Designing Social Interactions with a Humorous Robot Photographer

Timothy Adamson, C. Burton Lyng-Olsen, Kendrick Umstattd, and Marynel Vázquez
Computer Science Department, Yale University
New Haven, CT
tim.adamson@yale.edu

(a) (b) (c) (d) (e)

Figure 1: A person interacting with our robot photographer. The robot first offers to take a photo (a). During the countdown (b), it displays humorous content on its screen (c) to induce spontaneous smiles (d). Finally, it shows the subject's picture (e).

ABSTRACT

This paper describes our efforts to explore the design space of social interactions for a robot portrait photographer. Our human-centered design process involved professional and amateur photographers to better understand the social dimensions of subject-photographer interactions. This exploration then guided our design of a robot photographer, which employs humor to elicit spontaneous smiles during photography events. In a laboratory evaluation of our robot prototype, we found that the majority of the subjects considered the robot's humor to be comical and appreciated it. More spontaneous smiles were elicited by the robot when it delivered humorous content to its subjects than when it was not humorous. Our findings provide insights for the design of future social robot photographers.

CCS CONCEPTS

• **Human-centered computing → Collaborative and social computing devices**.

KEYWORDS

Human-Robot Interaction, Photography, Design, Humor

ACM Reference Format:

Timothy Adamson, C. Burton Lyng-Olsen, Kendrick Umstattd, and Marynel Vázquez. 2020. Designing Social Interactions with a Humorous Robot Photographer. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI '20), March 23–26, 2020, Cambridge, United Kingdom.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3319502.3374809

HRI '20, March 23–26, 2020, Cambridge, United Kingdom

ACM ISBN 978-1-4503-6746-2/20/03…$15.00
https://doi.org/10.1145/3319502.3374809

1 INTRODUCTION

Portrait photography is popular but complex. It brings together the lasting character of pictures with the dynamic nature of humans. While other areas of photography, such as landscape photography, seek to capture scenes that remain static for weeks, portrait photography captures people, who within seconds can change in countenance, pose, and body language. Further, portrait photography goes beyond the capturing of an image to include the nuanced interaction between the photographer and its subjects [11, 24, 37].

The proliferation of digital cameras and advancements in robotics have motivated the development of robot photographers in the past. In general, most prior research efforts in robot photography have focused on the problem of taking well composed pictures of people [6, 7, 34, 43]. Current commercial robot photographers can take portrait pictures in weddings [33] and other social events [23]. However, an important gap still exists on understanding human-robot interactions in photography contexts. What aspects of these interactions are most relevant for robot photographers? How should they engage with people when taking portrait photographs?

We followed a human-centered design process to better understand the design space of social interactions for a robot photographer. In particular, we explore the potential of creating a *humorous robot photographer that takes portrait pictures*. Our design builds on insights from both professional and amateur human photographers. It explores the potential of humor to trigger spontaneous smiles within the photography context in Human-Robot Interaction (HRI). Spontaneous smiles are important because they convey happiness and an inner glow that is often valued in photos [8, 18].

We evaluated our prototype of a portrait photographer in a controlled environment. Participants interacted with the robot on an individual basis, taking multiple pictures in a variety of positions as illustrated in Figure 1. Many discoveries were made during the evaluation and the design process, such as the complex interaction between posed and spontaneous smiles that can be elicited by a robot photographer. Based on these discoveries, we provide design recommendations for creating future robot portrait photographers.

2 RELATED WORK

Robot Photography. Most work on robot photography has focused on automatically capturing pictures based on well-established compositional rules, including the "rule of thirds" and the "no middle rule" [1, 6, 43]. In drone photography, prior efforts have proposed methods for computing flight trajectories that trade off shot smoothness, occlusion, cinematography guidelines, and motion dynamics [5, 26]. In more traditional social events, robot photographers have leveraged mobile navigation to achieve better portraits of people [1, 6, 7, 20, 43]. A common evaluation metric in this line of work are picture ratings by third-party evaluators [1, 6, 43].

Notably, Byers, Smart, Grimm and colleagues deployed their robot photographer at the SIGGRAPH 2002 conference, at a wedding event, and at other social gatherings [7, 34]. The robot successfully detected faces and avoided collisions with humans, although it was limited to a photo flash and camera click as the only means of externalizing itself. After observing people's reactions and behaviors in proximity to the robot, the authors suggested that autonomous robot photographers would benefit from a bimodal interaction system. Such a system would enable the robot to blend into the background for candid photo captures in crowds, and appropriately engage with humans when the opportunity arises. Worth noting, [34] provides system design recommendations for robot photographers. These recommendations include adaptability to the given social environment, and software modularity to facilitate customization. We consider these recommendations in the present work.

Commercial solutions exist today for robot photography [23, 33]. However, little is understood about how human-robot interactions should unfold in a photography context. Our work aims to advance our understanding of this design space.

Human Photography. The portrait photography literature provides insights into salient elements of the interactions between subjects and human photographers that can inspire social robot photography. For example, photography textbooks emphasize the importance of enlivening [17], engaging [16], and encouraging subjects [42] during photography events. In particular, B. Hurter [17] shares in his book about the importance of the photographer being an enlivening presence in the interaction. He provides strategies on how to enliven the photography experience, all with the intent of the subject having a good time. For instance, one strategy is "outrageous flattery and corn-ball humor." Furthermore, he argues that it is important to engage photography subjects through conversation and eye contact [16]. A subject's willingness to make eye contact can be used to gauge their social comfort. If they seem uncomfortable, one solution is getting the subject to share more about themselves in conversation. Photographers are also encouraged to identify the traits of their subjects which they genuinely appreciate, and then mention those traits in an affirming way [42]. This strategy can strengthen the photographer-subject relationship and help the subject feel more at ease.

Portrait Photography is complex because it involves social interactions, which generally have many nuances. This complexity suggests that to design a successful robot photographer one should consider the social context and the interactions that take place during the photographing event. These social aspects can influence the form of the pictures, i.e., type of composition that is adequate for a given photo. Additionally, these aspects drive the content, e.g., the emotions and sentiment transmitted by a photograph. These insights are in line with findings by Byers and colleagues [7] who discuss the importance of adaptability to social environments in robot photography. They motivated our research's focus on the social aspects of portrait photography.

3 DESIGN GOAL

Based on the insights gained from the related work described before, we set out to explore the design of a social robot photographer. As in prior work, we determined that our robot should be able to capture pictures autonomously. But more importantly, it should be able to take part in the social dynamics of situated photography events and leverage those social dynamics to take better pictures.

