants on how to *use* styles,

and in doing so we omitted showing participants that styles *existed*. This led to other downstream problems in the authoring tasks; when asked to make the robot happy, one participant added a listen for keyword block and typed in the keyword “happy”, then changed it to “promotion” to fit their dialog context, since they expected the robot would be happy if the user said they had received a




Figure 15: Users can now create a new style from a blank state, or use an existing style as a starting point, as desired. Options in the blank state appear unchecked, and reveal default values when enabled (e.g., middle slider values).

promotion at work. For this session, then, we eventually showed users how to interact with the toolbox tabs if they could not find it through exploration (nobody did), revealed the existence of the style blocks toolbox for them, and for the following session we iterated on our procedure to more explicitly note that there are more available blocks in the toolbox through the tab bar. Once users were aware of styles, they successfully applied them to customize their robot dialogs in the workspace.

Discoverability was not the only challenge that styles presented for users. One major problem we unearthed was style *management*: the workflow for creating and editing styles was unintuitive for users. To create a new style, the user was required to drag a preset style into the Your Styles area to use as a starting point (regardless if they clicked the Plus (+) button). This was unexpected and unintuitive behavior for our users. For instance, P1 attempted to apply the happy style to their dialog by dragging happy into the Your Styles section, thinking it would apply to the robot, but were instead met with the style customization pane, at which point they

hesitated and attempted to undo their action. Thus, we modified the style creation process to be more intuitive, supporting users to start from a blank slate style (with no selected behaviors, and parameters set to defaults when a behavior is checked, for instance, sliders set to the middle value). Users can still drag a preset to start from if they desire, indicated by the labelled drag and drop area (Figure 15).

Users also struggled with the flow of editing a preset style. If the user wanted to make the happy style include more frequent nodding, they would click on the edit button on the happy style block, which would initialize the process of creating a *new* style, since we wanted to preserve the preset styles, yet allow the user to create modified copies as desired. This proved to be unnecessarily obtuse and unintuitive, and so we updated the preset styles to explain this process more clearly with a modal (Figure 16).




Figure 16: We clarified how attempting to edit a preset style (top) preserves the style, and users are invited to create a new style with the preset as a starting point (bottom).

4.3.2 Second Workshop: More Style Management Issues

In the second workshop, we discovered that styles were still difficult to find, especially the tab bar in the toolbox. Given that siloed toolboxes are idiomatic to block-based programming interfaces, and we are not attempting to evaluate how well Blockly translates to paper, we




Figure 17: Our final paper interface prototype, deployed in the final workshop.

gathered that this challenge may be a limitation of paper interfaces and opted to maintain our Blockly-like structure, but made the tabs more obvious for a paper setting by using separated manila folder tabs. Further, we renamed the “Blocks” tab to “Actions” to better represent its contents. Further, we observed that the new modal popup for editing preset styles (Figure 16) was still unintuitive, and thus converted this to a simpler, informative tooltip, along with automatically creating a new custom style for them, named after the preset, for instance, “copy of Happy”. This helped especially with the case where users would attempt to edit preset styles already placed in the workspace, where the modal’s verbiage to “drag above” did not make sense for the workspace context. These changes resulted in our final paper prototype, shown in Figure 17.

4.3.3 *Final Workshop: Simplification*

After the third and final workshop, we observed that we were reaching saturation with the insights from the paper prototyping process: for many of our notes on user challenges, it had become difficult to discern the degree to which paper interfacing was contributing to the problem over actual interface design shortcomings. Given this point of diminishing return, we decided to conclude the iterative user testing paper prototyping stage.

Before developing software, one major issue was still clear across all our prototyping: the style management workflows were cumbersome, unintuitive, and tended to hinder authoring. Rather than forcing the user to manage their own styles, or continually tweaking this feature, we opted for a major simplification: we removed the ability to create additional styles as new blocks. Instead, we provide the preset styles as the only available blocks and enable editing of these styles within the workspace only. If the user changes the happy style, it applies only to that copy of the block. Users can still reuse their customized happy style by leveraging Blockly's copy and duplicate features. This way, users do not corral styles outside of the workspace (i.e., management), and can thus focus on applying and tweaking styles as they see fit (i.e., authoring).

Another confusion point was customization of styles for speaking behaviors (see Figure 5)—while we enabled styles to be applied to listening (as backchanneling) or speaking behaviors to follow our *Richer Dialog: Beyond Backchanneling* design recommendation, the customization screen we provided for backchanneling did not match the speaking behaviors (e.g., verbal utterances while speaking would not make sense). We removed customizability of speaking styles to enable us to focus on backchanneling, leaving this for future work.

4.4 Chapter Summary

In this chapter, we engaged our understanding of non-experts' mental models of robot backchanneling, as investigated in the elicitation study of Chapter 3. Through iterative prototyping with this foundation of knowledge and input from both HRI researchers and non-experts, we addressed our research question of identifying an interaction design and layout that enables non-experts to successfully author backchanneling for robots (RQ3a), with extra attention to detail on how we use emotion or personality to present (and thus, how we allow users to modify) backchanneling. Through formative user testing, we additionally engaged how users might engage such an interface (RQ3b) and identified challenges encountered in early prototype interaction designs (RQ3c). We take our findings from this chapter forward into Chapter 5, where we evaluate a high-fidelity software prototype with non-experts, whose design is heavily and directly informed by our prototyping results.

5 High-Fidelity Content Creation:

A Tool for Designing Robot Backchanneling

Following our iterative prototyping stage, as described in Chapter 4, we have the following pieces in place: first, we possess an understanding of how non-experts engage backchanneling conceptually, which aspects they struggle to understand, and how they might consider backchanneling for robots, including designing robot backchanneling with paper crafts (RQ1, RQ2). Second, we created a prototype interface design grounded in these mental models, in expert HRI researcher ideas, and in formative user testing with paper versions. Equipped with this knowledge, we proceed to address our third research question, specifically, how our interaction design and layout supports users to author robot backchanneling creatively, which behaviors they create, and any potential unforeseen challenges or limitations of our interface design (RQ3). In this chapter, we detail our high-fidelity software prototype and a user evaluation with non-experts, and the limitations of our approach, completing this thread of identifying how we can support users to create nuanced robot dialog behaviors. We begin by detailing our interface and its features (5.1), followed by our evaluation workshop where we observed non-experts using our tool to author robot backchanneling (5.2), and then our results (5.3), concluding with a discussion, including the contributions and limitations of our evaluation (5.4).

5.1 Interface Design

We developed a simple interface that enables users to create simple human-robot dialogs including backchanneling (Figure 18); this is a limited design that matches our goal of testing our integration of backchanneling into dialog, but in practice our interface would be integrated

into a more complete system. Users can drag and drop elements from the tool palette on the left onto the workspace, where they can snap together in specific ways to form a full dialog design.

The user can build linear dialogs by adding a Robot Dialog block (Figure 19, blue block on the left), and then adding and nesting within it a series of action blocks (the next three) arranged sequentially top to bottom, as shown in (Figure 18). To remove elements, they simply drag them out onto the workspace, where they are not used, or back into the palette, and they disappear.

We provided the following action blocks: “Say” blocks (green, Figure 19) which users can customize by clicking on “type here” and entering text for the robot to say; “Listen while person speaks” which indicates that the robot should wait until the human is done talking (until silence is detected); “Listen for Keywords” blocks which listen until user-specified keywords are recognized, which users can specify by clicking on the block and typing in the keywords. Users can press the play button (upper right corner, Figure 18) to send their behavior to the robot and test it out. For playback we used a VStone Sota small humanoid robot on the desktop next to the user’s computer (Figure 3).

Finally, inspired by the personalities created by participants in earlier workshops, we provide four different “Style” blocks that users can drag and attach to their action blocks: *happy*, *sad*, *excited*, and *bored* (Figure 19) to specify *how* the robot should complete that action. When attached to listen blocks, this specifies the robot’s backchanneling behavior. When attached to say blocks, the styles currently have no effect given the goal of our prototype to author backchanneling, thus we leave this feature for future work.

We programmed the robot to use nodding, utterances, and gaze behaviors for backchanneling, and we designed the initial styles to map to specific configurations of these. For instance, excited




Figure 18: Our interface, based on Blockly, with available blocks in the tool palette (left) and user creations in the workspace (right). The "Play" button on the upper right lets the user test their designs on a connected, physical robot.

would make the robot nod its head with high frequency and intensity, and emit utterances “wow”, “yes” and “right”. We did not evaluate or validate the specific backchanneling behaviors themselves but rather designed them to roughly match the personalities, given the research focus on the user interface and not the quality of the resulting robot dialogs themselves.

Users can fine-tune or edit a style by clicking the edit icon on the style block, whereby a panel pops out (Figure 21). This panel provides a full range of fine-tuned controls of the backchanneling behavior, including enabling the user to control which backchanneling behaviors and modalities are enabled (nodding, utterances, and gaze behaviors) and fine-tune parameters for these behaviors. For instance, the user can adjust the frequency and intensity of nodding behaviors




Figure 19: A diagram showing the blocks available to the user for authoring human-robot dialog with backchanneling.

and the direction of these head movements (vertical or horizontal), adjust the frequency of utterances and select specific words to use from a list, and adjust how much the robot maintains eye contact and how frequently it shifts its gaze.

Overall, this interface design falls in line with our initial recommendations from the elicitation workshop. Presenting the behaviors as coarse-grained styles first, followed by enabling optional fine-tuned control of specific parameters.

Further, by specifying the style per action we enable users to author backchanneling contextually, customizing each segment of dialog with the desired behavior. Finally, although not implemented at this time the interface itself does map the styles to utterances as well, as our approach

followed our recommendation to provide *Rich Dialog Beyond Backchanneling* and this made sense for the style metaphor.

For our implementation we built a user-facing web client, a server to coordinate user clients and robots and send commands between the two, and a thin-client running on our robots. We used VStone Sota robots as simple, small social robots designed for dialog interaction with people. The web client interface uses Google Blockly (v11.2.2) and React (v19.2.3) for interface elements. The server uses Python3 and listens for connections from clients and robots, generates text to speech content using the Google Translate API and performs keyword recognition using Alpha




Figure 20: Users can optionally edit backchanneling styles to fine-tune parameters for individual behaviors. Changes are saved only to the modified block; other blocks of the same style remain unaffected.

Cephei Vosk, which runs locally. The thin client on the Sota robot, written in Java 8, simply interprets and executes commands received from the server and manages socket connections for sending microphone data and receiving and playing back audio.

5.2 Prototype Evaluation Workshop

We conducted a prototype evaluation workshop to assess how our interface and its conceptual model for backchanneling supports non-expert users to explore and author human-robot dialog that contains backchanneling behaviors. To reflect on the alignment of our interface with our intended conceptual model for presenting backchanneling behaviors to non-expert authors, we recruited participants from previous phases (Table 1), evaluating for the general effectiveness of our interface in enabling users to create rich robot dialog with customized backchanneling.

We invited the same participants from our elicitation workshop and formative prototype tests to attend a 2-hour workshop in a computer lab at our institution, where they each had their own computer and robot workstation. After consent forms (Appendix C.1), we tasked participants with creating robot dialog behaviors for a VStone Sota robot using our interface (Figure 3). After a pilot session for technical testing and debugging, we held two sessions—we had enough workstations to host a maximum of 7 participants per session. Participants were compensated \$35 CAD at the start of the workshop. This research was approved by our institution’s research ethics board.

5.2.1 Tasks and Procedure

Each participant sat at a computer station with a mouse and keyboard, the interface loaded in a web browser, and a VStone Sota robot placed on the table. The author first explained backchanneling (5 minutes, see Appendix C.2), supported by two short video demonstrations of

human-human dialog and human-robot dialog. We guided participants through a brief interactive tutorial to create a simple dialog, and then to add backchanneling, covering all interface features: first make the robot say “Hi, how was your day at work?”, then listening attentively in a *happy* style; we included this tutorial given that the goal of our workshop was to test how participants engaged with and understood the human-robot dialog creation with backchanneling, and how they are represented conceptually in the interface, rather than the learnability of the Blockly interface itself [45, 88, 101, 102]. Following the tutorial, we began an extended, open-ended creation task (60+ minutes) where participants designed robot dialog and backchanneling behaviors. We emphasized to participants that they were not being evaluated for performance at any point during the workshop, but rather, they were helping us to test our system. The primary and secondary author circulated to take notes, inquire about participants’ designs, answer questions and resolve technical difficulties.

5.2.2 *Analysis Plan*

We collected qualitative data including researcher observations and notes in addition to digital artifacts of robot dialogs authored by participants. The author and their advisor conducted a qualitative thematic analysis on these data to assess user understanding and engagement in designing robot backchanneling using our interface, and the range of dialogs and behaviors they authored using the tool. Our primary analysis centered around how participants engaged with and understood backchanneling while designing robot dialog, and how our interface design’s incorporation of backchanneling supported user creation.

The author first organized all researcher observations and notes pertaining to how participants understood and engaged with backchanneling while designing with the interface, the range of

created behaviors, issues or uncertainty they may have encountered with respect to the interface's presentation of backchanneling, and feedback received. The author then analyzed artifacts of participants' creations to assess how our approach enabled users to author custom robot dialog with personalized backchanneling. Below we present key themes synthesized collaboratively by the author and their advisor.

5.3 Results

Recruitment resulted in 13 individuals who self-identified as being interested in designing robot behaviors (8 women, 5 men) to participate in 2-hour workshops, 6 in one and 7 in another. Most participants were aged 18-29 (9 of 13), with one aged 30-39, two aged 40-49, and one aged 60-69. Our participants mostly had little robot programming or design experience; three participants had programmed robots before (two in university courses, one with hobby kits/toys), and two others had general programming experience (one software developer and one student), but no robotics expertise. Nine individuals had previously participated in the elicitation study and two in paper prototype sessions (Table 1).

Our thematic analysis resulted in the following dominant themes:

5.3.1 *Low Friction and Sustained Engagement*

We found that all participants engaged with the task, with all actively authoring robot dialogs for the entire 60-minute open-ended creation activity and creating a wide range of behaviors. Participants did not encounter significant problems or barriers while using the interface. All participants successfully created dialog with backchanneling (e.g., see Figure 22) using each aspect of the interface without requiring researcher assistance, and all participants included backchanneling behaviors in their dialog designs during the open-ended task, all explored

varying configurations and customizations, and all continued to use the backchanneling in their creations throughout the workshop.

Participant problems were predominantly technical errors with our system, for example, occasionally needing to reset a robot, a bug where the side-to-side backchanneling would not work, or a bug where the system allowed for the “robot dialog” block to be nested within others, resulting in an error message. We provided small clarifications on these and other small implementation issues, including if the robot would enunciate question marks (it did not).

Related to this, participants asked general clarifying questions about features and functionality, such as whether the robot was listening to them (P25) and whether the robot would ever finish listening for keywords if they did not say a keyword (P14). A few participants asked mechanical

questions about the backchanneling elements, for example, if they could rename a style upon customizing it (P30), but no participant asked for assistance on the backchanneling elements.

Users explored the space of interaction possibilities with our interface, investigating and testing the backchanneling options and system capabilities. For instance, P11 wrote quick tests to see how different backchanneling parameters affected the robot’s behaviors, making the robot say “Nodding. I’m nodding. Maximum frequency and intensity nodding” with an attached *happy* style customized with frequency and intensity set to 100 percent. They discovered a bug that the



Figure 21: Users authored a range of interactions for their robots and applied backchanneling to them. Pictured is a snapshot of a portion of P30’s workspace and their joke-telling robot.

robot would not nod its head side to side when the style option was selected, testing with robot speech: “I am trying to turn my head side to side. The will to do so is not there.” (P11).




Figure 22: P6 used backchanneling styles contextually while authoring dialog, including an *excited* style that repeatedly says "wow".

5.3.2 From High-Level Structure to Tuning

From our in-session observations and analysis of saved workspaces, we observed a common pattern: our users would build their dialog first with robot speech and keywords, perhaps run an initial playback, then add styles and run it again, and tweak styles, running repeatedly as they went. For instance, P30 iterated on a dialog design to come up with a new style they called “intense excitement”, a frenetic listening style with nodding frequency set to the maximum and eye contact set to ten percent. They began with a *bored* style, tested it on the robot, and then heavily customizing the style to drastically change the robot’s behavior. P1 and P11 explored the style customization menu, testing behaviors out with incremental changes. P16 typically used the excited preset but then modified most instances to expand the range of utterances, selecting different combinations for different parts of the dialog to create contextual utterance behaviors (Figure 23). P29 and P36 also frequently edited the utterance list for their styles, for instance, P29

enabled utterances for the *bored* style (this style has utterances turned off by default) and fittingly added just one utterance to the list, “sigh”.

