

Jillian Greczek
greczek@usc.edu

Amin Atrash
atrash@usc.edu

Maja Mataric
mataric@usc.edu

What is Graded Cueing?

Graded cueing is a process of behavior shaping that uses increasingly specific cues, or prompts, to help improve people's skills at everyday tasks during recuperative therapy. It is used in treatment for individuals who have lost skills, such as through a brain injury, or need to learn new skills, such as social skills of individuals with autism spectrum disorders (ASD).

In this process, the therapist asks the patient to perform a task, and prompts the patient with increasing specificity the more the patient struggles with the task. The goal of graded cueing is to increase both patient task performance and patient autonomy in performing the task, through minimal therapist intervention.

Model Implementation as the Copy-Cat Game

Our probabilistic representation of graded cueing is implemented in a "Copy Cat" imitation game played between a NAO robot and a child with ASD. In the game, the robot poses its arms and asks the child to copy its pose. If the child successfully copies the robot's pose, the robot gives positive verbal feedback, nods, and flashes its eyes green. If the child does not successfully copy the robot's pose, then the robot gives a starting prompt as determined by the model.

During the very first interaction, the starting prompt specificity is P1, but for each subsequent round of interaction over the entire study, the starting prompt is determined by the model. From there, if the child requires further prompting, the robot moves up within the four levels of prompt specificity in this particular implementation, corresponding to five states:

P0: No prompts are given (success)