The robot would operate in the context of an indoor, semi-public university space without strict noise constraints, e.g., the entrance of a university's publicly accessible library. Such a space would provide protection from the weather and vandalism. Also, it would allow for interactions with a diverse array of people, including campus visitors, students, staff, and faculty. As visitors come to explore the campus, the robot would serve as a local photographer who can capture pictures of them for free. It would provide them with both a fun experience and with a photo souvenir of their visit. For the regular members of the community, the robot would serve as a tool to release stress and document relevant moments of their life on campus. Our robot would thus inhabit the social role of a fun, comfortable, and helpful service supplier and entertainer.

Throughout our design process, we narrowed our goal to *creating a social robot photographer that aims to elicit spontaneous smiles while taking portrait pictures*. Spontaneous smiles, also known as Duchenne Smiles [12], are often valued by the subjects of portraits [18]. They have been shown to positively affect mental state [38] and improve one's perceived attractiveness [3]. Work in marketing suggests that photos of people with Duchenne Smiles are preferred over photos of people without them [32].

4 DESIGN PROCESS

We started our iterative design process seeking to explore the design of a social robot photographer in a university environment. By the end, we had narrowed our design to be that of a humorous robot photographer that tries to elicit spontaneous smiles in subjects. The next sections describe how we arrived at this goal through interviews and an initial formative study. Further, they describe how we adapted the behavioral capabilities of our robot based on the findings from our design process.

4.1 Informal Interviews

We started our design process by conducting informal interviews with three professional photographers. At this point, we had not narrowed our design on a humorous photographer, so we did not ask any humor-specific questions. Two of the interviews were conducted over the phone, and one was in person while engaging in an unrelated task. The photographers were informed of the robot's design context at the beginning of the interview, so that their recommendations were tailored to our specific application. The two primary questions of the interview were:

(1) *What would be the best way for us to study professional photographers in action so as to glean valuable insights that could be applied to a social robot photographer?* The interviewees suggested having photographers come to our lab to take pictures of ourselves as well as others, while filming them, analyzing their interactions, and asking questions along the way. This would allow us to ascertain subtle yet important decisions that photographers make which might otherwise be overlooked if the photographer was simply reporting on a previous photo shoot or filling out a survey.

(2) *What aspects of the photography experience are expected to be especially salient to the subjects?* The responses to this question varied. Some of the salient features that were mentioned were unrelated to the social aspects of the photo-taking experience. Though these were not our primary interest, we did still find them useful because of their importance in producing high-quality photos. The suggestions included having an interesting background that added to the picture, taking pictures in a controlled environment where the lighting was consistent, and making use of whatever camera worked best on the robot without requiring the camera to be of the highest quality. When asked if a high-definition image captured by a standard RGB-D camera would suffice at taking portrait photos, one of the photographers said "*I don't think it would be that bad*." This gave us confidence in using a machine vision camera for the robot to perceive the environment and to take pictures of subjects. The social aspects of the photography experience that the interviewees mentioned focused on two areas: what the photographer says, and what the photographer does in order to elicit genuine smiles. The elicitation of genuine smiles especially caught our attention, because one of the photographers claimed that one of the most critical parts of his job was getting his clients to feel comfortable in front of the camera, and to smile genuinely when the photo was being taken. He explained how he had worked diligently at improving this aspect of his practice and saw it as vital to his career.

4.2 Robot Hardware and Software

Because we decided to focus our research on the social dimensions of robot photography, we did not spend our efforts iterating over the robot hardware. Instead, we made purposeful hardware design decisions at the beginning of our research and used the chosen hardware to explore our interaction design space.

4.2.1 Physical Embodiment. First, we decided to use a fixed tabletop platform for our research. This choice simplified power requirements and allowed for heavy compute in comparison to mobile platforms, which are a common choice for robot photographers [1, 6, 33, 43]. Second, we determined that the robot should be able to detect and orient itself towards subject's faces in order to engage socially with them. Thus, we opted for using a Trossen Robotics WidowX Robot Arm Kit for the body. We modified the arm to add a custom-made robot head with a camera (Fig. 1). The head had a screen face, which we could use to render any kind of facial expressions or other media. The head effectively turned the arm into a social robot and provided the means to capture pictures. We called the robot *Shutter* because it serves as a robot photographer.

4.2.2 Sensing & Computational Resources. The camera on the robot was a D435i RealSense RGB-D camera for general perception and photography tasks. For verbal communication, we initially used the built in speakers and microphone on a gaming laptop. But as the complexity of the robot's software progressed, the limited computational power and USB port availability of the laptop became a constraining issue. Therefore, we switched to using a more powerful desktop computer. With this change, we also explored using a UMA-8 microphone array for audio capture and a Logitech Z200 Stereo Speakers for crisp audio output (Fig. 1e).

4.2.3 Software. Following the suggestions from [34], we implemented the robot's software with a modular design. The software was a collection of Robot Operating System (ROS) processes, also known as nodes. The node in charge of the main interaction logic commanded gaze shifts and verbal utterances for the robot based on the state of the interaction, the subject's position, and subject's verbal input. For estimating users' position, another node visually detected faces relative to the robot using the OpenCV library. For verbal communication, two nodes provided text-to-speech and speech-to-text services using Google Cloud APIs. The supplementary material provides more details about our robot's software.

4.3 Informative Study

The feedback from the interviews (Sec. 4.1) motivated us to conduct a lab study to better understand subject-photographer interactions. In the study, participants were photographed by other participants and by Shutter. The photos were taken in similar conditions to facilitate transferring of insights from the human photographers to the robot. An advantage of the controlled environment is that it allowed human photographers to focus their attention on their interpersonal interactions with their subjects, instead of worrying about environmental distractions or other factors such as lighting.

4.3.1 Participants. We recruited groups of adults for two study sessions identified as A and B. In session A, one professional photographer participated along with three lab members who were not involved in the project. All the participants were male. The age of the professional photographer was 40, while the lab members were 19, 20, and 21 years old. In Session B, there was one female and one male participant, recruited from the local community. They were 21 and 20 years old, respectively. They reported having intermediate experience with robots (M=4, SE=1.0) and photography (M=3.5, SE = 0.5) on 7-pt Likert items. Session B allowed us to investigate interactions between amateur photographers who take pictures for fun and passion. It helped us discern some of the key skills that professional photographers develop through their practice.

4.3.2 Procedure. The study began by consenting the participants, having them complete a demographics survey, and informing them of the robot's design context to maximize the usefulness of their suggestions for Shutter. The study then continued with a Human-Human (HH) interaction phase in which humans took photos, followed by a Human-Robot (HR) interaction phase in which an initial prototype of the robot served as the photographer. We did not counterbalance the order of the phases because we did not want to bias the way that human interactions unfolded with our robot prototype at this stage of the design process.

In the HH phase, the professional photographer took pictures of the other participants in Session A (as shown in Fig. 2), while a randomly chosen participant served as the photographer in B. In general, the photographers were limited to: using a simple digital camera; not using manual options for zoom, focus, or sensitivity; and standing in specific places to take pictures. These limitations were imposed to prevent the photographers from focusing on technical aspects of the camera when taking pictures because the robot's camera provides few manual controls. Additionally, the placement of the photographer ensured that (s)he took pictures at a social distance [13], which we expected participants to keep from the robot based on prior work in human-robot proxemics [36, 39, 40]. At the end of the HH phase, the participants completed a short survey about their experience.