5.3.3 *Exploration of Rich and Nuanced Behaviors*

Participants spent considerable time and effort to author rich and nuanced behaviors throughout the workshop, exploring possibilities and fine-tuning their creations. Each participant made extensive use of the predefined styles, attaching them to robot speech and listening behaviors to customize their dialogs at a basic level.

Participants designed backchanneling behaviors to match the specific desired interaction contexts and narrative, for example, P6 authored a dramatic robot that urgently requested the user’s assistance, “there isn’t time, I don’t have any hands, will you help me?” using the *excited* style while listening to match. P30 created a dialog where the robot was sad and asked the user for permission to sing a song to be *happy*: if the user said yes, the robot sang and used a modified “shy” style (annotated with the blockly comment feature) while listening for feedback on their performance.

It was also common for participants to vary their robot’s use of backchanneling throughout a dialog instance, as part of their robot’s character design. For example, a “Joke bot,” which the participant annotated as “trying to make a sassy robot”, using the *happy*, *bored*, and *excited* styles to shape the narrative while telling a joke (P30, Figure 22), and a compassionate listening robot that alternated between *excited* and *happy* listening styles (P25). Some even built this into user response (although our system did not support branching, for example, P7 authored a robot that wanted to take over the world, and when the user replied, “no way”, the robot backchanneled

with a custom *sad* style: shaking its head with maximum frequency while emitting the utterance “sigh”. P35 made a robot say that it was sad, followed by listening with a *sad* style.

Fine-tuning styles—Participants also fine-tuned styles to match roles and context. For instance, P14 created a healthcare-focused companion robot and fine-tuned its use of the *happy* and *excited* styles, such as when the robot greeted the user and introduced itself as a “personal healthcare companion”, they increased the nodding frequency and intensity from 60% and 70% to 95% and 99%, respectively. P6 made their robot tell a joke—“want to hear my Owen Wilson impression?”—followed by a listening action with a customized style so the robot only said the word “wow” with frequent repetition (Figure 23).

These results support two key findings: first, that presenting backchanneling behaviors at a high-level as “styles” that reflect an overall emotion or disposition was approachable; second, that users heavily engaged in customizing styles for their robots throughout the workshop, highlighting the importance of user customization of robot backchanneling.

Authoring Backchanneling Minimally—There were participants who used backchanneling minimally, for example, keeping one style throughout a design, such as a robot playing a word game while using the same *happy* style (P30), or making casual conversation (P35, P19). One participant achieved a similar effect by simply attaching the *happy* style to the parent “robot dialog” block to control active listening throughout (although this participant, P1, changed the speaking styles). Finally, P29 created a companion without any backchanneling, placing styles only on robot speech.

We also noted some appropriation and unexpected use of our interface, for example, by using emoticons to alter the robot’s speech (the text-to-speech system pronounced them), and accented

characters to fine-tune the robot's pronunciation of words (P7). Two participants seated adjacently designed dialogs that made their robots talk to each other and listen for each other's keywords (P19, P34).

Desire for additional backchanneling customization—Participants identified and shared myriad additional features, not already included in the system, that they felt would aid their design process or provide more options and control during design of backchanneling. One participant wanted additional preset styles to use in certain contexts, e.g., disagreement (P14). Another felt the timing options for utterance backchannels could be improved (P11) and thus wanted more control over the delay before the robot gives its first utterance, as they felt the robot backchanneled too quickly after finishing speech and letting the user begin to speak, which was not desirable. Further, many participants mentioned hand movements as a key part of backchanneling that they wanted to incorporate in their design (P1, P7, P11, P14, P30). P11 wanted to add hand movements for expressivity, citing engagement, compassion and more “meaningful” interaction. P14 desired to customize the hand movements to match the cultural context, or more generally to specify cultural context so the robot's styles accurately represented backchanneling in conversation. For backchanneling customization, P30 asked if they could rename their style block so that it could accurately describe their customized style. This functionality was not supported by our prototype, so they used the comments feature of Blockly to annotate the style block with “not bored!!!” (P30).

Desire for other authoring tool features—A few participants (P1, P7, P14 and P30) separately mentioned wanting branching dialog based on what the user says. Other desired features included adjustable delay when checking for keywords before giving up and continuing; support

for key phrases in addition to keywords; and in-line placeholder values for speech, giving the example “the current time is {now}” (P7). Overall, users desired more controls and features for authoring backchanneling and other elements of rich dialog.

5.4 Discussion

We found non-expert users to understand backchanneling while designing robot dialog and engage thoroughly in authoring backchanneling behaviors throughout the workshop. Participants made extensive use of styles, demonstrating ready engagement with the abstraction of backchanneling into high level personality design synthesized from the elicitation workshop findings. Non-experts’ natural engagement with and creative authoring of backchannel behaviors demonstrated that the conceptual model underlying the design (coarse styles with optional fine-tuning) aligned with their expectations and supported meaningful use. No participant encountered difficulty in applying styles to robot actions, understanding which parts of dialog they would affect, or using the style customization menu to specify backchanneling behaviors. Our findings illustrate how our interface design and conceptual model, derived from our elicitation study, successfully presented customization of backchanneling to users in an approachable, understandable manner, and non-experts engaged in customizing their robot backchanneling designs using this model. This finding illustrates and supports our design recommendation from the elicitation study to *Design Dialog Authoring Tools to Include Backchanneling* (section 3.3.1).

Users engaged heavily in authoring at the high-level using styles to shape robot behavior. This tendency aligns with other work which demonstrated that, for robot behavior authoring, “moderate customization without excessive granularity” may lead to better engagement in

authoring tasks [106]. All our users began authoring at a high level, with some remaining at this level and others exploring and fine-tuning behaviors. This tendency to engage in a top-down manner, creating dialog first with general personality styles, followed by fine-grained adjustments to individual backchanneling behaviors, aligns with our findings from the elicitation study where participants engaged in authoring robot personalities before crafting and tweaking dialog and specific backchanneling behavior. Our participants' engagement with customizing fine-grained backchanneling parameters shows that non-expert users are able and willing to fine-tune backchanneling to achieve a desired result. This finding supports our design recommendation on granularity to *Support Personality-Level Control* and *Behavior-Level Fine-Grained Control* of backchanneling.

Our results show how users leveraged our interface and conceptual model to author personalized and context-specific backchanneling for robots, for instance, employing different styles when authoring humorous, dramatic or sincere dialog designs, or requesting contextual control over culture-specific backchanneling behaviors, which aligns with backchanneling literature [15]. This result speaks to the importance of flexibility and power in end-user authoring interfaces—enabling users to author tailored backchanneling behaviors for robot dialog resulted in strong engagement and a range of highly specific, creative designs that would have been quite different without the nuanced backchanneling style. This matches what we learned in our elicitation study and is further reflected in the literature [17, 38, 53]. This finding supports our design recommendation to *Support Contextual Customizability* for authoring robot backchanneling.

One result we did not anticipate was the degree to which our non-expert users wanted to exercise *even more control* over nuances of robot backchanneling, for instance, including controlling

specific timing intervals, authoring hand gestures, fine-tuning robot pronunciation, custom names for styles, and dialog branching. This desire for control underscores our finding from the elicitation study, that users can—and may want to—exercise a high degree of control over their robot backchanneling designs. This finding further supports our elicitation study result that non-expert users of robot behavior authoring tools may desire to provide exact, specific behaviors where needed, which may be difficult to achieve for automatic or predefined behavior generation. Thus, these results overall lend support for coarse-to-fine authoring of robot backchanneling. Robot behavior authoring tools should consider supporting not only backchanneling, but fine-grained control over robot dialog behaviors.

There are several ways in which our interface could be improved. Participants in our second session expressed desire for a cancel button to stop execution, in cases where they had a long behavior and wanted to quickly adjust a parameter and play it again without needing to wait. This feature would have enabled a faster feedback loop for the rapid and iterative prototyping our participants engaged in throughout the workshop. Further, our focus on supporting user authoring of backchanneling limited our users to author simple dialogs without branching—a feature several users requested. While we leave to future work integration of our model into more robust robot dialog authoring tools, our model could be effective in such systems to support users to author backchanneling in branching dialog, and authoring contextual expressive speech with styles.

5.4.1 Limitations

A key limitation with our work, specifically in the final prototype evaluation study, was how we explained backchanneling to participants as part of our methodology. While this step was

important in earlier studies to help ground and direct users, we realized after the fact that this may have primed our participants too strongly toward focusing on backchanneling behaviors. It would have been better if we excluded this education, and rather just demonstrated the feature itself, to gain a better sense of how naturally participants may gravitate toward these features. However, despite this our findings on participant ability and use of the tools are still significant.

While our findings from the initial elicitation workshops indicated that participants tend to think broadly about backchanneling, lending to coarse-grained approaches, in our interface if the user wants to create a new preset (e.g., annoyed), they *must* interact with the fine-grained behavior controls to design this behavior. Currently, this prototype interface does not support coarse-grained customization of “styles”. Future work could investigate high-level customization of these styles, so the user can avoid interacting with fine-grained behaviors throughout the authoring process. For instance, allow the user to describe a completely new “style” using natural language input (evaluated by LLMs and translated to the underlying parameters), or provide a 2-dimensional graph representing, for example, valence and arousal [84] on which the user selects and drags a point to their preferred location, to approximate the desired behavior.

Further, our prototype interface lacked features that many participants expressed desire for, namely branching behaviors. Given our research goal to enable users to author backchanneling, we opted to forego branching in the interest of time to focus on providing backchanneling customization instead, but more robust behavior authoring tools featuring branching may still benefit from our interface design.

5.5 Chapter Summary

In this chapter, we addressed our third research question regarding how to incorporate non-experts' mental models of backchanneling into a novel interface design that they can use to design robot backchanneling. We demonstrated that our layout and interaction design, formulated through elicitation (Chapter 3) and prototyping stages (Chapter 4), successfully enabled non-experts to engage in authoring robot backchanneling creatively (RQ3a). We found that users engaged with our high-level, personality-based authoring paradigm, while also digging deeper to finely tune specific behaviors, creating a variety of backchanneling implementations throughout our evaluation workshops (RQ3b). We also noted the limitations of our design, including barriers faced by our participants and potential alleviations in future, more robust tools (RQ3c). These insights can inform future robot behavior authoring tools that seek to enable end users to design rich robot dialog, including backchanneling behaviors. In the next chapter, we consider holistically the contributions of this work in all its stages, its limitations, and considerations for future work.

6 Conclusion

We designed and presented a novel interface for user authoring of robot backchanneling behaviors in dialog, informed by insights from an elicitation study to learn about how non-experts conceptualize backchanneling, and iterative prototyping through a user-centered design process. Our findings demonstrate how users engage with authoring robot backchanneling creatively for robot dialog and use a tool to create a range of dialog designs. Further, we demonstrate how users leverage both coarse-grained abstractions and fine-grained controls while authoring rich robot dialog with backchanneling. Robot behavior authoring tools—particularly those targeted to non-experts—can leverage our interface design, based on user understanding of backchanneling, to support user authoring of these behaviors.

In this final chapter, we return to our initial research questions (Chapter 1.2) to contextualize our contributions with respect to our research goal. We outline recommendations for robot behavior authoring tools, based on our findings, to inform future research and development of authoring tools that support end-users in creatively designing robot backchanneling or other behaviors. We discuss the limitations of our approach and findings and suggest avenues for future work to continue in this thread of end-user authoring of robot backchanneling. Following, we close this dissertation with a concluding summary.

6.1 Contributions

We investigated how to enable non-expert end users of robots to participate in designing robot backchanneling. This included engaging non-experts in the topic of backchanneling to understand how they conceptualize these behaviors, how they consider backchanneling in the

context of robots, and how they approach describing and authoring backchanneling for robots. It also included iterative interaction and interface design to identify how we could present backchanneling to users such that our authoring tool is compatible with their mental models, and an evaluation of our interface design through a workshop study where non-experts designed backchanneling hands-on. In this section, we outline our contributions with respect to our guiding research questions.

RQ1: How readily and comfortably do non-experts engage with backchanneling conceptually?

Through our elicitation workshop study (Chapter 3), we uncovered that non-experts readily understand backchanneling concepts, the impacts and importance of backchanneling for human conversations, and how backchanneling could apply to human-robot dialog. Non-experts readily understand a range of behaviors and modalities, and consider backchanneling to be an important component of both dialog and individual personality. When asked to brainstorm backchanneling behaviors, non-experts may consider behavior beyond the strict and specific boundaries of what is considered ‘backchanneling’, taking a more holistic view of dialog behavior.

RQ2: How might non-experts describe backchanneling elements of robot dialog?

When describing robot backchanneling in our elicitation workshops, non-experts frequently employ human characteristics to illustrate behavior. Most commonly, we observed these descriptions to rely on components of human personality, such as a “shy” or “encouraging” robot, though non-experts also employed emotions (e.g., “happy”) and supportive behaviors (“attentively listening”) in their descriptions. Non-experts also described individual backchanneling behaviors at a more granular level, and understand how to engage in authoring these individual, fine-grained behaviors.

RQ3: How do we incorporate non-expert user mental models of robot backchanneling into the design of a novel interface for authoring robot backchanneling behavior?

Through our iterative prototyping and user testing (Chapter 4) and evaluation study (Chapter 5), we uncovered that using coarse-grained personality *styles* to describe overall robot backchanneling behavior was compatible with how our non-expert participants conceptualized robot backchanneling, and supported their successful authoring of these behaviors. Incorporating styles into the approachable block-based programming paradigm—by letting users append styles to any part of dialog—enabled users to quickly experiment with applying different styles to their dialogs, and to quickly edit or remove styles as they saw fit. This interaction design leverages a combination of our novel understanding of user mental models and an empirical, intuitive user interface to support non-experts in prototyping rich robot dialogs with backchanneling.

Given our findings that non-experts also understand and engage individual behaviors, providing fine-grained control complemented the styles approach by giving non-experts the ability to engage backchanneling in a granular fashion. In our evaluation study, participants explored and utilized both coarse- and fine-grained controls to creatively author robot backchanneling.

Through our multi-phased approach, we addressed our research questions proposed in Chapter 1.2. The sum of this work gives us an understanding of how non-experts conceptualize and engage with robot backchanneling, and how we can design interactions and interfaces that enable their participation and support their creativity in authoring robot backchanneling. We turn now to recommendations for robot behavior authoring tools going forward.

6.2 Recommendations

Based on our elicitation study, we synthesized a set of design recommendations for researchers and designers of robot behavior authoring tools to consider (Chapter 3.4). Considering all phases of this work and their outcomes holistically, we present the following abbreviated and encompassing recommendations as an overarching message to system designers:

1. Design Dialog Authoring Tools to Include Backchanneling

We found that non-expert end users of robot authoring tools understand and engage with backchanneling and apply backchanneling to their robot designs. Dialog authoring tools should expose backchanneling as an option for users to create rich robot dialog.

2. Leverage Personality as a Coarse-grained Means of Authoring Backchanneling

An overwhelming success of our interface design revolved around our presentation of backchanneling behaviors as personality ‘styles’, which mapped neatly to non-expert mental models of backchanneling and was effectively employed in their authoring. Personality may be used to describe backchanneling or, as explored by our participants, many facets of robot behavior.

3. Give End-users Granular Authoring Controls for Robot Behaviors

Many of our participants authored backchanneling using a combination of personality and adjustments of individual behaviors. Based on our results, non-experts are not afraid to engage individual behaviors when authoring robot backchanneling, and may prefer this level of access and control over the robot’s ultimate behavior.

6.3 Limitations

Our work possesses key limitations that future work should consider. First, our explanation of backchanneling to participants in the design workshops (prototype evaluation) may have primed them toward authoring certain behaviors or to focus on backchanneling. While in our elicitation workshops, explaining backchanneling helped ground participants' knowledge and guide their brainstorming and exploration (after an unprimed task to first elicit natural perceptions and awarenesses), we acknowledge that maintaining this step for the design workshop may have affected authoring results. However, our findings on participant ability to engage in creative authoring using our interface design are still significant.

Second, the nature of our user participation limits our ability to reflect on real world use. Our users were voluntary research participants who were compensated, and thus not necessarily intrinsically motivated in the manner an eventual real-world user may be. While our focus was on achieving an interface design that successfully leveraged user mental models and enabled them to create, future work is required to assess how such an interface design (or any manner of end-user creation for backchanneling) is used naturally in real world scenarios.