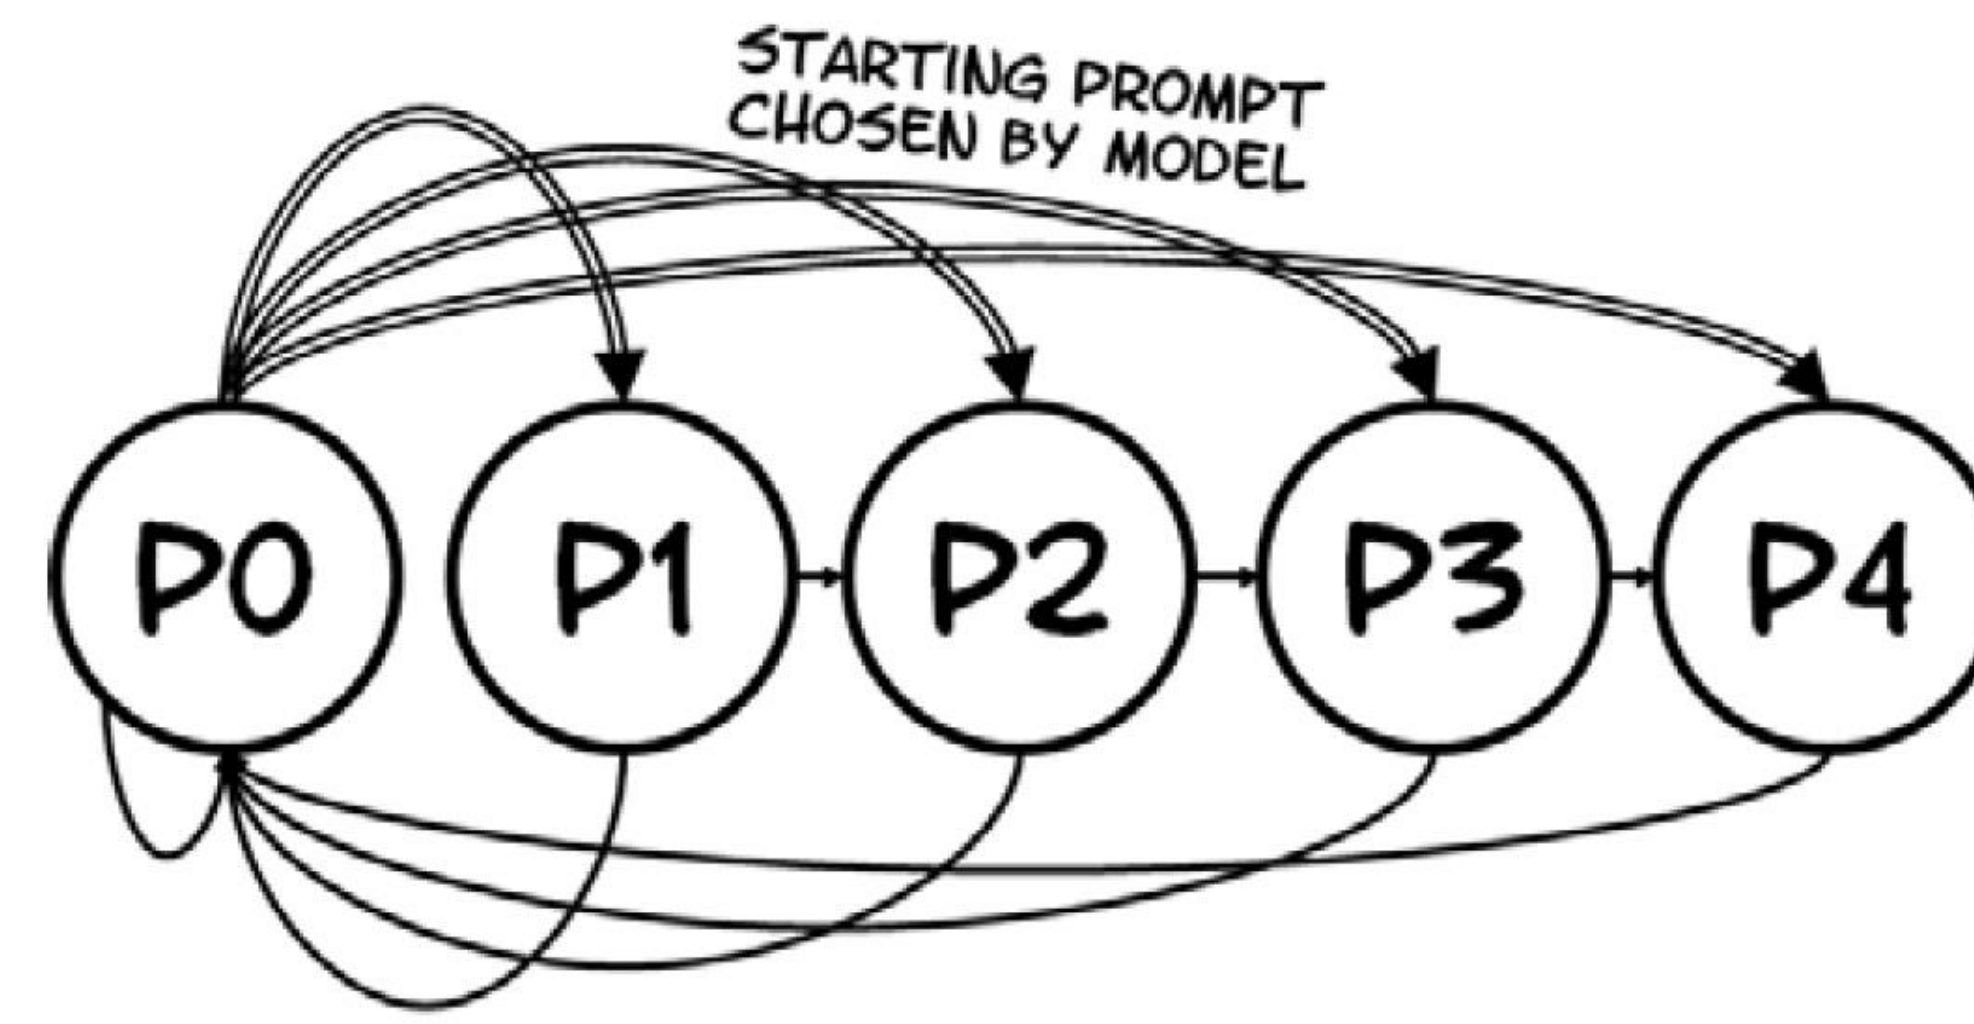
P1: words ("are you sure that's right?")