In the HR phase, the participants approached the robot one at a time. When the robot detected their faces, it offered to take a photo with the prompt *"Hi there, would you like me to take your picture?"* If participants responded in an affirmative manner, the robot proceeded to count down from 10 to take the photo. The picture was then displayed in a monitor next to the robot, and it said *"Wow, I think that picture looks great! If you don't like it though, we can retake it. Do you want a retake?"* The robot then kept retaking pictures until the participants were satisfied. Finally, the monitor displayed a QR code which could be scanned by the participants to retrieve their picture from a remote server.

Once individual pictures had been taken, the robot took a group picture of the participants. Next, they completed one more survey about their experience and discussed with the experimenter their impressions of their interaction with the robot. This group discussion provided an opportunity to gather feedback collaboratively.

Figure 2: A professional photographer (left) takes a picture of a participant (right) during our informative study.

The participants were compensated at the end of the study. Lab members received the photographer's pictures as compensation, while the rest of the participants were given $5 for their time. The protocol was approved by our Institutional Review Board (IRB).

4.3.3 Findings From the Human-Human Interaction Phase. The amateur and professional photographers gave directions for the subjects to follow. The amateur photographer told the subject to *"pose"*, while the professional photographer provided more sophisticated and explicit instructions. For example, the professional said *"look to each other as if you were speaking."* Additionally, he provided specific directions for the subjects to move or rotate. These directions evidenced creativity on the part of the photographer and mastery of composition rules. In general, clear directions seemed to be interpreted as an indication that the photographer was knowledgeable.

Another key feature of this phase of the study was the use of humor while photo shooting. Although not always the product of an overt joke, subjects of the amateur and professional photographer frequently laughed throughout the interaction. The benefits of humor on the photos were twofold. First, humor provided a means of eliciting genuine and more picturesque smiles from the subjects. Second, reactive smiles and subsequent laughter made participants more comfortable and improved their body language. One participant said that *"humor could be used to relax all the participants."*

When prompted to critique their photography experience, the most common criticism was the lack of an indication of when photos were taken. The participants were aware that three photos would be taken of them per session, however, there was no audible click or signal from the camera that was used by the photographers. One participant suggested that the photographer could give reassuring feedback during the photo shoot, such as *"you're doing great."*

4.3.4 Findings From the Human-Robot Interaction Phase. Although the robot prototype was in an early stage, we were able to gain valuable feedback by testing it with participants. In general, they were satisfied with the design of the interaction, claiming that Shutter was "explicit with its instructions and easy to understand." Another participant said "I liked the step-by-step process (of the photo shoot)." Further, we asked the participants if the robot should surrender some of its control by displaying the view from its camera on a screen before the photo is taken. However, most participants disagreed. They wanted the robot to remain in control as it provided a more interesting experience than a standard camera phone.

The robot's outward presence was important to the participants. In the surveys, it was mentioned that the friendliness of the robot made it more likeable. Further, multiple participants praised the design of the robot's eyes and gaze behavior. For example, a participant said that they liked the eye-tracking because it "*made me feel the robot was present.*" Users' perception of the robot's presence is reflective on the quality of the overall interaction.

Criticism of the social behavior of the robot was primarily rooted in participants feeling uncomfortable during the interaction. The most common complaint was the length of the countdown for a photo (e.g., "*the countdown was too long and made me feel awkward*"). Prolonged silences during session B also made the experience less natural than human-human photography events. Other participants also expressed desire for the robot to control the interaction, e.g., a person wanted it to "*tell[s] us to move for better positioning.*"

When asked what strategies the robot could use to make people smile, the professional photographer suggested that the robot tell a funny joke. This proposition was interesting as past studies suggest that humorous robots are perceived as more likable [25], and that humor can be used as a way to make people more comfortable in HRI [41]. A joke telling robot may elicit genuine smiles from its subject as well as improve the subject's overall experience.

4.4 Improving the Photography Experience

We modified the way that our robot photographer interacted with users based on the insights from our informative study (Sec. 4.3). First, we added numeric landmarks in front of the robot such that it could tell its subjects explicitly where to stand during photo shooting events, as depicted in the left image of Fig. 3. This change was motivated by participants' desire for clear directions from Shutter. Second, we shortened the countdown for taking a photo because counting from 10 took too long. The robot now counted from 5. Third, we modified Shutter's behavior so that it would stand still during the photography countdown. This change was implemented to prevent confusing the subject and taking blurry pictures. More specifically, Shutter would track the participant as they moved around and center the participant's face before beginning the countdown. Even when Shutter stood still, however, Shutter's eyes continued to move to track the participant. Fourth, we added humorous content to Shutter to improve its interaction with users and enable it to elicit spontaneous smiles. We expected these smiles to increase the pictures' value [18].

We explored several types of media formats for the robot to convey humor to its subjects. None of the content was offensive; but some of the humor made fun of the robot photographer. The four media formats that we considered were:

Memes. Humor that features text overlaid as commentary on a relevant image. The first of the three memes included a picture of Liam Neeson from the movie Taken with the quote *"If there is good light, I will find it and I will shoot it"*. The second meme was an image of a photographer in the process of taking a picture with text that reads "*Taking photos for a client when he asks: 'So do you have a real job on the side?'*" The third meme was of a man making a face at the camera with the quote *"If Monday had a face... This would be it"* (Fig. 1c). The memes were displayed on the robot's screen for 4, 6, and 4 seconds respectively, replacing the face of the robot.

GIFs. Humor in the form of short, soundless clips. The first GIF showed a dog sticking its head outside a car window at high speeds. The second GIF was an edited clip of the television character Mr. Bean sticking his tongue out. The third GIF was a video of a dog stuck in a toilet. The GIFs were displayed on the robot's screen in the place of the eyes for 3, 1.5, and 2.5 seconds, respectively.

Sound. Humor expressed through a short sound effect. We considered three sound samples: a goat bleating, a distorted "chipmunk" laugh, and a child laughing. The sound effects lasted 2.1, 1.3, and 2.3 seconds, respectively.

Speech. Humor conveyed through spoken phrases. The four phrases were: *"What am I doing again"*; *"2, 4, 3. Number number number"*; *"I'm really not very funny"*; and *"Say cheese."*

5 EVALUATION

After improving the design of our robot photographer, we conducted a laboratory study to further evaluate its interactions with subjects. The protocol was approved by our local IRB.