Further, while our participatory methodology prioritized including participants in multiple phases of the work to emphasize a focus on their mental models throughout the entire design process, including participants in multiple phases likely affected the results of our evaluation from the design workshop. Many participants in this final workshop—although happy to participate in such iterative research where they witnessed their designs come to life—had received our training on backchanneling behaviors in prior sessions, limiting our ability to reflect

on naïve learnability. Future work could evaluate tools for authoring backchanneling on completely inexperienced users to address any potential gaps left by our methodology.

Finally, we note technical limitations arising from our focus on enabling participants to author backchanneling using mechanisms that align with their mental models. Our prototype was incomplete and lacked the polish of a commercial tool, making it difficult to examine real, natural use given its current limitations. For instance, our prototype lacked conditional branching—a key feature of authoring tools that was also highly requested by our participants. Further, participants were required to manipulate the granular behavior controls to create new personality styles, which runs counter to our focus on letting the user engage at whichever layer they desired. Future work should incorporate authoring backchanneling as a component of personality in more robust authoring tools.

6.4 Conclusion

We present a novel interface for end users to intuitively and creatively author robot backchanneling. Through an iterative, participatory, user-centered design process, we discovered that non-expert end users intuitively understand robot backchanneling and readily engage in authoring these behaviors to create customized human-robot dialog. Our work demonstrates the range of possibilities when end-user participation in authoring rich, nuanced robot behaviors is enabled by leveraging the mental models of non-experts in interface design. We hope that our insights inform future user-centered design of robot behavior authoring tools to support user creation and personalization of rich and nuanced dialog behaviors.

Appendices

Appendix A: Elicitation Workshop Study

A.1 Recruitment Poster



University
of Manitoba



Interested in robots? We need your ideas!

We are looking for people to help us create tools to design robots that better suit people's needs.

Looking for those who are interested in designing robots, even if you don't have any technical skills.
Must be 18 years old or older.

Participation involves
one 2-hour session at
the University of Manitoba
Fort Garry Campus, including:

- Brainstorming in small groups
- Collaborative design activities
- Group discussions

Compensation
Receive \$35 (CAD) for your time.

If interested, please contact
Sam Barrett
barret15@myumanitoba.ca



Principal Investigator
Dr. James E. Young
young@cs.umanitoba.ca
University of Manitoba

This research has been approved by the Research Ethics Board at the University of Manitoba, Fort Garry campus.
The Research Ethics Board can be reached by phone 204-474-7122 or email humanethics@umanitoba.ca

A.2 *Recruitment Email Template*

Variants of this email template were used for all studies.

Hello <name>,

My name is Sam and I will be your primary contact for communication about our study. Thank you for being interested in this study. By being a part of this research, you will help us learn about how people think about talking robots and how we could design a tool for anyone to create a robot's talking behaviours.

In this email we will provide you with detailed information about the study, please feel free to ask any questions you may have.

First, we'd like to be clear that this study is not a test. You will not be evaluated nor scored for any sort of performance, rather, we are looking to gain insight into how people think about robots, to learn from you and develop ideas for how we can make better ways of designing robots.

If you choose to participate, you will attend a workshop located on the University of Manitoba's Fort Garry campus with 5 to 9 other participants, myself, and an assistant moderator. The session will be scheduled for 2 hours. Here's an overview of what will happen in the session:

We will start the session with refreshments and an introduction from myself, and a quick icebreaker activity. Then, we will learn about robots that talk, including a video demonstration of a talking robot. We will do a few activities in small groups that involve thinking about talking robots. Then, using paper, markers, and other basic crafting supplies, you will do an activity to think about how you might make a robot talk in the way you desire. At the end, we will have a quick discussion about our thoughts on talking robots.

Our study is designed for anyone who is interested in robots. You do not need to have any technical, robotic, or computer skills or experience to participate in this study.

Please note that we will not force anyone to speak or share their thoughts at any time, we welcome you to share your ideas only if you feel comfortable doing so.

Additionally, please find attached an Informed Consent Form detailing more information about the study. If you have any questions at any time, please reach out to us at this email. If you decide to participate, you will be given a copy of this Informed Consent Form to sign upon your arrival at the workshop.

In appreciation of your time, you will receive a cash honorarium of \$35 during the session.

If you would like to participate in this study, please reply to this email letting us know, and choose 2 or 3 of the following dates and times that work for you. We will do our best to match you to your preferred time.

Available times:

<a list of times: day of week, month, date, time>

Thank you for your time,

<researcher sign-off>

A.3 *Informed Consent Form*



**University
of Manitoba** Faculty of Science

Department of Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

Informed Consent Form

Study Title Participatory Design of End-User Tools for Human-Robot Interaction

Principal Investigator

Dr. James E. Young, Professor, Department of Computer Science, Faculty of Science,
young@cs.umanitoba.ca.

Co-investigators

James M. Berzuk, PhD Student, Department of Computer Science,
berzukj@myumanitoba.ca.

Samuel Barrett, MSc Student, Department of Computer Science,
barret15@myumanitoba.ca.

Conflicts of Interest and Undue Influence

None of the researchers have any conflicts of interest in this study/research.

Invitation to Participate

You are being invited to participate in a research study. This consent form, a copy of which has been given to you, is only part of the process of informed consent. If you want more details about something mentioned here, or information not included here, feel free to ask any of the people named above. Please take the time to read this document and any accompanying information carefully. It is very important that you understand:

- what is being asked of you
- what the risks and benefits of participation are, and
- how the information you provide will be used and stored.

Purpose of the Study

This research is being conducted under the supervision of Dr. James Young. This study is about exploring how people think about robots that talk with people. We are interested in how people think a talking robot should act, and how they would personally design a robot's talking behaviour. Ultimately, we hope to learn how we can best create a tool that helps people design their own talking robots.

Study Procedures

If you decide to take part in the study, we will meet in this room for this workshop for the next 2 hours, where you will learn about talking robots and complete individual and group-based robot design activities. Your ability to participate in future studies will not be impacted by your decision to take part or not in this study.



Department of Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

With your consent, the workshop may be audio recorded on a digital audio recording device for later analysis. Any recordings will not be shared outside of the research team. If you or another participant do not consent to be audio recorded, we will simply proceed without audio recording and instead take written notes.

First, you will be asked to complete a short demographics questionnaire on paper.

You will learn about talking robots via a video demonstration and an explanation from the Researcher. Then, you will do a group brainstorming activity to spot the difference between how a robot talks and how a person talks and come up with a list of potential things a robot might do during a conversation.

Then, you will work with paper crafting supplies (pens, markers, sticky notes, scissors and glue) to describe how you might design a robot to act during a conversation.

At the end of the session, a group discussion will provide you with an opportunity to share any feedback or ideas you may have.

Study Risks

This study poses minimal risk, meaning that the risk is no greater than what you might encounter in your daily life.

Study Benefits

Participating in this study may provide the opportunity to learn about talking robots, and how we can make robots have better conversations with people by making them do certain things.

This benefit is not guaranteed as it depends on your level of interest in the topic.

Compensation

You will be given \$35 cash to compensate you for time and expenses for participating in this study.

Confidentiality

All the information you provide as a participant in this study is confidential, which means we as the research team must keep it safe. We will do our best, however, it is not possible to guarantee absolute confidentiality. We will only share your personal information if required through court order or law.

Please do not share what other people say during the workshop with anyone else. We ask this to respect everyone's privacy. It's important to remember that we can't control what people may share outside of the group. Therefore, you will not be required to speak or share your thoughts at any point, and we invite you to share only what you are comfortable sharing.



Department of Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

We will ask that you put away your smartphone or other electronic devices for the duration of the workshop today.

Use and Storage of Information

Your information will be stored on a secure platform approved by the University of Manitoba. Physical artifacts, like this consent form, will be stored in a locked cabinet in the office of the Principal Investigator. The research team will see your personal information.

After we collect your information, we will take out anything that might identify you. We will also replace your name with a code (a participant number). The research team keeps a file that links each participant's code with their name and contact information. We keep this file separate from your other research information.

Please note that audio recordings of your voice may be identifiable information that cannot be coded. However, we will not publicly disseminate any audio recordings, and will only use them internally for data analysis.

We may share some of your information as quotes. The quotes will not include your name or information that may identify you.

We will destroy any identifiable information like your name, audio recordings, and contact information by June 2027.

Dissemination

We will share the results of the study through 1) possible inclusion in student theses; 2) publication to academic journals or conference presentations; 3) a summary report of the study results for interested participants, which we expect to send by March 2026.

Withdrawing

Your participation in this research study is voluntary. You can choose to do only the activities you are comfortable with. You may withdraw your participation from the study for any reason. You do not have to explain why. You will not be penalized in any way. The honorarium is yours to keep. If you withdraw partway through the study, we will do our best to destroy your information. It may not be possible where some of your answers are already combined with other participants'.

To withdraw from the study while attending the workshop, please simply tell the facilitator and leave.

To withdraw from the study after attending the workshop, please contact Samuel Barrett at barret15@myumanitoba.ca. Withdrawal must be requested by September 2025 (MMYY=0925) at the latest, after which date we will be unable to remove your data from analysis and potential submissions.



Department of Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

Questions or Concerns

Designated University of Manitoba personnel may look at your information to see that the research is being done in a safe and proper way.

This study has been reviewed and approved by a Research Ethics Board at the University of Manitoba. However, this does not mean that participation is risk-free. If you have questions about your rights as a research participant, you may contact the Office of Human Research Ethics at humanethics@umanitoba.ca or (204) 474-7122.

If you have any questions, concerns, or complaints about this study, you may contact any members of the research team listed on the first page or the Office of Human Research Ethics.

Consent

By signing this document, I agree that:

- I have read the above information or had it read to me.
- I have had the opportunity to ask and have had all of my questions answered.
- I understand what is being asked of me.
- I will be taking part in a research study.
- I may freely stop or leave the research study activities at any time.
- I understand what the researchers will do with the information I give them.
- I do not waive my legal rights by participating in the study.

Notice Regarding Collection, Use, and Disclosure of Personal Information

Your personal information is being collected under the authority of The University of Manitoba Act. The University of Manitoba is committed to preserving your right to privacy. The information you provide will be used by the University to support our research. Your personal information will not be used or disclosed for other purposes, unless permitted by The Freedom of Information and Protection of Privacy Act or The Personal Health Information Act. If you have any questions about the collection of personal information: Ph: 204-474-9462 or Email: fiippa@umanitoba.ca



Department of Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

I agree to participate in this study.	☐ YES ☐ NO
---------------------------------------	--

Summary of Results

If you would like to receive a summary of the study results, please provide an email address.

Email address: _____

I agree to be audio-recorded in this study.	☐ YES ☐ NO
I agree to be contacted by a member of this research team for future phases related to this research study. <i>By agreeing to be contacted, you are not required to participate in any future phases. We will keep your contact information for the duration of this multi-phase study. If you wish to withdraw your name from this list, please contact Samuel Barrett at barret15@myumanitoba.ca.</i>	☐ YES ☐ NO
I agree that the research team can use photos and videos of sketches, drawings, or other artifacts that I have produced for this study when they share the study results.	☐ YES ☐ NO

_____ Name of Participant	_____ Participant's Signature	_____ Date
------------------------------	----------------------------------	---------------

A.4 Demographic Questionnaire



**University
of Manitoba**

Questionnaire

This sheet will be collected and securely stored for analysis purposes only. If you would rather this sheet be destroyed please place in the shred bin.

Participant #

Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

My gender identity is [select all that apply]:

☐ Woman

☐ Man

☐ Non-Binary/Trans/Gender Diverse

☐ Or Please Specify: _____

☐ Prefer not to answer

My age range is:

☐ 18-29

☐ 30-39

☐ 40-49

☐ 50-59

☐ 60-69

☐ 70-79

☐ 80-89

☐ 90+

Please briefly describe your experience with, or understanding of, robots:

A.5 Elicitation Workshop Protocol Script

Overview script for opening and closing the workshop, including informed consent process.

Intro/Outro Overview Script

The moderator will follow the below script as a guiding reference for the structured activities and discussions, adapting as needed to participant interaction and questions. [Blue text](#) designates lines to be spoken by the moderator. Phase-specific scripts are noted and link to external files.

Script: Introduction

Participants arrive to a university room with a boardroom table, or the desks arranged to support participants to work in groups; simple refreshments including canned soft drinks, water, and pastries (e.g., Tim Hortons donuts or local bakery goods as available) will be placed on a table. In the room will be robots (the SoftBank Pepper, the SoftBank Nao, the Sony aibo, and the SnuggleBot), currently turned off, that will be used as support props during the workshops, which participants will be able to see right away.

The moderator welcomes everyone and asks them to take a seat wherever they like. As they take their seats, the consent forms are passed to each participant.

Hello everyone, we're going to get started.

First, welcome to the University of Manitoba campus if you haven't been here before, and if you have, welcome back.

My name is [<moderator name>](#) I'm with the Human-Robot Interaction laboratory, the work we do is all about how people and robots interact with each other. This is [<assistant moderator name and role>](#), will be sitting in today to assist as necessary.

We really appreciate that each of you are here today and have devoted your time to make this research possible, we really couldn't do what we do without you and your commitment. I'm excited to hear from you and I think we're going to have fun today.

I'll start with an acknowledgement that we are here today on the original lands of the Anishinaabeg, Ininiwak, Anisiniwuk, Dakota Oyate, Dene, and Inuit, and the National Homeland of the Red River Métis. I believe it's important for myself to reflect on how this research, education, the infrastructure around us today, these are all tied to this land and the Treaties that were made on it, and so I acknowledge this reality, I acknowledge the harms and mistakes of the past and present, and I dedicate myself to moving forward with this awareness in Reconciliation.

Today you're taking part in a study we're conducting to figure out what people think about robots that talk to people, and more importantly, how people might be able to design how a

robot talks all by themselves, even if they have not had any experience with robots before. We're interested in this because we want to make tools that enable people such as a hobbyist or someone just interested in robots, to create their own talking robot that behaves in the way that the person wants. To build something like that, we need to explore some ideas with people who are interested in robots, people like you. Of course, you don't need any experience or expertise with robots or computers for today's activities.

Today we're going to do some activities all about talking robots. You'll get to watch a video of a talking robot, we'll do some brainstorming about what a robot might do, we'll talk about our thoughts and ideas about talking robots, and you'll also get to design your own imaginary talking robot.

I'd like to emphasize that you are not being evaluated for any kind of performance, or score, this is not a test. We're interested in exploring your ideas with you, we want to see what you think. This means that all your ideas are entirely welcome here, if you have an idea but you think you're not sure if it's a good one, I encourage you to try it anyway, the more the merrier. We welcome all sorts of feedback, and I encourage you to voice any comments or questions or concerns you have.

You are welcome to share all of your thoughts on the things we'll be talking about today, I encourage you to speak up and speak your mind, but only if you feel comfortable doing so. Only share if you're comfortable doing so, I won't call anyone out and pressure you to speak.

I would also ask that we respect others' ideas and contributions, the research team is very open to all ideas, so there are no wrong or bad ideas today.

Finally, please refrain from using our smartphones or any other electronic devices for today's study, so please no pictures, videos or audio recordings. If you would like to take photos or videos of the robots, we can provide an opportunity after the study. We will be keeping an audio record of our discussions here today, and will be taking some photos of the ideas we come up with together.

Any questions or concerns before we begin?

Script: Obtain Informed Consent

Consent forms (06-ConsentForms.pdf) will be randomly numbered and distributed based on seating position of the participants. Further, we will assign participants a separate participant number that they will be asked to note on all artifacts they create; we will write the participant number on a sticky note, placed on the table in front of their seat, so they do not have to recall their number from memory. We will create a coding sheet that links the consent form number to the participant number, which is necessary for identifying data to be destroyed if

the participant withdraws from the study. Any artifacts the participant does not want the researchers to collect can be placed in a “shred bin” on the table, the contents of which will be destroyed using the Department of Computer Science’s secure shred bin immediately following the session.

If all of this sounds good, please take a moment to carefully read the consent form, fill the noted fields, and return it to me. Once you hand in your consent form, I will give you an envelope with your honorarium and we will continue the study.

We’re going to use what we learn in this study to develop a prototype tool. If you’re interested in trying out our prototype, you may check the box at the end of the consent form that allows us to reach out to you later for participation in the follow-up study as well. It’s not required.