P2: words + gesture ("are you sure?" + indicative movement)

P3: specific words + gesture (specify right arm + movement)


P4: specific words + specific gesture ("like this")

After using the most specific prompt, the robot moves on to a new pose with an encouraging response such as, "Let's try an easier one," or "Let's try a different pose." The responses we used were derived from consultations with an ASD therapist and observations of a classroom with ASD children.



Our model of graded cueing informs how specific of feedback each participant needs based on their performance in previous rounds of imitation.

The Microsoft Kinect sensor is instrumental to this process. With it, we can keep track of the pose of the participant in real time.



The Aldebaran NAO robot models the pose for the child to imitate and gives verbal and nonverbal feedback.

Our work is motivated by the evidence that children with ASD are often behind in their development of imitative behavior, and that practicing through repeated interactions with a therapist can improve imitation abilities. Our model approximates those effects with a socially assistive robot in order to use technology to broaden access to ASD therapy.

Using Probabilistic Models to Make Up for Limited Data

This model consists of N states for N-1 levels of prompt specificity, plus a success state (P0). The model must infer the correct user ability level and select the most ideal prompt based on that sensed ability level. Higher levels of ability are associated with decreased specificity. The goal of our model is to select the first prompt to give to the user depending on the user's ability level. From that initial prompt, any subsequent prompt required is the next most specific prompt. During execution, the robot uses the child's responses to maintain a distribution over the ability level, and then to select appropriate actions based on a policy calculated over the model.

To compensate for few available data points that are typical of HRI contexts, our model must be able to adapt quickly to the changing user state. To achieve this rapid adaptation, we use a Bayesian approach, wherein the model maintains a distribution over the states, updates based on the responses, and selects actions by sampling from the distribution.


The Experiment




A pilot study was conducted to validate the described model with children with ASD at a local elementary school. There were 12 participants, drawn from an ASD-only class of students between the ages of 7 and 10.

The participants are split into two groups: one receives graded cueing feedback from the robot (N=6), and the other (control group) receives constant feedback, P4, from the robot in each round (N=6). The constant-P4 condition represents a lack of cueing, which was expected to result in a lack of generalization of the imitation skills learned by the child despite descriptive feedback. The study was conducted in 5, 15-minute sessions over 2.5 weeks. The robot's prompts were recorded along with video of the interactions.

Pilot Study Results



The above graphs show example rounds (prompts for a single pose) for participant T in the **graded cueing group**. The plot shows the average error (X, Y, and Z distance from ideal) after receiving prompts P1-P4. These graphs are only for rounds where prompting was necessary, which were about 1/5 of all rounds overall.



The above graphs show example sessions for participant D in the **control group**. The majority of participants in the control group got slightly worse with each additional P4 prompt. The majority of participants in the graded cueing group either stayed the same or slightly improved their performance, although not enough to pass the threshold for successful imitation.

Overall, the participants had little difficulty with the imitation task. Because participants "got it in one" for the majority of poses, the probabilistic model of prompt choice did not get enough data to adjust any participant's starting prompt beyond P1. However, as seen on the left, rounds that did require prompting saw different behavior on average from the graded cueing group than the control group. Participants receiving graded cueing tended to either stay the same or get better, while participants receiving no cueing tended to get worse. This suggests that varied feedback was more helpful for the participants in this pilot study.

Participants in the control group showed 4 instances of frustration with the robot over all sessions, while the graded cueing group showed only 2 instances. However, frustration was not necessarily correlated with poor performance at the task, suggesting that another benefit of graded cueing was in the qualitative experience for the participants.