5.1 Environmental Setup

Figures 1 and 3 show the environment in which the study was conducted. The robot was placed at a table with a monitor next to it, which was used to display photos. Behind the robot was a Logitech BRIO camera, and far to its left side was a Stereolabs ZED camera. The former camera recorded high-resolution images of users' while they interacted with the robot; the latter camera recorded auxiliary footage. Participants generally stood in front of the robot during their interactions with it. They moved between the numeric landmarks on the floor, based on the robot's directions. Meanwhile, the experimenter sat behind the robot at another computer, waiting for the photo shooting session with the robot to finish.

5.2 Procedure

After providing consent and completing a demographics survey, participants were introduced to the robot and asked to stand on the number one landmark, approximately 0.77 meters away from the robot (Fig. 3 left). When the participant's face was detected by Shutter, the robot initiated the interaction by saying, "*Hi there, I'm Shutter, a robot photographer. Would you like me to take your picture?*" When the participant agreed, the robot started to take the first of five photos. To start, Shutter said "*Great! I will count down from 5, and then take the picture. Get ready.*" The robot then counted from five to one and played the sound of a camera shutter, signaling that a picture was taken. This picture was displayed on a monitor next to Shutter for the participant to view (Photo Monitor in Fig. 3). Shutter rotated towards the monitor and back to the participant to simulate as if it was looking at the photo, and directed the user to move to the next position. For example, it said "*That picture looks pretty good! I think I'll take another! How about you move over to the number 5 now so that I can get a different angle of you.*"

The second through fifth photos were taken similarly to the first, except that (1) the participant stood at a different location for each

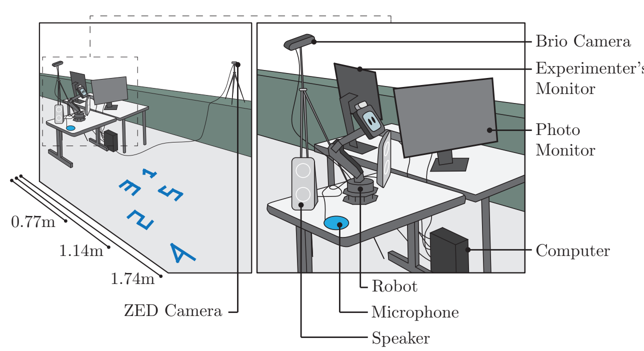

Figure 3: Left: Interaction space. The floor had 5 landmarks for subjects to position themselves during photo shooting (their distance to the robot is in the left-bottom corner). Right: Close-up view of the robot and computer peripherals.

photo based on the robot's directions, and (2) Shutter tried to elicit spontaneous smiles through humor. The same sequence of locations was used for all participants. Humorous content was conveyed after the numbers 5, 3, and 1 in the countdown. The specific content used in each case was chosen randomly from 14 options: no humor, and 13 different humorous contents of the aforementioned types (Meme, GIF, Sound or Speech). The countdown paused for each humorous content and none was repeated to the same participant. After the last photo, the robot said goodbye to the participant.

After the photo shooting session with the robot, the participant was instructed to sit at the experimenter's computer to view and label his or her reactions to the humorous content delivered by Shutter. The purpose of this labeling task was to get the participant's opinion on whether the robot induced him or her to smile spontaneously. The reactions to humor could be labeled as: *Unsure if smiling or not*, *Not smiling*, *Posed smile*, *Spontaneous smile*, and *Can't tell if smile is spontaneous or posed*. For consistency, the experimenter used a script to explain the meaning of each label.

A post-study survey and an open-ended discussion about the photography experience followed the data labeling task, the details of which can be found in the supplementary material. At the end of the study, each participant received $5 for every thirty minutes of participation. Note that the participants were not informed of the design context of Shutter until partway through the discussion because we were interested in observing how participants naturally reacted to the robot and its humorous content.

5.3 Participants

We recruited 25 participants (9 females, 16 males) through flyers and word of mouth. Participants ranged in age from 18 to 71 years old (M=26.8, SD=11.2). Most of the participants were students of diverse disciplines. Generally, they reported not having much experience with photography (M=3.0, SE=0.26) nor robots (M=2.6, SE=0.32) on 7-pt Likert responding formats.

6 RESULTS

We analyzed survey responses both quantitatively and qualitatively. We also considered transcriptions of the final discussion that took place in the study in qualitative analyses. These analyses began by inspecting the data and organizing it into two big categories: one about general aspects of the interaction and photography, and another one related to humor. For the former category, we found themes organically through an affinity diagram. Then, we further categorized ideas into positive feedback, negative feedback, and suggestions for improvement. For the humor category, we followed a similar process in general, but paid special attention to the humor format. The main groups of ideas that resulted from this process are further discussed in the following sections, along with quantitative results where appropriate. For all statistical tests discussed next, we verified their respective assumptions.

6.1 Appearance

Participants first noticed the appearance of Shutter when they started their interactions with it. The most defining characteristic of Shutter was its eyes. The large eyes along with the robot's small stature made the participants believe that the robot was approachable and "*cute*" (P10). In the final survey, they reported that the robot made them feel comfortable (M=5.72, SE=0.29) and was friendly (M=5.96, SE=0.20) on a 7-pt Likert responding format.

Based on the qualitative data, Shutter succeeded in creating a comfortable photo-taking environment for the participants. For instance, they reported that the robot was "*very polite*" (P25) and were "*pleasantly surprised by how personable it was*" (P19). Some participants assigned a personality to the robot, describing it as outgoing and friendly. While physical appearance certainly contributed to this perceived personality, other features such as the robot's behaviors and humor also played a critical role.

6.2 Verbal Behaviors

Participants generally agreed that Shutter's verbal directions were "*quick, simple, and easy to follow,*" (P23) or "*quick and to the point*" (P16). They rated Shutter's professionalism (M=5.32, SE=0.31) and timeliness (M=5.36, SE=0.33) positively on a 7-pt Likert responding format. Interestingly, Shutter's compliment for the photos that it took led to mixed impressions. Some participants found the post-picture commentary to be affirming and engaging, but others criticised it, e.g., because they found it to be "*insincere*" (P16). In the study, it was suggested that the robot "*be more honest*" (P5) and responded differently when it captured a photo with poor quality. Taken together, this feedback indicates that robot photographers should not only be capable of taking good pictures, but should also use their understanding of photo quality to acknowledge mistakes and recover from failures.

Beyond the content of Shutter's speech, participants were generally supportive of the robot's voice. We selected the default female voice from the Google Cloud Text-to-Speech API. This voice was found to be easy to understand, which helped orchestrate the interaction. Some participants found the voice to be "*comforting*" (P4) and were pleased by its feminine characteristics (P14 and P21). Others however, complained that Shutter's voice was monotonous. Participants were especially critical of the robot's lack of intonation during speech humor. Although it may be difficult for a robot photographer to not have a "robotic" (P9) voice, our findings highlight the importance of the robot's tone of voice in the interaction.