On the table in front of you is an informed consent form, please read through it carefully, take your time, and I’m happy to go over any part of it if you have questions. If you decide to consent to participating today, I’ll collect your form once signed and you’ll get a copy of the form to keep.

The moderator waits for participants to complete the forms and then collects them.

Thank you, now before we begin, I’ll now give you your honorarium, this is yours to keep.

The moderator disburses honorariums to participants.

Script: Demographics Questionnaire Procedure

Demographics questionnaire forms (09-Questionnaire.pdf) are passed out to each participant.

First, I’d like you to complete this demographics questionnaire form. It asks about your basic demographic information and briefly about your experiences with robots or your thoughts about them. Please do not write your name, but instead write your participant number, located on the sticky note in front of you. Please take a moment to fill out the form and return it to me, if you have any questions feel free to ask.

Wait for participants to fill the forms and collect them.

Script: Group Formation and Icebreaker [Phase 1 Only]

Okay, let’s start with a little warm-up. I will start by forming you into groups based on who you are sitting with, please let me know if you’d rather move for any reason. We’re going to be doing groups of 2-3 people, and throughout this study you’ll be working in this groups to come up with ideas before we all discuss those ideas together.

The moderator forms groups simply by how people are seated, into 3 groups of 3. In the case that fewer than 9 participants showed up remaining participants will be placed in groups of 2 (e.g., a group of two if 8 participants show up, two groups of two and a group of three if 7).

To help everyone get acquainted and familiar, I'd like you all to take 5 minutes to briefly introduce yourselves to your new group members, and to tell each other something you like to do in your free time. I'll give you a few minutes for that now.

Script: Core Tasks

Now we complete the core tasks of each respective phase variant, as detailed in the respective files:

- 02-Protocol-Phase1a.pdf
- 03-Protocol-Phase1b.pdf
- 04-Protocol-Phase2a.pdf
- 05-Protocol-Phase2b.pdf

Script: Closing

Okay, so that concludes our session today. It was nice to meet all of you, I had a great time today, and once again I want to thank all of you sincerely for dedicating your time to this study, this has been valuable to us, and it helps us make a better tool that people will find useful in the future, maybe when they get to design their first robot and they use our tool or a similar one.

When the study is concluded you'll receive a summary of our findings by email which will tell you all about what we do next and how your insights today helped us make this research possible. We will deliver this by February 2026 (MMYY=0226).

Please help yourself to more refreshments on your way out today if you'd like. Thank you.

Following is the script for the remainder of the workshop (tasks).

Protocol: Phase 1a Variant Details

Protocol

The core activities of this variant begin with a short lecture (~15 minutes) from the moderator including two short video demonstrations. Following, we will conduct a group brainstorming activity about the ideas from the lecture (10 minutes). Next, we will conduct a second short lecture where the moderator explains the idea of backchanneling (10 minutes), followed by a second brainstorm and standard card sort task (20 minutes). A design exercise using standard paper prototyping methods follows (30 minutes), before a final group discussion (15 minutes).

Lecture: Compare Dialogs

Note that this script continues from file [01-ProtocolOverview.pdf](#).

We will administer a task where participants watch two videos of dialog (explained below) and work together to examine the differences between them. The goal here is to help participants warm up to the tasks, and to get them to think about dialog and everything that goes into a conversation.

Okay, it seems like we're all acquainted with one another, glad to see we're getting along.

Our first activity today will be to watch two short videos, and while we're at it, we're going to try and spot as many differences as we can. The first video is going to show a person having a conversation with a robot, and the second video is going to show two people having a conversation with each other. What I want you to do is pay attention to see if you can notice any differences between how the people and robots are communicating during their conversations. The topic of conversation will be the same.

We will all do this individually but of course you should feel free to talk to your group members if you wish. After we watch the videos, we'll spend a few minutes in our groups to discuss what we saw.

I will hand out papers for you to work on. Every paper we write on today will have a spot in the upper corner where I'll ask you to write down your participant number, please write this down on every piece of paper you write on today. This is just to help us collect and analyze the data afterward, and to connect different things you may write. I'll circulate and collect the papers when we're done with each one, but if you would prefer we not look at it later,

you can place it in the shred bin on the table <point to bin>, and we will shred that today, right after we finish the study.

Hand out the blank data collection sheets and ensure participants know where to write their number.

The first video is of a person talking with a robot.

Show Video 1. This is a roughly 60 second video where a person talks with a VStone Sota robot about the weather (Figure 1). The robot will speak to the person emotively and gesture its arms, but while listening to the person speak back to it, the robot will remain still and silent, i.e., it will not perform any backchanneling behaviors.



Figure 1: A person talking to a VStone Sota robot.

Okay, so that was a person talking with a robot. Now let's see the next video of two people talking to each other.

Show video 2. Following, the moderator will show a second video of two people talking. Both people make use of typical backchanneling behaviors as they listen.

And that was a conversation between two people.

Put the videos on a screen looping without sound, to be used during the next task.

Activity: Brainstorm / Spot the Difference

Participants will be asked to work in their groups to study the two videos and to try and find differences between how the robot communicates and the person, given 5 minutes.

Now let's play spot the difference. Let's turn to our groups and now we'll work as a team. I'll pass around some paper and pencils so your groups can write down all the differences as you discuss them, and we also have bins of stationary supplies in the center of the table. We're going to spend 10 minutes brainstorming, plenty of time, there's no rush.

At the end, we're going to share the ideas that our groups had, so if you'd like to have maybe a representative from your group share your group's ideas, anyone is invited to speak or not speak according to their comfort. Any member of the group is welcome to volunteer an idea at any time, and no group or members will be pressured to speak or volunteer ideas if you do not want to. Any time you have a question during the activity, feel free to ask me.

Any questions about this one before we begin?

Let the group play spot the difference. Help guide the activity as needed.

Get a dry erase marker and write Differences on the whiteboard.

Okay everyone, let's bring it back. At this point I want to collect the ideas that people gathered and put them on the whiteboard. What are some of the differences we found? Would anyone like to go first?

Accept all ideas. Proceed when no further suggestions are made.

Okay, is that mostly everything we noticed? Do we feel okay with this list?

Okay, these are some great ideas, I really like the range we got. Now that we've seen that and thought about it, I want to focus in on some of these, the things that the robot is doing while it's listening to the person speak. We're especially interested in this part.

If any backchanneling behaviours were identified, circle them.

Lecture: Learning About Backchanneling

The moderator will explain the idea of backchanneling to participants to bring the group to a shared understanding of the topics being explored, as many participants are unlikely to be familiar with the concept or may not be focusing on this.

Let's zero in on what's happening when someone listens during a conversation.

When the people were listening to the other person speak, what were they doing?

Right, they weren't really talking, but they were still engaged in the conversation, we could tell that they were actively listening and were showing this in a range of ways.

What did the robot do while it was listening?

That's right, the robot did nothing. It just sat there, entirely still and devoid of any reactions, or emotions, or anything that might indicate that it understands what's going on, whether it wants to hear more, etc.

When the two people were having a conversation, whoever was listening to the other was really active, they were nodding their head, saying little things like "mhm", "yeah", they made faces. These are all things we often do when we're listening to someone.

These little things we do are small, simple, but they mean so much. They help the other person know that we're paying attention, we understand what you're saying, or even if we shake our head, we can show we disagree. If we make a weird face, we can show that we're confused or surprised.

It even helps with taking turns, so we don't interrupt each other so much. We can signal to each other when we'd like to speak. Ever been on a phone call, and you keep getting interrupted, or accidentally interrupting the other person? On the phone there's a bit of delay, and we can't see them, it can be really difficult to tell when they're *about* to speak.

These actions are a huge part of our communication overall. Since robots talk with people too, we're interested in how robots can do this too.

This general idea is called "backchanneling." Don't worry, you do not need to memorize this, but I will use the term to refer to this part of communication.

Flip to the next slide listing the various functions of backchanneling.

It turns out, it can be a bit difficult to figure out what a robot should do if we want it to backchannel. We're interested in exploring how you think of robots that do this, and how you might design a robot that does these things.

First, we need to make a list of everything that you think a person or a robot might do to backchannel to another person. So, let's do that.

Activity: Backchanneling Brainstorm and Card Sort

Participants will be asked to brainstorm again in their groups, focusing on backchanneling specifically. This will help us understand how people think about backchanneling, what kinds of aspects they tend to think about (or not), and what language they use, to inform our understanding of their mental model.

Let's turn back to our groups, to brainstorm as many things as we can that a person or a robot might do while listening to someone. Everything and anything counts, it could be saying a word, or multiple words, or breathing in a certain way, or any body language or action, any facial expression, anything you can think of that someone might do while

listening to another person to indicate any of the backchanneling purposes mentioned earlier.

For this one, I'm going to hand out index cards, and as your group thinks of things the robot could do, I want you to write one thing on each card. If you want to, you can consider designating one group member to write the cards while the others help discuss ideas. We have tons of cards, so use as many as you like. Again, please put your number in the corner of each card you write. We will use them for the next part of this activity.

We're going to spend about 10 minutes doing this, any questions before we start?

Let participants write down their ideas for 10 minutes.

Alright, now that we've got a bunch of ideas written down, I would like us to work together to categorize all these things we can do. I'm going to place some labels around the table, and I want us to get up and walk around, and place each of our cards in the category where you feel it's a good fit.

For example, if you have "nod head" and the category is "show that I understand", and you think that nodding your head is a way to show that you understand, put it down.

We also have extra labels, so if you don't think that your card fits any of the existing groups, then you can make your own!

If you really aren't sure, we have an "I don't know" category, so if you get really stuck on one, put it there.

You might notice that one of your cards fits in two or more categories. That's okay, take another card and write down the same thing so you can place them in both categories.

Spread the cards out instead of stacking them so we can see all the cards at once, and when we're done, <assistant moderator> will take a picture of the cards.

There's no rush for this, we'll just go until we're done, when we've all found places for all our cards and written up more cards where needed.

Let participants categorize their cards. Circulate and assist as needed.

Alright, now that we've organized these actions into their respective purposes, and we've thought about the different purposes of backchanneling, I want you to try and think of any other action that you haven't thought of yet. Perhaps you realized something as you were walking around, or you noticed a category that you didn't fill as much. Let's use a different colour of index card for this one, please use the <color> index cards and take a few minutes to try and come up with a couple more things.

Hand out colored index cards, and let participants complete the task for a few minutes.

Activity: Design Human-Robot Dialogs

Participants will design their own robot dialog to get into the headspace of what it is like to create dialog.

Okay, now that we've got a good idea of what kinds of things a robot or a person could do to backchannel, we're going to try designing a potential dialog between a person and a robot. To do this, we're going to work from a scenario as our starting point.

Here's a robot, its name is Nao, and it's talking to our colleague.

The scenario we'll be working from is a robot that asks a person questions about their day. Start here, think of what the robot might ask the person, what the person might say back, and create a conversation between the two.

Walk through an example on the projector of filling in the dialog according to the theme of a robot asking a person about their day.

I'm going to hand out sheets of printed paper that have speech bubbles for a person and a robot. I will give you a scenario or theme to work from, and I'd like you to fill in what the person and robot might say in that scenario.

There's also a Robot description sheet. You don't need to use this, but you can put things here about the robot's general behaviour, if you want to or find it helpful to keep track. For example, I can say that my robot talks in a funny cat voice and it always sounds just a little bit angry.

Show an example of using the Robot data sheet on the projector.

I'm going to pass out a bunch of stationery – pens, markers, sticky notes, scissors and glue, and colored construction paper. Feel free to grab anything you want to use.

There are no right or wrong answers here, I just want us to get creative and try stuff.

I will circulate and help, brainstorm ideas with you. Talk to each other and share your ideas,, I encourage that if you're comfortable, but it's not mandatory.

If you feel stuck or unsure, I have some sample sheets you can see and work from.

You can create multiple dialogs, or use multiple sheets to write a single dialog, you're not limited to just one single idea.

As before, please write down your number on each sheet.

Are there any questions before we begin? I will circulate and help as we go.

Hand out dialog sheets, robot data sheets, and put crafting supplies on the boardroom table.

Circulate and help with the activity.

Activity: Annotate Dialog Designs with Backchanneling

Participants will take their created dialog artifacts and attempt to add backchanneling to their designs. This will be done using stationery supplies to annotate their dialogs with a description of robot's backchanneling behaviour. This helps us identify how participants think about backchanneling, e.g., whether they think of the robot's behaviour in terms of its individual actions or its personality.

Now that we've created our own conversations between a person and a robot, and we've thought about how our robot might behave, what I'd like you to do is to add some of the cues we explored earlier to our dialog for the robot to show that it is listening – we will add backchanneling.

You can add this or describe this in any way that you want. You can write it out specifically, or grab markers, sticky notes, and modify your sheets. Feel free to use any of the cues we talked about earlier or think of new ones, get creative with it. Feel free to use your Robot Data sheet as inspiration if you made one.

Any questions before we begin?

Circulate and help with the activity.

Activity: Group Discussion

The final activity consists of a group discussion to elicit any further thoughts participants may have on robot backchanneling and how they can design robots that backchannel. We will probe mental models using open-ended questions in a loose discussion format that allows participants to suggest ideas and build on each other's ideas.

Alright, that went well, I'm glad to see everyone trying different things, this is valuable for us to see.

As our final activity I'd like to have a group discussion or brainstorm about backchanneling, this is a time where you get to share any thoughts you've had today, whether it's about robots that do backchanneling, or how you might design backchanneling robots, this is the place to talk about it. I have a few questions to ask you as a group, anyone is free to share

their thoughts, but no one is required to speak if they don't want to. At the end, I'll open the floor to hear from you any further comments we may not have covered yet.

Choose from the following sample questions to ask the group, depending on available time remaining, and ask follow-up questions where relevant based on participant answers.

A list of probing questions that the moderator follows:

1. How did that activity go?

Standard focus group procedure. May elicit problems with the procedure or limitation that we did not foresee or observe.

2. Were there any problems, anything that you wished you could have done or would have been done differently?

Standard focus group procedure.

3. Which backchanneling behaviours would you want to make a robot do?

To elicit insight on which kinds of backchanneling participants gravitate toward for robots.

4. If you were using a digital tool for specifying how a robot backchannels, how would you like to describe the robot's behaviour to the tool?

To elicit ideas on how participants may expect to interact with such a tool. Examples for elaboration: would you prefer to say, "the robot should nod a lot", or perhaps describe more of a personality, like the character sheets?

5. Does anyone have any other thoughts they'd like to share on backchanneling robots?

Standard focus group procedure, to give participants the space to share any ideas we may have missed or failed to elicit in prior questions or activities.

Alright, so that concludes our discussion, and concludes the session today.

The moderator proceeds to closing remarks (01-ProtocolOverview.pdf).

A.6 *Authoring Worksheets*



**University
of Manitoba**







Data Collection Sheet

This sheet will be collected and securely stored for analysis purposes only. If you would rather this sheet be destroyed please place in the shred bin.

Participant #

Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

Dialog Sheet





**University
of Manitoba**

Data Collection Sheet

This sheet will be collected and securely stored for analysis purposes only. If you would rather this sheet be destroyed please place in the shred bin.

Participant #

Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

Robot Data Sheet



- Does your robot have a name?
- How does your robot act?
- How does your robot talk?
- What is your robot's character?
- What does your robot look like?

Feel free to use this paper to write, draw, or describe any aspect of your robot, thinking about how it may act in a conversation with a person.



University
of Manitoba

Data Collection Sheet

This sheet will be collected and securely stored for analysis purposes only. If you would rather this sheet be destroyed please place in the shred bin.

Participant #

Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

Robot Data Sheet



- Does your robot have a name?
- How does your robot act?
- How does your robot talk?
- What is your robot's character?
- What does your robot look like?

Feel free to use this paper to write, draw, or describe any aspect of your robot, thinking about how it may act in a conversation with a person.

Name: Penny

A stern, quiet, but compassionate robot,
who isn't afraid to show their real opinion.





**University
of Manitoba**

Data Collection Sheet

This sheet will be collected and securely stored for analysis purposes only. If you would rather this sheet be destroyed please place in the shred bin.

Participant #

Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

Dialog Sheet



How was your
commute today?



Oh, it was okay,
but construction
was bad...

I see. Was the
weather nice
at least?

It was fantastic,
the sun felt warm
on my face.

Appendix B: Formative Paper Prototype Evaluation

B.1 *Informed Consent Form*



Department of Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

Informed Consent Form

Study Title Participatory Design of End-User Tools for Human-Robot Interaction

Principal Investigator

Dr. James E. Young, Professor, Department of Computer Science, Faculty of Science,
young@cs.umanitoba.ca.