6.3 Countdown

The countdown between each photo was useful for signaling when the photos were being taken, but many participants felt that the 5 second timer was still too long. Furthermore, when the robot interrupted the countdown to deliver humorous content, some subjects found the countdown to be "*difficult to navigate*" (P22). Participants suggested that not only should the countdown duration be decreased, but also be more consistent. Their suggestions reinforce the idea that subjects value clear indications of when a photo is being taken during photography events.

6.4 Humor Perception

Research on humor has shown that a variety of factors impact humor perception, including culture [19], education [22], and gender [2], so it comes as no surprise that some participants found the humor conveyed by the robot to be extremely comical, while others did not share the same appreciation. In particular, thirteen participants expressed a positive perception of the humor, five expressed otherwise, and the rest made no specific comments. Positive perceptions from survey responses included phrases such as "*quite funny*" (P5) and "*amusing and unexpected*" (P9). In terms of the negative comments, three participants explicitly mentioned the humor as something that they disliked about their interaction with Shutter.

Mixed results were also found on the impact that humor had on participants' interaction with the robot as a whole. The many participants that expressed appreciation of the humor mentioned several ways it improved the interaction. For example, P14 expressed that the humor humanized Shutter, P21 felt that the humor put her at ease, and P3 wrote "*I liked that I was entertained and engaged throughout the session. It's hard to hold a smile for a while or to stay attentive during a traditional photography session.*" P9 even expressed that "*The jokes made me feel more comfortable*" – one of our primary goals. The participants that did not appreciate the humor mentioned that it extended the interaction and reduced the photographer's efficiency (P20). P10 wrote that "sometimes I felt too focused on the jokes." However, the other three participants who did not find the humor content funny never mentioned this aspect worsening the overall interaction with the robot. Overall, the effect of the humorous content was net positive.

6.5 Spontaneous Smile Elicitation

We used participant-labeled data to perform a chi-square test of goodness-of-fit to determine whether a spontaneous smile was equally observed when the robot delivered humorous content and when it did not during the countdown. Spontaneous smile responses were not equally distributed, $X^2(2, N = 297) = 4.5064, p < 0.05$. Humorous content preceded spontaneous smiles 40% of the time, while no humor preceded spontaneous smiles 14% of the time. The humor tended to elicit spontaneous smiles as we expected.

The amount of humor content that was found to be comical varied across participants. For some, all of it was funny; for others, none of it was. But for most, a few of the content pieces were fairly humorous. The success of the various content media at eliciting spontaneous smiles from participants can be seen in Figure 4. The post-study survey and discussion also revealed that for many participants, there was one particular medium which they enjoyed more than the rest. The most popular one was GIF.

We used a linear regression to predict a participant's rating of their overall experience base on the number of spontaneous smiles that (s)he had. After correcting for heteroskedasticity, the statistical analysis showed that a significant linear regression equation was found, F(1, 23) = 7.62, p = 0.011. The more spontaneous smiles participants showed, the higher their perception of their experience.

6.6 Photo Quality & Style

The participants reported split opinions about the quality of the photos taken by Shutter in the final survey. Thirteen participants rated the quality positively, ten rated it negatively, and two provided neutral ratings. A REML analysis on the Quality ratings with Participant ID as random effect and Gender as fixed effect resulted in significant differences for Gender (F[1, 23]=5.32, p=0.03). A Tukey-HSD post-hoc revealed that in general, females had a significantly higher

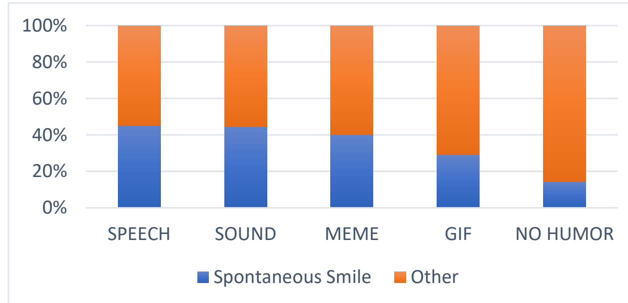

Figure 4: Shutter's success in using various humorous media formats to elicit spontaneous smiles. Y-axis corresponds to percentage of Spontaneous Smile or Other response.

perception of the quality of the photos (N=9, M=5.33, SE=0.40) in comparison to males (N=16, M=3.69, SE=0.48). The quality ratings were not significantly correlated with the participants' photography experience, how much they enjoy taking photos, nor how much they enjoy having their photo taken. However, a trend for a negative correlation was observed between the quality ratings and participants' experience with robots (p=0.052).

The main reason that participants expressed for disliking the pictures taken by shutter was the way in which they were framed or the resulting photo composition. Twelve participants suggested improving the framing of the photos. Worth noting, we generally assumed that people would want to appear smiling in the pictures taken by Shutter. However, there was one participant who expressed the contrary. The person said that though he found the jokes funny, he tried to control his smiling because he wanted a good picture, not with his mouth open. Two other people also indicated being unsure if they were supposed to smile. They originally thought that they had to be serious for the photos, such as when passport photos are taken. This feedback suggests that future versions of our robot photographer should discuss photo styles with subjects.

7 DISCUSSION

7.1 Subject-Photographer Interactions

Our evaluation showed that our prototype was in line with our design goals and confirmed several observations from our design process. First, we found humor to be a viable mechanism for our robot to elicit spontaneous smiles during photo shooting events. Second, humor, the appearance and the verbal behavior of the robot helped the participants feel more comfortable and at ease during human-robot interactions. This is important because we wanted to make our social robot photographer a fun and comfortable service supplier. Third, we found more evidence that suggested that subjects want the robot to control the interaction. This finding aligns with Smart and Grimm's suspicion that a robot photographer should drive an interaction once it engages with its subject(s) [34]. Fourth, we found a trend that suggested that the more experience users have with robots, potentially the more they expect for the quality of the photos that it takes. Lastly, our results suggest that not all photo styles are liked the same by potential subjects. Personalization of the photography experience may be beneficial.

7.2 The Complexity of Smiles

Our design process showed us that endowing robot photographers with the ability to reason about smiles is a complex but useful task. The task is complex because often times, once the countdown for a picture had started, participants felt motivated to pose smiles. They wanted their photo to convey a happy emotion. This meant that their facial expression transitioned from a posed smile to a spontaneous smile (and vice-versa). We expect these types of transitions to pose problems for existing spontaneous smile detectors, e.g., [10, 21, 27], because these systems are often built and evaluated using videos in which people transition between neutral facial expressions and spontaneous smiles. The latter type of transitions are easier to identify than those observed in our interactive setting.

Enabling humorous, robot photographers with the ability to distinguish spontaneous smiles is important for them to respond appropriately to different situations. For example, the robot could change its humor if failing to elicit spontaneous smiles. Or when a spontaneous smile is identified before the countdown finishes, the robot could save that photo as its final picture and stop counting down. This would accelerate photo shooting events.

7.3 Limitations

Our project is still in progress. Our evaluation of Shutter was limited in that we did not evaluate the pictures taken by the robot through third-party observers, as is typical within the robot photography literature [1, 6, 43]. Furthermore, we enabled our robot to take pictures under a centering composition rule, but more aesthetic compositions could be used instead, e.g., the rule of thirds. Lastly, we only considered interactions in our lab with a single subject at a time, but groups are relevant in our problem domain as well. We plan to improve Shutter in all these respects and test it in a semi-public, university environment in the future.

8 DESIGN RECOMMENDATIONS

The next paragraphs summarize key lessons from our design process and evaluation. They provide a guide to design future social robots that serve as portrait photographers:

Comfort is crucial, and can be realized using a soft voice and pleasant face. The importance of comfort was mentioned in the interviews with professional photographers, in the informative study, and throughout the evaluation of our final design. We recommend intentionally designing robot portrait photographers such that participants feel comfortable interacting with them. Though achieving this comfort requires a comprehensive approach, the robot's face and voice are significant components.

Use humor to improve the photography experience, but consider that it is not universal. People generally enjoy a humor-filled experience and prefer pictures in which they have a genuine smile. In our last study, the more spontaneous smiles participants had, the more they were satisfied with their photography experience. And because the vast majority of spontaneous smiles were the direct result of the robot conveying humor content, we recommend adding humor to robot portrait photographers. This can be done through a variety of media, as in our work. It is worth noting however that no one piece of humor will be funny to all audiences. Trying an array of humorous content or customizing the humor per subject can increase the chances of the humor being comical.

Make the photography countdown consistent and short. We shortened the countdown from 10 to 5 based on our initial design exploration, but even 5 took too long in our final evaluation. Thus, we suggest a shorter countdown, e.g., starting from 3. We also found that when humor interrupted the timing of the countdown, participants were displeased that they did not know exactly when the picture would be taken. Setting better expectations in this regard and keeping the countdown more consistent is beneficial.

Consider human-robot interaction dynamics for spontaneous smile detection. Spontaneous smile detectors are often trained with video data collected in very constrained environments. But this type of data is not necessarily representative of subject-photography interactions, in which transitions between posed smiles and spontaneous smiles are often observed. For this reason, we recommend considering data that is more representative of the photography scenario when building such detectors.

Acknowledge mistakes. When Shutter complimented participants on a poor-quality picture, the complement was generally perceived as being disingenuous. Thus, we suggest instead to enable photography robots to recognize mistakes and be honest about undesirable outcomes, such as poorly framed pictures. Honest responses can build rapport with the robot and strengthen the interaction. Our work reinforces the importance of human trust in social HRI [4, 9, 14, 28–31, 35] and highlights the need for robust, failure detection methods [15].

When adding humor to an interaction, consider the cost of time. Because it is challenging to find humor that everyone will find a funny, the time required to convey the humor should be seen as a cost borne by those who don't find it humorous. Thus, if possible, we recommend adding humor in parallel to the main interaction. For example, Shutter could show a funny GIF while simultaneously counting down to take the picture. The cost of humor is worth considering in other HRI scenarios in which a robot's quality of being comic plays a key role in the interactions.

9 CONCLUSION

We explored the design of a humorous robot that takes portrait pictures while trying to elicit spontaneous smiles from subjects. Our human-centered design process included both professional and amateur photographers. We engaged with them through interviews and an initial, informative study to better understand subject-photographer interactions and create a working prototype of our robot photographer. We then conducted an evaluation of our design and found that our robot was able to leverage humor to elicit spontaneous smiles from subjects. The more people smiled naturally as a result of the robot's humor, the higher they rated their photography experience with the robot. We discussed design recommendations based on the insights gained from this process.

ACKNOWLEDGMENTS

We thank Annie Gao for helping implement software components for our robot photographer.

REFERENCES

[1] Hyunsang Ahn, Dohyung Kim, Jaeyeon Lee, Suyoung Chi, Kyekyung Kim, Jinsul Kim, Minsoo Hahn, and Hyunseok Kim. 2006. A robot photographer with user interactivity. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 5637–5643.

[2] Marlène Aillaud and Annie Piolat. 2012. Influence of gender on judgment of dark and nondark humor. *Individual Differences Research* 10, 4 (2012), 211–222.

[3] Anne E Beall. 2007. Can a new smile make you look more intelligent and successful? *Dental Clinics of North America* 51, 2 (2007), 289–297.

[4] Deborah R Billings, Kristin E Schaefer, Jessie YC Chen, and Peter A Hancock. 2012. Human-robot interaction: developing trust in robots. In *Proceedings of the seventh annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 109–110.

[5] Rogerio Bonatti, Yanfu Zhang, Sanjiban Choudhury, Wenshan Wang, and Sebastian Scherer. 2018. Autonomous drone cinematographer: Using artistic principles to create smooth, safe, occlusion-free trajectories for aerial filming. *arXiv preprint arXiv:1808.09563* (2018).

[6] Zachary Byers, Michael Dixon, Kevin Goodier, Cindy M Grimm, and William D Smart. 2003. An autonomous robot photographer. In *Proc. 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Vol. 3. IEEE, 2636–2641.

[7] Zachary Byers, Michael Dixon, William D Smart, and Cindy M Grimm. 2004. Say cheese! Experiences with a robot photographer. *AI magazine* 25, 3 (2004), 37–37.

[8] Manuel G Calvo, Andrés Fernández-Martín, and Lauri Nummenmaa. 2012. Perceptual, categorical, and affective processing of ambiguous smiling facial expressions. *Cognition* 125, 3 (2012), 373–393.

[9] Munjal Desai, Poornima Kaniarasu, Mikhail Medvedev, Aaron Steinfeld, and Holly Yanco. 2013. Impact of robot failures and feedback on real-time trust. In *Proceedings of the 8th ACM/IEEE International Conference on Human-Robot Interaction*. IEEE Press, 251–258.

[10] Hamdi Dibeklioğlu, Albert Ali Salah, and Theo Gevers. 2015. Recognition of genuine smiles. *IEEE Transactions on Multimedia* 17, 3 (2015), 279–294.

[11] Dennis J Dunleavy. 2004. *In the age of the instant: The influence of the digital camera on the photojournalistic routines of productivity, empowerment and social interaction between subject and photographer*. Ph.D. Dissertation.

[12] Paul Ekman, Richard J Davidson, and Wallace V Friesen. 1990. The Duchenne smile: Emotional expression and brain physiology: II. *Journal of Personality and Social Psychology* 58, 2 (1990), 342.

[13] Edward Twitchell Hall. 1910. *The Hidden Dimension*. Vol. 609. Garden City, NY: Doubleday.