Co-investigators

James M. Berzuk, PhD Student, Department of Computer Science,
berzukj@myumanitoba.ca.

Samuel Barrett, MSc Student, Department of Computer Science,
barret15@myumanitoba.ca.

Conflicts of Interest and Undue Influence

None of the researchers have any conflicts of interest in this study/research.

Invitation to Participate

You are being invited to participate in a research study. This consent form, a copy of which has been given to you, is only part of the process of informed consent. If you want more details about something mentioned here, or information not included here, feel free to ask any of the people named above. Please take the time to read this document and any accompanying information carefully. It is very important that you understand:

- what is being asked of you
- what the risks and benefits of participation are, and
- how the information you provide will be used and stored.

Purpose of the Study

This research is being conducted under the supervision of Dr. James Young. This study is about exploring how people might use a tool to design the way a talking robot behaves in a conversation with a person. We hope to learn about how our tool helps people create talking robot designs, where the tool could be improved, and what kinds of robot behaviours people create with the tool.

Study Procedures

If you decide to take part in the study, we will meet in this room for this workshop for the next 1 hour, where you will learn about robots, about a tool for designing how robots talk, and you will have the opportunity to use the tool to make your own talking robot design. Your ability to participate in future studies will not be impacted by your decision to take part or not in this study.



Department of Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

With your consent, the workshop may be audio recorded on a digital audio recording device for later analysis. Any recordings will not be shared outside of the research team. If you or another participant do not consent to be audio recorded, we will simply proceed without audio recording and instead take written notes.

First, you will be asked to complete a short demographics questionnaire on paper.

You will learn about talking robots through a couple of quick video demonstrations and an explanation from the Researcher, including what kinds of things robots can do during conversations, and what behaviours we are interested in designing with you.

Then, you will interact with a provided paper tool to specify how a robot might talk in a conversation with a person. The moderator will give a tutorial to demonstrate using the tool.

You will then be invited to complete a short questionnaire about your experience with the tool and your thoughts on designing robot behaviors.

At the end of the session, a group discussion will provide you with an opportunity to share any feedback or ideas you may have.

Study Risks

This study poses minimal risk, meaning that the risk is no greater than what you might encounter in your daily life.

Study Benefits

Participating in this study may provide the opportunity to learn about talking robots, and how we can make robots have better conversations with people by making them do certain things.

This benefit is not guaranteed as it depends on your level of interest in the topic.

Compensation

You will be given \$20 cash to compensate you for time and expenses for participating in this study.

Confidentiality

All the information you provide as a participant in this study is confidential, which means we as the research team must keep it safe. We will do our best, however, it is not possible to guarantee absolute confidentiality. We will only share your personal information if required through court order or law.

Please do not share what other people say during the workshop with anyone else. We ask this to respect everyone's privacy. It's important to remember that we can't control what



Department of Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

people may share outside of the group. Therefore, you will not be required to speak or share your thoughts at any point, and we invite you to share only what you are comfortable sharing.

We will ask that you put away your smartphone or other electronic devices for the duration of the workshop today.

Use and Storage of Information

Your information will be stored on a secure platform approved by the University of Manitoba. Physical artifacts, like this consent form, will be stored in a locked cabinet in the office of the Principal Investigator. The research team will see your personal information.

After we collect your information, we will take out anything that might identify you. We will also replace your name with a code (a participant number). The research team keeps a file that links each participant's code with their name and contact information. We keep this file separate from your other research information.

Please note that audio recordings of your voice may be identifiable information that cannot be coded. However, we will not publicly disseminate any audio recordings, and will only use them internally for data analysis.

We may share some of your information as quotes. The quotes will not include your name or information that may identify you.

We will destroy any identifiable information like your name, audio recordings, and contact information by June 2027.

Dissemination

We will share the results of the study through 1) possible inclusion in student theses; 2) publication to academic journals or conference presentations; 3) a summary report of the study results for interested participants, which we expect to send by March 2026.

Withdrawing

Your participation in this research study is voluntary. You can choose to do only the activities you are comfortable with. You may withdraw your participation from the study for any reason. You do not have to explain why. You will not be penalized in any way. The honorarium is yours to keep. If you withdraw partway through the study, we will do our best to destroy your information. It may not be possible where some of your answers are already combined with other participants'.

To withdraw from the study while attending the workshop, please simply tell the facilitator and leave.

To withdraw from the study after attending the workshop, please contact Samuel Barrett at barret15@myumanitoba.ca. Withdrawal must be requested by January 2026 (MMYY=0126)



Department of Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

at the latest, after which date we will be unable to remove your data from analysis and potential submissions.

Questions or Concerns

Designated University of Manitoba personnel may look at your information to see that the research is being done in a safe and proper way.

This study has been reviewed and approved by a Research Ethics Board at the University of Manitoba. However, this does not mean that participation is risk-free. If you have questions about your rights as a research participant, you may contact the Office of Human Research Ethics at humanethics@umanitoba.ca or (204) 474-7122.

If you have any questions, concerns, or complaints about this study, you may contact any members of the research team listed on the first page or the Office of Human Research Ethics.

Consent

By signing this document, I agree that:

- I have read the above information or had it read to me.
- I have had the opportunity to ask and have had all of my questions answered.
- I understand what is being asked of me.
- I will be taking part in a research study.
- I may freely stop or leave the research study activities at any time.
- I understand what the researchers will do with the information I give them.
- I do not waive my legal rights by participating in the study.

Notice Regarding Collection, Use, and Disclosure of Personal Information

Your personal information is being collected under the authority of The University of Manitoba Act. The University of Manitoba is committed to preserving your right to privacy. The information you provide will be used by the University to support our research. Your personal information will not be used or disclosed for other purposes, unless permitted by The Freedom of Information and Protection of Privacy Act or The Personal Health Information Act. If you have any questions about the collection of personal information: Ph: 204-474-9462 or Email: fippa@umanitoba.ca



Department of Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

I agree to participate in this study.	☐ YES ☐ NO
---------------------------------------	--

Summary of Results

If you would like to receive a summary of the study results, please provide an email address.

Email address: _____

I agree to be audio-recorded in this study.	☐ YES ☐ NO
I agree to be contacted by a member of this research team for future phases related to this research study. <i>By agreeing to be contacted, you are not required to participate in any future phases. We will keep your contact information for the duration of this multi-phase study. If you wish to withdraw your name from this list, please contact Samuel Barrett at barret15@myumanitoba.ca.</i>	☐ YES ☐ NO
I agree that the research team can use photos and videos of sketches, drawings, or other artifacts that I have produced for this study when they share the study results.	☐ YES ☐ NO

_____ Name of Participant	_____ Participant's Signature	_____ Date
------------------------------	----------------------------------	---------------

B.2 Formative Paper Prototyping Protocol Script

Protocol: Phase 2as Variant Details

Protocol

The core activities of this variant begin with a short lecture to get participants up to speed on the concepts being explored (~10 minutes). Participants will then complete a standard paper prototyping task to become familiar with the task of designing robots on paper (10 minutes). Then, a tutorial on using the prototype will be given, and participants will interact with the paper tool to design the behavior of an imagined robot (20 minutes). Finally, we administer a standard questionnaire to elicit participants' experiences with the tool and their feedback (5 minutes), followed by a group discussion on the prototype (10 minutes).

Lecture: Learning about Backchanneling

Note that this script continues from file [01-ProtocolOverview.pdf](#).

The moderator will explain the idea of backchanneling to participants to bring the group to a shared understanding of the behaviours they will be designing for a social robot, as participants may be unfamiliar with the concept of backchanneling.

Our first activity today will be to understand a little bit about robots that talk with people, and the kinds of things they do.

We're going to start by watching two videos. The first video is of a person talking with a typical robot, and the second video shows two people talking with each other.

Show video 1.

Show video 2.

Let's zero in on what's happening in these two videos, and what makes those two conversations different.

When the people were listening to the other person speak, what were they doing?

Right, they weren't really talking, but they were still engaged in the conversation, we could tell that they were listening.

What did the robot do while it was listening?

That's right, the robot did nothing. It just sat there, entirely still and devoid of any reactions, or emotions, or anything that might indicate that it understands what's going on.

When the two people were having a conversation, whoever was listening to the other was really active, they were nodding their head, saying little things like "mhm", "yeah", they made faces. These are all things we often do when we're listening to someone, and this is sometimes called backchanneling.

These little things we do are small, simple, but they mean so much. They help the other person know that we're paying attention, we understand what you're saying, or even if we shake our head, we can show we disagree. If we make a weird face, we can show that we're confused or surprised.

It even helps with taking turns, so we don't interrupt each other so much. We can signal to each other when we'd like to speak. Ever been on a phone call, and you keep getting interrupted, or accidentally interrupting the other person? On the phone there's a bit of delay, and we can't see them, it's really hard to tell when they're *about* to speak.

These actions are a huge part of our communication overall. Since robots talk with people too, we're interested in how robots can do this too. So we designed the tool that you'll be using today, so that you can make a robot do the same thing that people do when we're listening to someone.

This general idea is called "backchanneling." Don't worry, you do not need to memorize this, but I will use the term to refer to this part of communication.

We're interested in exploring how you would choose a robot to behave, and whether the tool we made works well at enabling you to achieve that.

Activity: Design Backchanneling on Paper

Participants will complete a short design exercise to describe a robot's backchanneling behaviour on paper. This activity will help participants become familiar with thinking of robot dialog, and of robot feedback including backchanneling before moving to using the toolkit to design a real robot's backchanneling behaviour.

Before we use the toolkit, we're going to design how we want our robots to behave on paper. I'm going to pass around a sheet for describing the robot's behaviour. In the

That's right, the robot did nothing. It just sat there, entirely still and devoid of any reactions, or emotions, or anything that might indicate that it understands what's going on.

When the two people were having a conversation, whoever was listening to the other was really active, they were nodding their head, saying little things like "mhm", "yeah", they made faces. These are all things we often do when we're listening to someone, and this is sometimes called backchanneling.

These little things we do are small, simple, but they mean so much. They help the other person know that we're paying attention, we understand what you're saying, or even if we shake our head, we can show we disagree. If we make a weird face, we can show that we're confused or surprised.

It even helps with taking turns, so we don't interrupt each other so much. We can signal to each other when we'd like to speak. Ever been on a phone call, and you keep getting interrupted, or accidentally interrupting the other person? On the phone there's a bit of delay, and we can't see them, it's really hard to tell when they're *about* to speak.

These actions are a huge part of our communication overall. Since robots talk with people too, we're interested in how robots can do this too. So we designed the tool that you'll be using today, so that you can make a robot do the same thing that people do when we're listening to someone.

This general idea is called "backchanneling." Don't worry, you do not need to memorize this, but I will use the term to refer to this part of communication.

We're interested in exploring how you would choose a robot to behave, and whether the tool we made works well at enabling you to achieve that.

Activity: Design Backchanneling on Paper

Participants will complete a short design exercise to describe a robot's backchanneling behaviour on paper. This activity will help participants become familiar with thinking of robot dialog, and of robot feedback including backchanneling before moving to using the toolkit to design a real robot's backchanneling behaviour.

Before we use the toolkit, we're going to design how we want our robots to behave on paper. I'm going to pass around a sheet for describing the robot's behaviour. In the

speech bubbles, you'll write what you want the robot to say, and in the boxes, you'll indicate what the robot is doing while listening to the other person speak, like what kinds of backchanneling it might be doing.

I'm going to give you a sample scenario, and you'll write down what you want the robot to say and also indicate what the robot will do while listening to a person respond.

Change the slide on the projector to describe the scenario for the robot talking to a person.

The scenario is: a robot that asks questions about a person's day at the end of the day.

Think of a robot in this scenario and think of what it will say to the person. You do not need to know what the person might say back. Instead, the robot can just ask a bunch of questions to the person or maybe share a little bit about its day.

I will hand out stationery, including pens, coloured markers and sticky notes, you can write, draw, use colours, describe the robot however you'd like.

There are no right or wrong answers for what a robot should do, you're creating your own robot designs so feel free to be creative with it. The point here is not to make a great dialog, but just to practice what it is like to have to design a robot dialog.

Once we have our robot designs, then we'll move to the real robots and you'll use the computer to bring your designs to life.

Distribute dialog sheets, pens, markers and sticky notes.

Please take a moment to look over these sheets. You can see that on one side we have spaces to put words that the robot says, and that the person answers for some time. You also have space to specify how the robot should give feedback while listening.

Assist with the task, answer any questions and provide guidance where participants may be unsure of what to do.

Tutorial: Using a Prototype Tool to Design Robot Backchanneling

In this phase participants will use a prototype tool made of paper that includes standard interface elements like "buttons" and "dropdown menus", to allow specifying how a robot should backchannel in conversation (e.g., an option for "types of backchanneling" could feature a list of behaviors on which the participant places

small checkboxes next to the behaviors they wish the robot to use, such as “nodding” or “verbal acknowledgement”).

Introduce participants to the tool by showing examples on the projector.

In just a moment, we’re going get started designing our robots, so I’ll quickly introduce you to the toolkit and explain a little bit about how it works and how you can use it.

This is the toolkit. It works by arranging the paper elements to describe what the robot is saying. Once you’ve described what the robot will say, you can use this interface to specify how the robot will provide feedback while the person is talking. For example, I can “click” this gesture dropdown and I’ll see that I can choose <list gesture behaviours like nodding, shaking head> depending on what I want the robot to do.

Okay, any questions before you try it for yourselves? I’ll be circulating and helping as we go, feel free to ask me for help at any time.

Activity: Using the Toolkit

Participants will each sit in front of a complete paper prototype of the tool provided for them.

Let’s try this for ourselves. Now that we have an idea for the robot we’re designing, we will use the toolkit to recreate those designs on the real robot.

Please consider the conversation you worked on earlier, and design how you think the robot could behave in this instance. When you’re done, let me know and I’ll take a picture of the paper. We also have example conversations to build from, so when you’re done the first one, please take another.

You can call me over at any point, you don’t need to wait until the end.

I will circulate and help with any questions or difficulties you may have.

Circulate while participants use the toolkits.

Procedure: Questionnaire

We will administer a standard System Usability Scale¹ from Human-Computer Interaction to probe participant’s experiences using the toolkit. This Likert scale

¹ J. Brooke, “SUS: A ‘Quick and Dirty’ Usability Scale”, CRC Press, 1996.

questionnaire helps assess the subjective usability and effectiveness of a piece of software by its users and has been used in Human-Robot Interaction to assess user experiences with software toolkits for programming robots².

The questionnaire is attached in file [09-Questionnaires.pdf](#).

Okay, we'll stop there. Now that you've used the tool, I'd like you all to complete a brief questionnaire about your experience today using the tool and making your own talking robots. It's just a few quick questions, and we really appreciate the feedback that we get. One last time, please make sure to write your number on the sheet.

Pass out the questionnaire forms and wait for participants to complete them, then collect them.

Activity: Group Discussion

We will facilitate a group discussion to collect any further feedback not covered by the questionnaire.

Okay, now I want to have an open discussion about the tool and designing robot behaviors for conversations. This is a place to raise any questions, comments or other feedback that you may not have had the opportunity to say yet. We're looking for general feedback on the tool and your experience using it, anything that went well or anything that may have been challenging. Would anyone like to start?

Conduct unstructured group discussion if participants have any feedback, asking follow-up and clarification questions where relevant.

Proceed to closing remarks ([01-ProtocolOverview.pdf](#)).

² Cruz-Sandoval et al., "PODER: A Robot Programming Framework for Further Inclusion of People with Mild Cognitive Impairment in HRI Research," *Intl. Conf. Human-Robot Interaction, Melbourne, Australia, 2025*. <https://dl.acm-org.uml.idm.oclc.org/dol/10.5555/3721488.3721563>.

B.3 Task Instruction Script

Backchanneling Toolkit Paper Prototype

Procedure

I will first explain how this activity works: I will give you some instructions for you to create a dialog/conversation, and you'll use the interface to create it. There are a few steps that I'll walk you through, and then for the final task you'll get to create your own conversation without instructions from me. The main rule here is that I will ask you what to do, but I cannot tell you *how* to do it, so do what feels right to you.

1. Prescribed

Scenario: For this first task, the scenario is you're designing a robot that lives in someone's home. You'll make the robot's behaviors for when the person arrives home from work and is greeted by the robot. We are focusing on the dialog part of this behaviour.