[14] Peter A Hancock, Deborah R Billings, Kristin E Schaefer, Jessie YC Chen, Ewart J De Visser, and Raja Parasuraman. 2011. A Meta-Analysis of Factors Affecting Trust in Human-Robot Interaction. *Human Factors* 53, 5 (2011), 517–527.

[15] Shanee Honig and Tal Oron-Gilad. 2018. Understanding and resolving failures in human-robot interaction: Literature review and model development. *Frontiers in Psychology* 9 (2018), 861.

[16] Bill Hurter. 2004. *Portrait Photographer's Handbook*. Amherst Media, Inc.

[17] Bill Hurter. 2008. *The Best of Portrait Photography: Techniques and Images from the Pros*. Amherst Media.

[18] L. Jacobs. 2010. *The Art of Posing: Techniques for Digital Portrait Photographers*. Amherst Media.

[19] Tonglin Jiang, Hao Li, and Yubo Hou. 2019. Cultural Differences in Humor Perception, Usage, and Implications. *Frontiers in Psychology* 10 (2019), 123. https://doi.org/10.3389/fpsyg.2019.00123

[20] Myung-Jin Kim, Tae-Hoon Song, Seung-Hun Jin, Soon-Mook Jung, Gi-Hoon Go, Key-Ho Kwon, and Jae-Wook Jeon. 2010. Automatically available photographer robot for controlling composition and taking pictures. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 6010–6015.

[21] GA Rajesh Kumar, Ravi Kant Kumar, and Goutam Sanyal. 2017. Discriminating real from fake smile using convolution neural network. In *2017 International Conference on Computational Intelligence in Data Science (ICCIDS)*. IEEE, 1–6.

[22] Joseph C. Kush. 1997. Relationship Between Humor Appreciation and Counselor Self-Perceptions. *Counseling and Values* 42, 1 (1997), 22–29. https://onlinelibrary.wiley.com/doi/abs/10.1002/j.2161-007X.1997.tb00950.x

[23] CNN Maisy Mok. 2019. This robot photographer will take your profile pic for you. https://www.cnn.com/travel/article/worlds-first-robotic-photographer-intl-hnk-scli

[24] Andrew Mendelson. 2015. The construction of photographic meaning. *Handbook of Research on Teaching Literacy Through the Communicative and Visual Arts, Volume II: A Project of the International Reading Association* (2015), 27.

[25] Nicole Mirnig, Susanne Stadler, Gerald Stollnberger, Manuel Giuliani, and Manfred Tscheligi. 2016. Robot humor: How self-irony and Schadenfreude influence people's rating of robot likability. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 166–171.

[26] Tobias Nägeli, Lukas Meier, Alexander Domahidi, Javier Alonso-Mora, and Otmar Hilliges. 2017. Real-time planning for automated multi-view drone cinematography. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 132.

[27] Monica Perusquia-Hernández, Saho Ayabe-Kanamura, Kenji Suzuki, and Shiro Kumano. 2019. The Invisible Potential of Facial Electromyography: A Comparison of EMG and Computer Vision when Distinguishing Posed from Spontaneous Smiles. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 149.

[28] Paul Robinette, Ayanna M Howard, and Alan R Wagner. 2017. Effect of Robot Performance on Human-Robot Trust in Time-Critical Situations. *IEEE Transactions on Human-Machine Systems* 47, 4 (2017), 425–436.

[29] Maha Salem, Gabriella Lakatos, Farshid Amirabdollahian, and Kerstin Dautenhahn. 2015. Would you trust a (faulty) robot?: Effects of error, task type and personality on human-robot cooperation and trust. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 141–148.

[30] Tracy Sanders, Kristin E Oleson, Deborah R Billings, Jessie YC Chen, and Peter A Hancock. 2011. A model of human-robot trust: Theoretical model development. In *Proceedings of The Human Factors and Ergonomics Society Annual Meeting*, Vol. 55. SAGE Publications Sage CA: Los Angeles, CA, 1432–1436.

[31] Tracy L Sanders, Tarita Wixon, K Elizabeth Schafer, Jessie YC Chen, and PA Hancock. 2014. The influence of modality and transparency on trust in human-robot interaction. In *2014 IEEE International Inter-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*. IEEE, 156–159.

[32] Anne E Scanlon and Danielle C Polage. 2011. The Strength of a Smile: Duchenne Smiles Improve Advertisement and Product Evaluations. *Pacific Northwest Journal of Undergraduate Research and Creative Activities* 2, 1 (2011), 3.

[33] SLR Lounge Sean Lewis. [n.d.]. Eva The Robot Photographer-For-Hire Now Photographs Weddings. https://www.slrlounge.com/eva-the-robot-photographer-for-hire-now-photographs-weddings

[34] William Smart and Cindy Grimm. 2003. (Not) Interacting with a Robotic Photographer. In *AAAI Spring Symposium*. 181–186. Interaction observations with Lewis the robot photographer.

[35] Sarah Strohkorb Sebo, Margaret Traeger, Malte Jung, and Brian Scassellati. 2018. The ripple effects of vulnerability: The effects of a robot's vulnerable behavior on trust in human-robot teams. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 178–186.

[36] Leila Takayama and Caroline Pantofaru. 2009. Influences on proxemic behaviors in human-robot interaction. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 5495–5502.

[37] TJ Thomson and Keith Greenwood. 2017. Beyond Framing: Influences of subject–photographer interactions on visual journalism. *Journalism Practice* 11, 5 (2017), 625–644.

[38] Hitomi Tsujita and Jun Rekimoto. 2011. HappinessCounter: Smile-encouraging Appliance to Increase Positive Mood. In *CHI '11 Extended Abstracts on Human Factors in Computing Systems (CHI EA '11)*. ACM, New York, NY, USA, 117–126. https://doi.org/10.1145/1979742.1979608

[39] Marynel Vázquez, Elizabeth J Carter, Braden McDorman, Jodi Forlizzi, Aaron Steinfeld, and Scott E Hudson. 2017. Towards robot autonomy in group conversations: Understanding the effects of body orientation and gaze. In *Proc. of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 42–52.

[40] Michael L Walters, Kerstin Dautenhahn, Kheng Lee Koay, Christina Kaouri, R te Boekhorst, Chrystopher Nehaniv, Iain Werry, and David Lee. 2005. Close encounters: Spatial distances between people and a robot of mechanistic appearance. In *5th IEEE-RAS International Conference on Humanoid Robots*. IEEE, 450–455.

[41] Cornelia S Wendt and Guy Berg. 2009. Nonverbal humor as a new dimension of HRI. In *RO-MAN 2009-The 18th IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, 183–188.

[42] Paul Wilkinson and Sarah Plater. 2015. *Mastering Portrait Photography*. Ammonite Press.

[43] Manfredas Zabarauskas and Stephen Cameron. 2014. Luke: An autonomous robot photographer. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 1809–1815.

Designing Social Interactions with a Humorous Robot Photographer

Supplementary Material

A. Software

A diagram of the modular software architecture used for the Evaluation Study (Sec. 5) is found in Figure A.1. A simplified version of this architecture was used for the Informative Study (Sec. 4.3). The primary input sources to the system include audio data captured by the robot's microphone array, and RGB-D data from the robot's camera. The main outputs are the robot's body motion controlled by its joints, the visual displays of its screen and the auxiliary monitor used to display the photos that the robot took, and audio communicated through the robot's speakers. The next section details the main modules of the robot's software architecture.

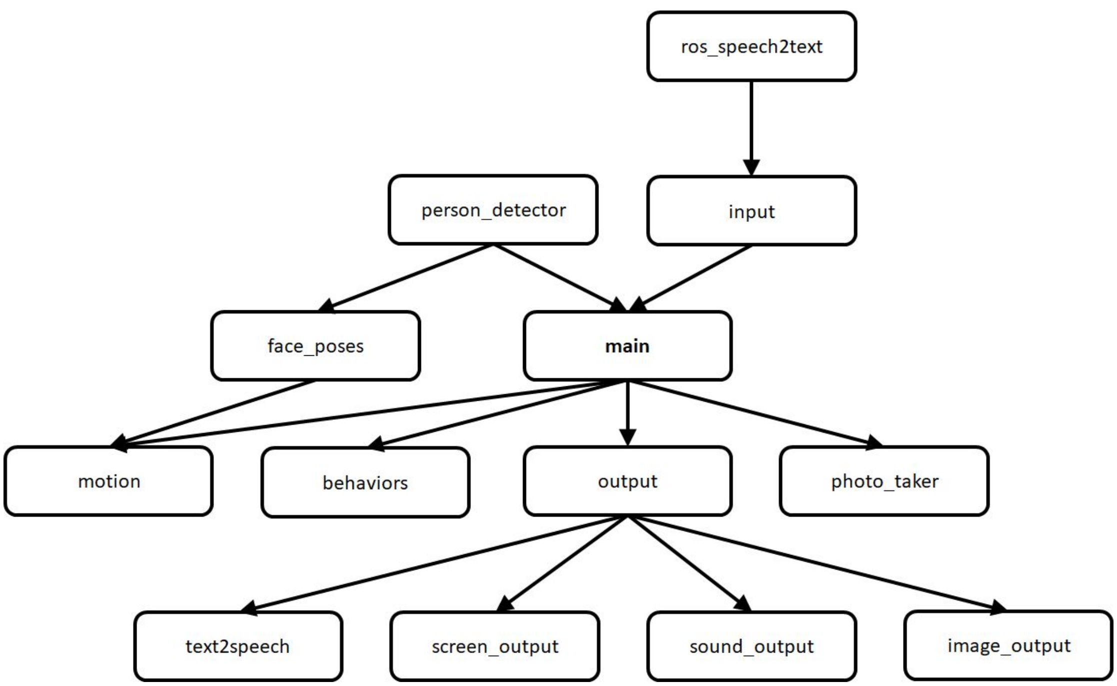

Figure A.1. Our robot's software architecture.

Software Modules

ros_speech2text: Converts speech spoken by the user to text using Google Cloud services.
person_detector: Detects faces found in the images acquired by the robot's camera with the OpenCV library.
input: Performs initial processing of the audio input. Responses from the user, such as "sure" or "no", are converted to a binary signal which represents either an affirmative response, or a negative response.
face_poses: Uses depth images captured by the robot's camera and the face detections to determine the 3D location of each detected face relative to the robot.
main: Handles the flow of the interaction. Orchestrates the robot's motion and other communication modalities in the photography context based on sensed inputs.
motion: Moves the joints of the robot to look at one of the detected faces. Also moves the eyes to look at a face as well.
behaviors: Moves the robot to perform a fixed behavior such as looking at the monitor.
output: Performs initial processing of the output. A configuration file and the output's type are used to determine where to send the output. For example, if the output is text, the configuration file is consulted to determine if the text should be spoken, displayed on the screen, or both.
photo_taker: Captures an image and saves it to the disk.
text2speech: Converts text to a speech.
screen_output: Displays text output on the screen.
sound_output: Plays the desired sound.
image_output: Displays an image on the screen.

B. Prior Experience Questions

The demographics questionnaire used in the Evaluation Study (Sec. 5) asked participants about their background and preferences regarding photography and robots. The specific wording of the questions is included below to provide a more thorough understanding of the study:

- How much experience with photography do you have?
- How much experience with robots do you have?
- How much do you enjoy taking photos?
- How much do you enjoy having your photo taken?

Quantitative results for the above questions are discussed in Sec. 6.6.

C. Post-Interaction Questions

After participants of the Evaluation Study (Sec. 5) interacted with the robot, they answered questions about their experience. The question formats included 7-point Likert format, typed short response, and transcribed discussions. We include the exact wording of the questions in this supplementary material for completeness.

Photography Satisfaction Survey

The Satisfaction Survey was inspired by a Customer Satisfaction Survey by Juhl Photography, a professional photography business in Watertown, WI. The survey had questions that each asked about one specific aspect of participants' experience with Shutter. Each question presented the name of the respective aspect that it referred to and a brief explanation. Participants responded using a 7-point Likert responding format (with 1 being Very Unhappy and 7 being Very Happy). We only included in the paper's Results (Sec. 6) the Satisfaction Survey questions that related to our most significant findings. The complete list of questions is found below:

How would you rate Shutter on:

- *Timeliness* - How well did your photographer meet your expectations concerning length of session?
- *Friendliness* - The friendliness of your photographer?
- *Professionalism* - The professionalism of your photographer?
- *Knowledge* - The knowledge and skill of your photographer?
- *Comfort* - How well did your photographer make you feel comfortable during the session?
- *Attitude* - The attitude of your photographer?
- *Poses* - How happy are you about the number of different poses taken during your session?
- *Service* - If there were any problems or concerns, how well did your photographer handle them for you?
- *Quality* - The quality of your final portraits?

Open Ended Survey Questions

We designed open-ended questions to provide us with a qualitative understanding of the participants' experience. Responses can be found in Results (Sec. 6), and the exact wording of the questions is found below:

- What did you like about Shutter's direction?
- What would you change about Shutter's direction?
- What did you like about your interaction with Shutter?
- What did you dislike about your interaction with Shutter?

- What could Shutter have done differently to improve your experience?

Final Discussion Questions

Discussions varied depending on which aspects of the interaction each participant found most notable. There were a few questions that we often asked, however, because of their relevance to the study. The most common discussion questions were:

- What was you favorite humorous content and why?
- What was your favorite type of humorous content and why?
- If you could add more humorous content, what would you add?

The key findings of the discussion are found in the paper's Results (Sec. 6).