1.1 *Create a simple dialog.* First, you will create a dialog that includes speaking and listening. Remember, we are also specifying the style of how the robot interacts with a person as part of our designs. Please use the paper prototype to create the following behavior:

First, please program the robot to say "Hi, how was your day at work?" and then listen to the user's response. Please program this using a "happy" style. Now that you've designed the behavior, please test the behavior to see what it looks like when the robot does it.

1.2 *Change the robot's listening style.* Okay, now that you've created this conversation, please make the robot act excited while it listens to the person speak. Could you make the robot use an excited style?

1.3 *Modify a granular behavior.* Okay, now when you're testing out the robot's excited style, you notice that it's nodding its head quite a lot when it's

excited, and that may be distracting, so please make the robot nod its head less often when it's excited.

1.4 *Create a new personality style.* Okay, so now you've got the hang of it, let's say the robot is annoyed when the person comes home, so you need the robot to show that it's annoyed using a style. Please make an annoyed style and use it in the conversation when the robot greets the person.

Running only one of multiple blocks: If I have multiple sets of blocks here, and I want to make sure only one of them runs, I just need to disable the other ones. We can do this by right clicking on the blue Robot Dialog block, and then clicking disable.

2. Open-ended

2.1 *Create a dialog based on a scenario.* Alright, now for the last part let's switch it up, I'd like you to create an interaction without having specific instructions from myself. I'll give you a scenario, but what the robot does is up to you. The scenario is a cooking assistant robot, so the robot might be giving instructions, suggesting ingredients or teaching the person how to cook the meal.

Moderator tips

If the user asks how to do something, say you cannot tell them and will make a note of it to fix.

Appendix C: Evaluation Workshop

C.1 *Informed Consent Form*



Department of Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

Informed Consent Form

Study Title Participatory Design of End-User Tools for Human-Robot Interaction

Principal Investigator

Dr. James E. Young, Professor, Department of Computer Science, Faculty of Science,
young@cs.umanitoba.ca.

Co-investigators

James M. Berzuk, PhD Student, Department of Computer Science,
berzukj@myumanitoba.ca.

Samuel Barrett, MSc Student, Department of Computer Science,
barret15@myumanitoba.ca.

Conflicts of Interest and Undue Influence

None of the researchers have any conflicts of interest in this study/research.

Invitation to Participate

You are being invited to participate in a research study. This consent form, a copy of which has been given to you, is only part of the process of informed consent. If you want more details about something mentioned here, or information not included here, feel free to ask any of the people named above. Please take the time to read this document and any accompanying information carefully. It is very important that you understand:

- what is being asked of you
- what the risks and benefits of participation are, and
- how the information you provide will be used and stored.

Purpose of the Study

This research is being conducted under the supervision of Dr. James Young. This study is about exploring how people might use a tool to design the way a talking robot behaves in a conversation with a person. We hope to learn about how our tool helps people create talking robot designs, where the tool could be improved, and what kinds of robot behaviours people create with the tool.

Study Procedures

If you decide to take part in the study, we will meet in this room for this workshop for the next 2 hours, where you will learn about robots, about a tool for designing how robots talk, and you will have the opportunity to use the tool to make your own talking robot design. With your consent, computer screens will be screen-recorded for later analysis. Your ability to participate in future studies will not be impacted by your decision to take part or not in this study.



Department of Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

With your consent, the workshop may be audio recorded on a digital audio recording device for later analysis. Any recordings will not be shared outside of the research team. If you or another participant do not consent to be audio recorded, we will simply proceed without audio recording and instead take written notes.

First, you will be asked to complete a short demographics questionnaire on paper.

First, you will learn about talking robots through a couple of quick video demonstrations and an explanation from the Researcher, including what kinds of things robots can do during conversations, and what behaviours we are interested in designing with you.

Next, you will design a robot's talking behaviours in the way that you prefer, first on paper and then using a computer and a real robot. You will use the computer to access software that lets you make the robot talk. No programming experience is required.

Finally, you will be invited to complete a short questionnaire about your experience with the tool and your thoughts on designing robot behaviours.

Study Risks

This study poses minimal risk, meaning that the risk is no greater than what you might encounter in your daily life.

Study Benefits

Participating in this study may provide the opportunity to learn about talking robots, and how we can make robots have better conversations with people by making them do certain things.

This benefit is not guaranteed as it depends on your level of interest in the topic.

Compensation

You will be given \$35 cash to compensate you for time and expenses for participating in this study.

Confidentiality

All the information you provide as a participant in this study is confidential, which means we as the research team must keep it safe. We will do our best, however, it is not possible to guarantee absolute confidentiality. We will only share your personal information if required through court order or law.

Please do not share what other people say during the workshop with anyone else. We ask this to respect everyone's privacy. It's important to remember that we can't control what people may share outside of the group. Therefore, you will not be required to speak or share your thoughts at any point, and we invite you to share only what you are comfortable sharing.



Department of Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

We will ask that you put away your smartphone or other electronic devices for the duration of the workshop today.

Use and Storage of Information

Your information will be stored on a secure platform approved by the University of Manitoba. Physical artifacts, like this consent form, will be stored in a locked cabinet in the office of the Principal Investigator. The research team will see your personal information.

After we collect your information, we will take out anything that might identify you. We will also replace your name with a code (a participant number). The research team keeps a file that links each participant's code with their name and contact information. We keep this file separate from your other research information.

Please note that audio recordings of your voice may be identifiable information that cannot be coded. However, we will not publicly disseminate any audio recordings, and will only use them internally for data analysis.

We may share some of your information as quotes. The quotes will not include your name or information that may identify you.

We will destroy any identifiable information like your name, audio recordings, and contact information by June 2027.

Dissemination

We will share the results of the study through 1) possible inclusion in student theses; 2) publication to academic journals or conference presentations; 3) a summary report of the study results for interested participants, which we expect to send by March 2026.

Withdrawing

Your participation in this research study is voluntary. You can choose to do only the activities you are comfortable with. You may withdraw your participation from the study for any reason. You do not have to explain why. You will not be penalized in any way. The honorarium is yours to keep. If you withdraw partway through the study, we will do our best to destroy your information. It may not be possible where some of your answers are already combined with other participants'.

To withdraw from the study while attending the workshop, please simply tell the facilitator and leave.

To withdraw from the study after attending the workshop, please contact Samuel Barrett at barret15@myumanitoba.ca. Withdrawal must be requested by [September 2025](#) [January 2026](#) (MMYY=09250126) at the latest, after which date we will be unable to remove your data from analysis and potential submissions.



Department of Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

Questions or Concerns

Designated University of Manitoba personnel may look at your information to see that the research is being done in a safe and proper way.

This study has been reviewed and approved by a Research Ethics Board at the University of Manitoba. However, this does not mean that participation is risk-free. If you have questions about your rights as a research participant, you may contact the Office of Human Research Ethics at humanethics@umanitoba.ca or (204) 474-7122.

If you have any questions, concerns, or complaints about this study, you may contact any members of the research team listed on the first page or the Office of Human Research Ethics.

Consent

By signing this document, I agree that:

- I have read the above information or had it read to me.
- I have had the opportunity to ask and have had all of my questions answered.
- I understand what is being asked of me.
- I will be taking part in a research study.
- I may freely stop or leave the research study activities at any time.
- I understand what the researchers will do with the information I give them.
- I do not waive my legal rights by participating in the study.

Notice Regarding Collection, Use, and Disclosure of Personal Information

Your personal information is being collected under the authority of The University of Manitoba Act. The University of Manitoba is committed to preserving your right to privacy. The information you provide will be used by the University to support our research. Your personal information will not be used or disclosed for other purposes, unless permitted by The Freedom of Information and Protection of Privacy Act or The Personal Health Information Act. If you have any questions about the collection of personal information: Ph: 204-474-9462 or Email: fippra@umanitoba.ca



**University
of Manitoba**

Faculty of Science

Department of Computer Science
University of Manitoba
Winnipeg, Manitoba
Canada
R3T 2N2

I agree to participate in this study.

☐ **YES** ☐ **NO**

Summary of Results

If you would like to receive a summary of the study results, please provide an email address.

Email address: _____

I agree to be audio-recorded in this study.

☐ **YES** ☐ **NO**

I agree that the research team can use photos and videos of sketches, drawings, or other artifacts that I have produced for this study when they share the study results.

☐ **YES** ☐ **NO**

Name of Participant	Participant's Signature	Date

C.2 Evaluation Workshop Protocol Script

Protocol: Phase 2as Variant Details

Protocol

The core activities of this variant begin with a short lecture to get participants up to speed on the concepts being explored (~10 minutes). Participants will then complete a standard paper prototyping task to become familiar with the task of designing robots on paper (10 minutes). Then, a tutorial on using the prototype will be given, and participants will interact with the paper tool to design the behavior of an imagined robot (20 minutes). Finally, we administer a standard questionnaire to elicit participants' experiences with the tool and their feedback (5 minutes), followed by a group discussion on the prototype (10 minutes).

Lecture: Learning about Backchanneling

Note that this script continues from file [01-ProtocolOverview.pdf](#).

The moderator will explain the idea of backchanneling to participants to bring the group to a shared understanding of the behaviours they will be designing for a social robot, as participants may be unfamiliar with the concept of backchanneling.

Our first activity today will be to understand a little bit about robots that talk with people, and the kinds of things they do.

We're going to start by watching two videos. The first video is of a person talking with a typical robot, and the second video shows two people talking with each other.

Show video 1.

Show video 2.

Let's zero in on what's happening in these two videos, and what makes those two conversations different.

When the people were listening to the other person speak, what were they doing?

Right, they weren't really talking, but they were still engaged in the conversation, we could tell that they were listening.

What did the robot do while it was listening?

That's right, the robot did nothing. It just sat there, entirely still and devoid of any reactions, or emotions, or anything that might indicate that it understands what's going on.

When the two people were having a conversation, whoever was listening to the other was really active, they were nodding their head, saying little things like "mhm", "yeah", they made faces. These are all things we often do when we're listening to someone, and this is sometimes called backchanneling.

These little things we do are small, simple, but they mean so much. They help the other person know that we're paying attention, we understand what you're saying, or even if we shake our head, we can show we disagree. If we make a weird face, we can show that we're confused or surprised.

It even helps with taking turns, so we don't interrupt each other so much. We can signal to each other when we'd like to speak. Ever been on a phone call, and you keep getting interrupted, or accidentally interrupting the other person? On the phone there's a bit of delay, and we can't see them, it's really hard to tell when they're *about* to speak.

These actions are a huge part of our communication overall. Since robots talk with people too, we're interested in how robots can do this too. So we designed the tool that you'll be using today, so that you can make a robot do the same thing that people do when we're listening to someone.

This general idea is called "backchanneling." Don't worry, you do not need to memorize this, but I will use the term to refer to this part of communication.

We're interested in exploring how you would choose a robot to behave, and whether the tool we made works well at enabling you to achieve that.

Activity: Design Backchanneling on Paper

Participants will complete a short design exercise to describe a robot's backchanneling behaviour on paper. This activity will help participants become familiar with thinking of robot dialog, and of robot feedback including backchanneling before moving to using the toolkit to design a real robot's backchanneling behaviour.

Before we use the toolkit, we're going to design how we want our robots to behave on paper. I'm going to pass around a sheet for describing the robot's behaviour. In the

speech bubbles, you'll write what you want the robot to say, and in the boxes, you'll indicate what the robot is doing while listening to the other person speak, like what kinds of backchanneling it might be doing.

I'm going to give you a sample scenario, and you'll write down what you want the robot to say and also indicate what the robot will do while listening to a person respond.

Change the slide on the projector to describe the scenario for the robot talking to a person.

The scenario is: a robot that asks questions about a person's day at the end of the day.

Think of a robot in this scenario and think of what it will say to the person. You do not need to know what the person might say back. Instead, the robot can just ask a bunch of questions to the person or maybe share a little bit about its day.

I will hand out stationery, including pens, coloured markers and sticky notes, you can write, draw, use colours, describe the robot however you'd like.

There are no right or wrong answers for what a robot should do, you're creating your own robot designs so feel free to be creative with it. The point here is not to make a great dialog, but just to practice what it is like to have to design a robot dialog.

Once we have our robot designs, then we'll move to the real robots and you'll use the computer to bring your designs to life.

Distribute dialog sheets, pens, markers and sticky notes.

Please take a moment to look over these sheets. You can see that on one side we have spaces to put words that the robot says, and that the person answers for some time. You also have space to specify how the robot should give feedback while listening.

Assist with the task, answer any questions and provide guidance where participants may be unsure of what to do.

Tutorial: Using a Prototype Tool to Design Robot Backchanneling

In this phase participants will use a prototype tool made of paper that includes standard interface elements like "buttons" and "dropdown menus", to allow specifying how a robot should backchannel in conversation (e.g., an option for "types of backchanneling" could feature a list of behaviors on which the participant places

small checkboxes next to the behaviors they wish the robot to use, such as “nodding” or “verbal acknowledgement”).

Introduce participants to the tool by showing examples on the projector.

In just a moment, we’re going get started designing our robots, so I’ll quickly introduce you to the toolkit and explain a little bit about how it works and how you can use it.

This is the toolkit. It works by arranging the paper elements to describe what the robot is saying. Once you’ve described what the robot will say, you can use this interface to specify how the robot will provide feedback while the person is talking. For example, I can “click” this gesture dropdown and I’ll see that I can choose <list gesture behaviours like nodding, shaking head> depending on what I want the robot to do.

Okay, any questions before you try it for yourselves? I’ll be circulating and helping as we go, feel free to ask me for help at any time.

Activity: Using the Toolkit

Participants will each sit in front of a complete paper prototype of the tool provided for them.

Let’s try this for ourselves. Now that we have an idea for the robot we’re designing, we will use the toolkit to recreate those designs on the real robot.

Please consider the conversation you worked on earlier, and design how you think the robot could behave in this instance. When you’re done, let me know and I’ll take a picture of the paper. We also have example conversations to build from, so when you’re done the first one, please take another.

You can call me over at any point, you don’t need to wait until the end.

I will circulate and help with any questions or difficulties you may have.

Circulate while participants use the toolkits.

Procedure: Questionnaire

We will administer a standard System Usability Scale¹ from Human-Computer Interaction to probe participant’s experiences using the toolkit. This Likert scale

¹ J. Brooke, “SUS: A ‘Quick and Dirty’ Usability Scale”, CRC Press, 1996.

questionnaire helps assess the subjective usability and effectiveness of a piece of software by its users and has been used in Human-Robot Interaction to assess user experiences with software toolkits for programming robots².

The questionnaire is attached in file [09-Questionnaires.pdf](#).

Okay, we'll stop there. Now that you've used the tool, I'd like you all to complete a brief questionnaire about your experience today using the tool and making your own talking robots. It's just a few quick questions, and we really appreciate the feedback that we get. One last time, please make sure to write your number on the sheet.

Pass out the questionnaire forms and wait for participants to complete them, then collect them.

Activity: Group Discussion

We will facilitate a group discussion to collect any further feedback not covered by the questionnaire.

Okay, now I want to have an open discussion about the tool and designing robot behaviors for conversations. This is a place to raise any questions, comments or other feedback that you may not have had the opportunity to say yet. We're looking for general feedback on the tool and your experience using it, anything that went well or anything that may have been challenging. Would anyone like to start?

Conduct unstructured group discussion if participants have any feedback, asking follow-up and clarification questions where relevant.

Proceed to closing remarks ([01-ProtocolOverview.pdf](#)).

² Cruz-Sandoval et al., "PODER: A Robot Programming Framework for Further Inclusion of People with Mild Cognitive Impairment in HRI Research," *Intl. Conf. Human-Robot Interaction, Melbourne, Australia, 2025*. <https://dl.acm-org.uml.idm.oclc.org/dol/10.5555/3721488.3721563>.

References

- [1] Gopika Ajaykumar, Maia Stiber, and Chien-Ming Huang. 2021. Designing user-centric programming aids for kinesthetic teaching of collaborative robots. *Robotics and Autonomous Systems* 145, (November 2021), 103845. <https://doi.org/10.1016/j.robot.2021.103845>
- [2] Margherita Andrao, Federica Gini, Francesco Greco, Alessandro Cappelletti, Giuseppe Desolda, Barbara Treccani, and Massimo Zancanaro. 2025. “React”, “Command”, or “Instruct”? Teachers Mental Models on End-User Development. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*, April 25, 2025. Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3706598.3713234>
- [3] Agnes Axelsson, Hendrik Buschmeier, and Gabriel Skantze. 2022. Modeling Feedback in Interaction With Conversational Agents—A Review. *Front. Comput. Sci.* 4, (March 2022). <https://doi.org/10.3389/fcomp.2022.744574>
- [4] Minja Axelsson, Raquel Oliveira, Mattia Racca, and Ville Kyrki. 2021. Social Robot Co-Design Canvases: A Participatory Design Framework. *J. Hum.-Robot Interact.* 11, 1 (October 2021), 3:1-3:39. <https://doi.org/10.1145/3472225>
- [5] Tewodros W. Ayalew, Jennifer Wang, Michael L. Littman, Blase Ur, and Sarah Sebo. 2025. Enabling End Users to Program Robots Using Reinforcement Learning. In *Proceedings of the 2025 ACM/IEEE International Conference on Human-Robot Interaction (HRI '25)*, March 04, 2025. IEEE Press, Melbourne, Australia, 767–777.
- [6] Rubén Baena-Perez, Iván Ruiz-Rube, Juan Manuel Dodero, and Miguel Angel Bolivar. 2020. A Framework to Create Conversational Agents for the Development of Video Games by End-Users. In *Optimization and Learning*, 2020. Springer International Publishing, Cham, 216–226. https://doi.org/10.1007/978-3-030-41913-4_18

- [7] Gagan Bansal, Besmira Nushi, Ece Kamar, Walter S. Lasecki, Daniel S. Weld, and Eric Horvitz. 2019. Beyond Accuracy: The Role of Mental Models in Human-AI Team Performance. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing* 7, (October 2019), 2–11. <https://doi.org/10.1609/hcomp.v7i1.5285>
- [8] Christine M. Barber, Robin J. Shucksmith, Bruce MacDonald, and Burkhard C. Wünsche. 2010. Sketch-based robot programming. In *2010 25th International Conference of Image and Vision Computing New Zealand*, November 2010. 1–8. <https://doi.org/10.1109/IVCNZ.2010.6148868>
- [9] David Bau, Jeff Gray, Caitlin Kelleher, Josh Sheldon, and Franklyn Turbak. 2017. Learnable programming: blocks and beyond. *Commun. ACM* 60, 6 (May 2017), 72–80. <https://doi.org/10.1145/3015455>
- [10] James M. Berzuk and James E. Young. 2022. More Than Words: A Framework for Describing Human-Robot Dialog Designs. In *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction (HRI '22)*, March 07, 2022. IEEE Press, Sapporo, Hokkaido, Japan, 393–401.
- [11] Narjes Bessghaier, Mohammed Sayagh, Ali Ouni, and Mohamed Wiem Mkaouer. 2023. What Constitutes the Deployment and Runtime Configuration System? An Empirical Study on OpenStack Projects. *ACM Trans. Softw. Eng. Methodol.* 33, 1 (November 2023), 5:1-5:37. <https://doi.org/10.1145/3607186>
- [12] Elin A. Björling, Emma Rose, Andrew Davidson, Rachel Ren, and Dorothy Wong. 2020. Can We Keep Him Forever? Teens’ Engagement and Desire for Emotional Connection with a Social Robot. *Int J of Soc Robotics* 12, 1 (January 2020), 65–77. <https://doi.org/10.1007/s12369-019-00539-6>
- [13] Elin A. Björling, Wendy M. Xu, Maria E. Cabrera, and Maya Cakmak. 2019. The Effect of Interaction and Design Participation on Teenagers’ Attitudes towards Social Robots. In *2019 28th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, October 2019. 1–7. <https://doi.org/10.1109/RO-MAN46459.2019.8956427>

- [14] Peter Blomsma, Gabriel Skantze, and Marc Swerts. 2022. Backchannel Behavior Influences the Perceived Personality of Human and Artificial Communication Partners. *Front. Artif. Intell.* 5, (March 2022). <https://doi.org/10.3389/frai.2022.835298>
- [15] Peter Blomsma, Julija Vaitonyté, Gabriel Skantze, and Marc Swerts. 2024. Backchannel behavior is idiosyncratic. *Language and Cognition* (February 2024), 1–24. <https://doi.org/10.1017/langcog.2024.1>
- [16] Christine L. Borgman. 1986. The user’s mental model of an information retrieval system: an experiment on a prototype online catalog. *International Journal of Man-Machine Studies* 24, 1 (January 1986), 47–64. [https://doi.org/10.1016/S0020-7373\(86\)80039-6](https://doi.org/10.1016/S0020-7373(86)80039-6)
- [17] Stephen Brade, Sam Anderson, Rithesh Kumar, Zeyu Jin, and Anh Truong. 2025. SpeakEasy: Enhancing Text-to-Speech Interactions for Expressive Content Creation. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI ’25)*, April 25, 2025. Association for Computing Machinery, New York, NY, USA, 1–19. <https://doi.org/10.1145/3706598.3714263>
- [18] Crystal Chao and Andrea L. Thomaz. 2013. Controlling social dynamics with a parametrized model of floor regulation. *J. Hum.-Robot Interact.* 2, 1 (February 2013), 4–29. <https://doi.org/10.5898/JHRI.2.1.Chao>
- [19] Lettie Y. Conrad and Virginia M. Tucker. 2018. Making it tangible: hybrid card sorting within qualitative interviews. *Journal of Documentation* 75, 2 (November 2018), 397–416. <https://doi.org/10.1108/JD-06-2018-0091>
- [20] Dagoberto Cruz-Sandoval, Michele Murakami, Alyssa Kubota, and Laurel D. Riek. 2025. PODER: A Robot Programming Framework to Further Inclusion of People with Mild Cognitive Impairment in HRI Research. In *Proceedings of the 2025 ACM/IEEE International Conference on Human-Robot Interaction (HRI ’25)*, March 04, 2025. IEEE Press, Melbourne, Australia, 599–609.
- [21] Amy Csizmar Dalal, Jackie Chan, and Kirby Mitchell. 2019. A Preliminary Study of the Role of Language in Home Network Troubleshooting. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems (CHI EA ’19)*, May 02, 2019.

- Association for Computing Machinery, New York, NY, USA, 1–6.
<https://doi.org/10.1145/3290607.3312856>
- [22] Ronald Cumbal, Daniel Alexander Kazzi, Vincent Winberg, and Olov Engwall. 2022. Shaping unbalanced multi-party interactions through adaptive robot backchannels. In *Proceedings of the 22nd ACM International Conference on Intelligent Virtual Agents (IVA '22)*, September 06, 2022. Association for Computing Machinery, New York, NY, USA, 1–7. <https://doi.org/10.1145/3514197.3549680>
- [23] Chandan Datta, Chandimal Jayawardena, I Han Kuo, and Bruce A MacDonald. 2012. RoboStudio: A visual programming environment for rapid authoring and customization of complex services on a personal service robot. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, October 2012. 2352–2357. <https://doi.org/10.1109/IROS.2012.6386105>
- [24] David L. Morgan. 1997. *The Focus Group Guidebook*. SAGE Publications, Inc, Thousand Oaks. Retrieved May 9, 2025 from <http://uml.idm.oclc.org/login?url=https://search.ebscohost.com/login.aspx?direct=true&db=e000xna&AN=478002&site=ehost-live>
- [25] Swaib Dragule, Thorsten Berger, Claudio Menghi, and Patrizio Pelliccione. 2021. A survey on the design space of end-user-oriented languages for specifying robotic missions. *Softw Syst Model* 20, 4 (August 2021), 1123–1158. <https://doi.org/10.1007/s10270-020-00854-x>
- [26] Saad Elbeleidy, Michelle Lieng, and Tom Williams. 2024. A Robot Dialogue Authoring Interface with Smart Capabilities. In *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI '24)*, March 11, 2024. Association for Computing Machinery, New York, NY, USA, 442–446. <https://doi.org/10.1145/3610978.3640555>
- [27] Olov Engwall, Ronald Cumbal, and Ali Reza Majlesi. 2023. Socio-cultural perception of robot backchannels. *Front. Robot. AI* 10, (January 2023). <https://doi.org/10.3389/frobt.2023.988042>

- [28] Carol Figueroa, Marcel De Korte, Magalie Ochs, and Gabriel Skantze. 2024. Mhm... Yeah? Okay! Evaluating the Naturalness and Communicative Function of Synthesized Feedback Responses in Spoken Dialogue. In *Proceedings of the 25th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, 2024. Association for Computational Linguistics, Kyoto, Japan, 544–553. <https://doi.org/10.18653/v1/2024.sigdial-1.46>
- [29] Gerhard Fischer. 1991. The Importance of Models in Making Complex Systems Comprehensible. In *Human Factors in Information Technology*, M. J. Tauber and D. Ackermann (eds.). North-Holland, 3–36. <https://doi.org/10.1016/B978-0-444-88602-6.50005-0>
- [30] Gerhard Fischer. 2001. User Modeling in Human–Computer Interaction. *User Modeling and User-Adapted Interaction* 11, 1 (March 2001), 65–86. <https://doi.org/10.1023/A:1011145532042>
- [31] Gerhard Fischer and Antonio Piccinno. 2017. Revisiting and Broadening the Meta-Design Framework for End-User Development. In *New Perspectives in End-User Development*, Fabio Paternò and Volker Wulf (eds.). Springer International Publishing, Cham, 61–97. https://doi.org/10.1007/978-3-319-60291-2_4
- [32] Daniela Fogli, Luigi Gargioni, Giovanni Guida, and Fabio Tampalini. 2022. A hybrid approach to user-oriented programming of collaborative robots. *Robotics and Computer-Integrated Manufacturing* 73, (February 2022), 102234. <https://doi.org/10.1016/j.rcim.2021.102234>
- [33] Christopher Frauenberger, Judith Good, Geraldine Fitzpatrick, and Ole Sejer Iversen. 2015. In pursuit of rigour and accountability in participatory design. *International Journal of Human-Computer Studies* 74, (February 2015), 93–106. <https://doi.org/10.1016/j.ijhcs.2014.09.004>
- [34] Helena Anna Frijns, Darja Stoeva, Oliver Schürer, and Margrit Gelautz. 2024. Programming Robot Animation Through Human Body Movement. In *2024 IEEE International Conference on Advanced Robotics and Its Social Impacts (ARSO)*, May 2024. 273–279. <https://doi.org/10.1109/ARSO60199.2024.10557907>

- [35] Luigi Gargioni and Daniela Fogli. 2024. Integrating ChatGPT with Blockly for End-User Development of Robot Tasks. In *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI '24)*, March 11, 2024. Association for Computing Machinery, New York, NY, USA, 478–482. <https://doi.org/10.1145/3610978.3640653>
- [36] Cristina Gena, Alberto Lillo, Claudio Mattutino, and Enrico Mosca. 2022. Wolly: an affective and adaptive educational robot. In *Adjunct Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization*, July 04, 2022. ACM, Barcelona Spain, 146–150. <https://doi.org/10.1145/3511047.3537684>
- [37] Javi F. Gorostiza and Miguel A. Salichs. 2011. End-user programming of a social robot by dialog. *Robotics and Autonomous Systems* 59, 12 (December 2011), 1102–1114. <https://doi.org/10.1016/j.robot.2011.07.009>
- [38] Saul Greenberg. 2007. Toolkits and interface creativity. *Multimed Tools Appl* 32, 2 (February 2007), 139–159. <https://doi.org/10.1007/s11042-006-0062-y>
- [39] Saul Greenberg and Chester Fitchett. 2001. Phidgets: easy development of physical interfaces through physical widgets. In *Proceedings of the 14th annual ACM symposium on User interface software and technology (UIST '01)*, November 11, 2001. Association for Computing Machinery, New York, NY, USA, 209–218. <https://doi.org/10.1145/502348.502388>
- [40] Samantha Harper and Stephen Dorton. 2019. A Context-Driven Framework for Selecting Mental Model Elicitation Methods. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 63, 1 (November 2019), 367–371. <https://doi.org/10.1177/1071181319631422>
- [41] Paul Hatch and Joseph Michaelis. 2023. “Wait, Wait, Go Back”: Investigating Social Supports for Homework During Do-design Sessions with Teachers. *Annual Meeting of the International Society of the Learning Sciences (ISLS)* (June 2023). Retrieved September 30, 2025 from <https://par.nsf.gov/biblio/10435463-wait-wait-go-back-investigating-social-supports-homework-during-do-design-sessions-teachers>

- [42] Paul Hatch, Md Anisur Rahman, and Joseph E. Michaelis. 2023. PATHWiSE: An Authoring Tool to Support Teachers to Create Robot-Supported Social Learning Experiences During Homework. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW1 (April 2023), 144:1-144:23. <https://doi.org/10.1145/3579620>
- [43] Jennifer Morey Hawkins. 2017. Thematic Analysis. In *The SAGE Encyclopedia of Communication Research Methods*. SAGE Publications, Inc, 1757–1760. <https://doi.org/10.4135/9781483381411>
- [44] Guy Hoffman, Gurit E. Birnbaum, Keinan Vanunu, Omri Sass, and Harry T. Reis. 2014. Robot responsiveness to human disclosure affects social impression and appeal. In *Proceedings of the 2014 ACM/IEEE International Conference on Human-Robot Interaction (HRI '14)*, 2014. Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/2559636.2559660>
- [45] Robert Holwerda and Felienne Hermans. 2018. A Usability Analysis of Blocks-based Programming Editors using Cognitive Dimensions. In *2018 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*, October 2018. 217–225. <https://doi.org/10.1109/VLHCC.2018.8506483>
- [46] Eric Horvitz. 1999. Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems (CHI '99)*, May 01, 1999. Association for Computing Machinery, New York, NY, USA, 159–166. <https://doi.org/10.1145/302979.303030>
- [47] Erzhen Hu, Yanhe Chen, Mingyi Li, Vrushank Phadnis, Pingmei Xu, Xun Qian, Alex Olwal, David Kim, Seongkook Heo, and Ruofei Du. 2025. DialogLab: Authoring, Simulating, and Testing Dynamic Human-AI Group Conversations. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology*, September 28, 2025. ACM, Busan Republic of Korea, 1–20. <https://doi.org/10.1145/3746059.3747696>
- [48] Soshi Iba, Christiaan J. J. Paredis, and Pradeep K. Khosla. 2005. Interactive Multimodal Robot Programming. *The International Journal of Robotics Research* 24, 1 (January 2005), 83–104. <https://doi.org/10.1177/0278364904049250>

- [49] Takeo Igarashi, Satoshi Matsuoka, and Hidehiko Tanaka. 1999. Teddy: a sketching interface for 3D freeform design. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques (SIGGRAPH '99)*, July 01, 1999. ACM Press/Addison-Wesley Publishing Co., USA, 409–416. <https://doi.org/10.1145/311535.311602>
- [50] Carlos T. Ishi, ChaoRan Liu, Hiroshi Ishiguro, and Norihiro Hagita. 2010. Head motions during dialogue speech and nod timing control in humanoid robots. In *Proceedings of the 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI '10)*, 2010. IEEE Press, Osaka, Japan, 293–300.
- [51] John M. Carroll, Judith Reitman Olson. 1988. Mental Models in Human-Computer Interaction. In *Handbook of Human-Computer Interaction*. North-Holland, 45–65. <https://doi.org/10.1016/B978-0-444-70536-5.50007-5>
- [52] Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. 2017. *Research Methods in Human-Computer Interaction*. Morgan Kaufmann, Cambridge, MA. Retrieved May 10, 2025 from <http://uml.idm.oclc.org/login?url=https://search.ebscohost.com/login.aspx?direct=true&db=e000xna&AN=1158797&site=ehost-live>
- [53] Sabrina Lakhdhir, Chehak Nayar, Fraser Anderson, Helene Fournier, Liisa Holsti, Irina Kondratova, Charles Perin, and Sowmya Somanath. 2024. GlucoMaker: Enabling Collaborative Customization of Glucose Monitors. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*, May 11, 2024. Association for Computing Machinery, New York, NY, USA, 1–21. <https://doi.org/10.1145/3613904.3642435>
- [54] Divesh Lala, Koji Inoue, Tatsuya Kawahara, and Kei Sawada. 2022. Backchannel Generation Model for a Third Party Listener Agent. In *Proceedings of the 10th International Conference on Human-Agent Interaction (HAI '22)*, December 05, 2022. Association for Computing Machinery, New York, NY, USA, 114–122. <https://doi.org/10.1145/3527188.3561926>
- [55] Divesh Lala, Pierrick Milhorat, Koji Inoue, Masanari Ishida, Katsuya Takanashi, and Tatsuya Kawahara. 2017. Attentive listening system with backchanneling, response generation and

- flexible turn-taking. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, August 2017. Association for Computational Linguistics, Saarbrücken, Germany, 127–136. <https://doi.org/10.18653/v1/W17-5516>
- [56] Jin Joo Lee, Fei Sha, and Cynthia Breazeal. 2020. A Bayesian theory of mind approach to nonverbal communication. In *Proceedings of the 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI '19)*, 2020. IEEE Press, Daegu, Republic of Korea, 487–496.
- [57] Sau-lai Lee, Ivy Yee-man Lau, S. Kiesler, and Chi-Yue Chiu. 2005. Human Mental Models of Humanoid Robots. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, April 2005. 2767–2772. <https://doi.org/10.1109/ROBOT.2005.1570532>
- [58] Stephen C. Levinson and Francisco Torreira. 2015. Timing in turn-taking and its implications for processing models of language. *Front. Psychol.* 6, (June 2015). <https://doi.org/10.3389/fpsyg.2015.00731>
- [59] Yifang Li, Nishant Vishwamitra, Hongxin Hu, and Kelly Caine. 2020. Towards A Taxonomy of Content Sensitivity and Sharing Preferences for Photos. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*, April 23, 2020. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3313831.3376498>
- [60] Henry Lieberman, Fabio Paternò, Markus Klann, and Volker Wulf. 2006. End-User Development: An Emerging Paradigm. In *End User Development*, Henry Lieberman, Fabio Paternò and Volker Wulf (eds.). Springer Netherlands, Dordrecht, 1–8. https://doi.org/10.1007/1-4020-5386-X_1
- [61] Chaoran Liu, Carlos T. Ishi, Hiroshi Ishiguro, and Norihiro Hagita. 2012. Generation of nodding, head tilting and eye gazing for human-robot dialogue interaction. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI '12)*, 2012. Association for Computing Machinery, New York, NY, USA, 285–292. <https://doi.org/10.1145/2157689.2157797>

- [62] Lorenz Cuno Klopfenstein. 2019. Loquacious Blockly: a proposal for conversational interface building blocks. *Chatbot Research and Design – CONVERSATIONS* (2019).
- [63] Monica G. Maceli. 2017. Tools of the Trade: A Survey of Technologies in End-User Development Literature. In *End-User Development*, 2017. Springer International Publishing, Cham, 49–65. https://doi.org/10.1007/978-3-319-58735-6_4
- [64] Karthik Mahadevan, Blaine Lewis, Jiannan Li, Bilge Mutlu, Anthony Tang, and Tovi Grossman. 2025. ImageInThat: Manipulating Images to Convey User Instructions to Robots. In *Proceedings of the 2025 ACM/IEEE International Conference on Human-Robot Interaction (HRI '25)*, March 04, 2025. IEEE Press, Melbourne, Australia, 757–766.
- [65] Nicolai Marquardt, Robert Diaz-Marino, Sebastian Boring, and Saul Greenberg. 2011. The proximity toolkit: prototyping proxemic interactions in ubiquitous computing ecologies. In *Proceedings of the 24th annual ACM symposium on User interface software and technology (UIST '11)*, October 16, 2011. Association for Computing Machinery, New York, NY, USA, 315–326. <https://doi.org/10.1145/2047196.2047238>
- [66] Damien Masson, Sylvain Malacria, G ry Casiez, and Daniel Vogel. 2024. DirectGPT: A Direct Manipulation Interface to Interact with Large Language Models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*, May 11, 2024. Association for Computing Machinery, New York, NY, USA, 1–16. <https://doi.org/10.1145/3613904.3642462>
- [67] Louis McCallum and Peter W. McOwan. 2015. Face the Music and Glance: How Nonverbal Behaviour Aids Human Robot Relationships Based in Music. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI '15)*, March 02, 2015. Association for Computing Machinery, New York, NY, USA, 237–244. <https://doi.org/10.1145/2696454.2696477>
- [68] Mona A. Mohamed, Joyram Chakraborty, and Josh Dehlinger. 2017. Trading off usability and security in user interface design through mental models. *Behaviour & Information Technology* 36, 5 (May 2017), 493–516. <https://doi.org/10.1080/0144929X.2016.1262897>

- [69] Masahiro Mori, Karl F. MacDorman, and Norri Kageki. 2012. The Uncanny Valley [From the Field]. *IEEE Robotics & Automation Magazine* 19, 2 (June 2012), 98–100. <https://doi.org/10.1109/MRA.2012.2192811>
- [70] Terran Mott, Alexandra Bejarano, and Tom Williams. 2022. Robot Co-design Can Help Us Engage Child Stakeholders in Ethical Reflection. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, March 2022. 14–23. <https://doi.org/10.1109/HRI53351.2022.9889430>
- [71] Michael Murray, Nick Walker, Amal Nanavati, Patricia Alves-Oliveira, Nikita Filippov, Allison Sauppe, Bilge Mutlu, and Maya Cakmak. 2022. Learning Backchanneling Behaviors for a Social Robot via Data Augmentation from Human-Human Conversations. In *Proceedings of the 5th Conference on Robot Learning*, January 11, 2022. PMLR, 513–525. Retrieved November 10, 2025 from <https://proceedings.mlr.press/v164/murray22a.html>
- [72] Alfred Mutanga. MENTAL MODELS OF ROBOTIC PROGRAMMING.
- [73] Bilge Mutlu, Toshiyuki Shiwa, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. 2009. Footing in human-robot conversations: how robots might shape participant roles using gaze cues. In *Proceedings of the 4th ACM/IEEE International Conference on Human Robot Interaction (HRI '09)*, 2009. Association for Computing Machinery, New York, NY, USA, 61–68. <https://doi.org/10.1145/1514095.1514109>
- [74] Don Norman. 2013. *The Design of Everyday Things: Revised and Expanded Edition*. Basic Books.
- [75] Donald A. Norman. 1983. Some Observations on Mental Models. In *Mental Models*. Psychology Press.
- [76] Hae Won Park, Mirko Gelsomini, Jin Joo Lee, and Cynthia Breazeal. 2017. Telling Stories to Robots: The Effect of Backchanneling on a Child’s Storytelling. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, March 06, 2017. ACM, Vienna Austria, 100–108. <https://doi.org/10.1145/2909824.3020245>

- [77] David J. Porfirio, Laura Stegner, Maya Cakmak, Allison Sauppé, Aws Albarghouthi, and Bilge Mutlu. 2021. Figaro: A Tabletop Authoring Environment for Human-Robot Interaction. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, May 06, 2021. ACM, Yokohama Japan, 1–15. <https://doi.org/10.1145/3411764.3446864>
- [78] Aditi Ramachandran, Chien-Ming Huang, Edward Gartland, and Brian Scassellati. 2018. Thinking Aloud with a Tutoring Robot to Enhance Learning. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction (HRI '18)*, 2018. Association for Computing Machinery, New York, NY, USA, 59–68. <https://doi.org/10.1145/3171221.3171250>
- [79] Preeti Ramaraj. 2021. Robots that Help Humans Build Better Mental Models of Robots. In *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction (HRI '21 Companion)*, March 08, 2021. Association for Computing Machinery, New York, NY, USA, 595–597. <https://doi.org/10.1145/3434074.3446365>
- [80] Daniel J. Rea, Takeo Igarashi, and James E. Young. 2014. PaintBoard: prototyping interactive character behaviors by digitally painting storyboards. In *Proceedings of the second international conference on Human-agent interaction*, October 29, 2014. ACM, Tsukuba Japan, 315–322. <https://doi.org/10.1145/2658861.2658886>
- [81] Luis Rodríguez-Gil, Javier García-Zubia, Pablo Orduña, Aitor Villar-Martinez, and Diego López-De-Ipiña. 2019. New Approach for Conversational Agent Definition by Non-Programmers: A Visual Domain-Specific Language. *IEEE Access* 7, (2019), 5262–5276. <https://doi.org/10.1109/ACCESS.2018.2883500>
- [82] Matthew Rueben, Jeffrey Klow, Madelyn Duer, Eric Zimmerman, Jennifer Piacentini, Madison Browning, Frank J. Bernieri, Cindy M. Grimm, and William D. Smart. 2021. Mental Models of a Mobile Shoe Rack: Exploratory Findings from a Long-term In-the-Wild Study. *J. Hum.-Robot Interact.* 10, 2 (June 2021), 1–36. <https://doi.org/10.1145/3442620>
- [83] Matthew Rueben, Maja J. Matarić, Eitan Rothberg, and Matthew Tang. 2020. Estimating and Influencing User Mental Models of a Robot’s Perceptual Capabilities: Initial Development

- and Pilot Study. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI '20)*, April 01, 2020. Association for Computing Machinery, New York, NY, USA, 418–420. <https://doi.org/10.1145/3371382.3378392>
- [84] James Russell. 1980. A circumplex model of affect. *Journal of Personality and Social Psychology* 39, 6 (December 1980), 1161–1178. Retrieved January 20, 2026 from <https://oce.ovid.com/article/00005205-198012000-00015>
- [85] Harvey Sacks, Emanuel Schegloff, and Gail Jefferson. 1974. A Simplest Systematics for the Organization of Turn-Taking for Conversation. *Language* 50, 4 (December 1974), 696–735.
- [86] Daisuke Sakamoto, Koichiro Honda, Masahiko Inami, and Takeo Igarashi. 2009. Sketch and run: a stroke-based interface for home robots. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*, April 04, 2009. Association for Computing Machinery, New York, NY, USA, 197–200. <https://doi.org/10.1145/1518701.1518733>
- [87] Allison Sauppé and Bilge Mutlu. 2014. Design patterns for exploring and prototyping human-robot interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, April 26, 2014. ACM, Toronto Ontario Canada, 1439–1448. <https://doi.org/10.1145/2556288.2557057>
- [88] Mazyar Seraj, Eva-Sophie Katterfeldt, Kerstin Bub, Serge Autexier, and Rolf Drechsler. 2019. Scratch and Google Blockly: How Girls' Programming Skills and Attitudes are Influenced. In *Proceedings of the 19th Koli Calling International Conference on Computing Education Research (Koli Calling '19)*, November 21, 2019. Association for Computing Machinery, New York, NY, USA, 1–10. <https://doi.org/10.1145/3364510.3364515>
- [89] Jesper Simonsen and Toni Robertson (Eds.). 2012. *Routledge International Handbook of Participatory Design*. Routledge, London. <https://doi.org/10.4324/9780203108543>
- [90] Gabriel Skantze. 2021. Turn-taking in Conversational Systems and Human-Robot Interaction: A Review. *Computer Speech & Language* 67, (May 2021), 101178. <https://doi.org/10.1016/j.csl.2020.101178>

- [91] Gabriel Skantze and Bahar Irfan. 2025. Applying General Turn-taking Models to Conversational Human-Robot Interaction. In *Proceedings of the 2025 ACM/IEEE International Conference on Human-Robot Interaction (HRI '25)*, March 04, 2025. IEEE Press, Melbourne, Australia, 859–868.
- [92] Jennifer Smith, Aaron Spaulding, Harry Bratt, Dimitra Vergyri, Girish Acharya, Kristin Precoda, Andreas Kathol, and Colleen Richey. 2022. Towards Conversationally Intelligent Dialog Systems. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, April 27, 2022. ACM, New Orleans LA USA, 1–7. <https://doi.org/10.1145/3491101.3519842>
- [93] Brian Still and Kate Crane. 2017. *Fundamentals of User-Centered Design: A Practical Approach*. CRC Press, Boca Raton. <https://doi.org/10.4324/9781315200927>
- [94] Ivana Storjak, Ana Sovic Krzic, and Tomislav Jagust. 2022. Elementary School Pupils' Mental Models Regarding Robots and Programming. *IEEE Transactions on Education* 65, 3 (August 2022), 297–308. <https://doi.org/10.1109/TE.2022.3158472>
- [95] Sarah Strohkorb Sebo, Ling Liang Dong, Nicholas Chang, and Brian Scassellati. 2020. Strategies for the Inclusion of Human Members within Human-Robot Teams. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, March 09, 2020. ACM, Cambridge United Kingdom, 309–317. <https://doi.org/10.1145/3319502.3374808>
- [96] Dag Sverre Syrdal, Kerstin Dautenhahn, Kheng Lee Koay, Michael L. Walters, and Nuno R. Otero. 2010. Exploring human mental models of robots through explication interviews. In *19th International Symposium in Robot and Human Interactive Communication*, September 2010. 638–645. <https://doi.org/10.1109/ROMAN.2010.5598688>
- [97] Paige Tuttösi, Shivam Mehta, Zachary Syvenky, Bermet Burkanova, Gustav Henter, and Angelica Lim. 2025. *EmojiVoice: Towards long-term controllable expressivity in robot speech*. <https://doi.org/10.48550/arXiv.2506.15085>

- [98] Zining Wang, Paul Reisert, Eric Nichols, and Randy Gomez. 2024. Ain't Misbehavin' -- Using LLMs to Generate Expressive Robot Behavior in Conversations with the Tabletop Robot Haru. <https://doi.org/10.48550/arXiv.2402.11571>
- [99] Michel Weike, Kai Ruske, Reinhard Gerndt, and Tobias Doernbach. 2024. Enabling Untrained Users to Shape Real-World Robot Behavior Using an Intuitive Visual Programming Tool in Human-Robot Interaction Scenarios. In *Proceedings of the 2024 International Symposium on Technological Advances in Human-Robot Interaction (TAHRI '24)*, May 09, 2024. Association for Computing Machinery, New York, NY, USA, 38–46. <https://doi.org/10.1145/3648536.3648541>
- [100] David Weintrop, Afsoon Afzal, Jean Salac, Patrick Francis, Boyang Li, David C. Shepherd, and Diana Franklin. 2018. Evaluating CoBlox: A Comparative Study of Robotics Programming Environments for Adult Novices. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*, April 21, 2018. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3173574.3173940>
- [101] David Weintrop and Uri Wilensky. 2015. To block or not to block, that is the question: students' perceptions of blocks-based programming. In *Proceedings of the 14th International Conference on Interaction Design and Children (IDC '15)*, June 21, 2015. Association for Computing Machinery, New York, NY, USA, 199–208. <https://doi.org/10.1145/2771839.2771860>
- [102] David Weintrop and Uri Wilensky. 2017. Comparing Block-Based and Text-Based Programming in High School Computer Science Classrooms. *ACM Trans. Comput. Educ.* 18, 1 (October 2017), 3:1-3:25. <https://doi.org/10.1145/3089799>
- [103] Sheida White. 1989. Backchannels across Cultures: A Study of Americans and Japanese. *Language in Society* 18, 1 (1989), 59–76. Retrieved January 18, 2026 from <https://www.jstor.org/stable/4168001>
- [104] James E. Young, Takeo Igarashi, and Ehud Sharlin. 2008. Puppet Master: designing reactive character behavior by demonstration. In *Proceedings of the 2008 ACM*

- SIGGRAPH/Eurographics Symposium on Computer Animation (SCA '08)*, July 07, 2008. Eurographics Association, Goslar, DEU, 183–191. Retrieved March 11, 2026 from <https://dl.acm.org/doi/10.5555/1632592.1632619>
- [105] James Young, Kentaro Ishii, Takeo Igarashi, and Ehud Sharlin. 2012. Style by demonstration: teaching interactive movement style to robots. In *Proceedings of the 2012 ACM international conference on Intelligent User Interfaces*, February 14, 2012. ACM, Lisbon Portugal, 41–50. <https://doi.org/10.1145/2166966.2166976>
- [106] Alex Wuqi Zhang, Rafael Queiroz, and Sarah Sebo. 2025. Balancing User Control and Perceived Robot Social Agency through the Design of End-User Robot Programming Interfaces. In *Proceedings of the 2025 ACM/IEEE International Conference on Human-Robot Interaction (HRI '25)*, March 04, 2025. IEEE Press, Melbourne, Australia, 899–908. Retrieved March 17, 2025 from <https://dl-acm-org.uml.idm.oclc.org/doi/10.5555/3721488.3721598>
- [107] T. Zimmermann. 2016. Card-sorting: From text to themes. In *Perspectives on Data Science for Software Engineering*, Tim Menzies, Laurie Williams and Thomas Zimmermann (eds.). Morgan Kaufmann, Boston, 137–141. <https://doi.org/10.1016/B978-0-12-804206-9.00027-1>
- [108] Blockly. *Google for Developers*. Retrieved February 7, 2025 from <https://developers.google.com/blockly>
- [109] Card Sorting: Uncover Users' Mental Models. *Nielsen Norman Group*. Retrieved September 26, 2025 from <https://www.nngroup.com/articles/card-sorting-definition/>