Why We Should Build Robots That Both Teach and Learn

Timothy Adamson*
Yale University
tim.adamson@yale.edu

Debasmita Ghose*
Yale University
debasmita.ghose@yale.edu

Shannon C. Yasuda
Yale University
shannon.yasuda@yale.edu

Lucas Jehu Silva Shepard
Yale University
lucas.shephard@yale.edu

Michal A. Lewkowicz
Yale University
michal.lewkowicz@yale.edu

Joyce Duan
Yale University
joyce.duan@yale.edu

Brian Scassellati
Yale University
brian.scassellati@yale.edu

ABSTRACT

In this paper, we argue in favor of creating robots that both teach and learn. We propose a methodology for building robots that can learn a skill from an expert, perform the skill independently or collaboratively with the expert, and then teach the same skill to a novice. This requires combining insights from learning from demonstration, human-robot collaboration, and intelligent tutoring systems to develop knowledge representations that can be shared across all three components. As a case study for our methodology, we developed a glockenspiel-playing robot. The robot begins as a novice, learns how to play musical harmonies from an expert, collaborates with the expert to complete harmonies, and then teaches the harmonies to novice users. This methodology allows for new evaluation metrics that provide a thorough understanding of how well the robot has learned and enables a robot to act as an efficient facilitator for teaching across temporal and geographic separation.

CCS CONCEPTS

• **Human-centered computing** → **Collaborative and social computing systems and tools**.

KEYWORDS

Human-robot interaction, human-robot collaboration, musical robot, robot tutoring, robot learning

ACM Reference Format:
Timothy Adamson, Debasmita Ghose, Shannon C. Yasuda, Lucas Jehu Silva Shepard, Michal A. Lewkowicz, Joyce Duan, and Brian Scassellati. 2021. Why We Should Build Robots That Both Teach and Learn. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction (HRI '21), March 8–11, 2021, Boulder, CO, USA.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3434073.3444647

*Both authors contributed equally to the paper

HRI '21, March 8–11, 2021, Boulder, CO, USA

ACM ISBN 978-1-4503-8289-2/21/03…$15.00
https://doi.org/10.1145/3434073.3444647

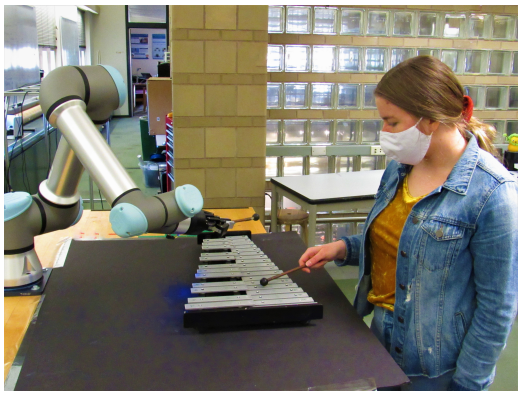

Figure 1: This robot first learns how to play three-note harmonies on a glockenspiel from an instructor, then plays on its own, or with the instructor, and finally becomes the instructor to teach a novice what it has learned.

1 INTRODUCTION

Advancements in machine learning and artificial intelligence have given robots the ability to gain knowledge, complete tasks, and teach humans like never before. Yet, to the best of our knowledge, there does not exist a robot system capable of all three. In traditional learning systems, a robot's learning process ends once the robot demonstrates that it can successfully complete a task. In conventional tutoring systems, a robot's teaching process does not allow for the robot to learn new skills from human experts. We argue that robot systems capable of learning a skill, performing it, and teaching it to a novice have advantages that can improve robot learning, collaborating, and teaching in significant ways.

In traditional learning models, evaluation of a robot's knowledge is limited to its performance. A robot is considered to have mastered a skill if it can complete a task to a desired accuracy level. While good performance may often be the primary goal of teaching a robot, building a truly robust system may require more comprehensive evaluation metrics to ensure there are no gaps in the robot's knowledge. By requiring a robot to teach what it has learned to a novice, we can measure the robot's level of understanding in novel

ways. While it may not be necessary for a robot to teach well to perform a task well (even highly skilled humans can be bad teachers), teaching provides a new viewpoint to evaluate understanding.

One of the most valuable abilities humans possess is our ability to pass the knowledge we have learned to others [60]. Much research has gone into creating robots that teach people new skills. However, none of these robots are designed to gain and adapt to new knowledge provided by a human teacher. As a result, knowledge transmission is stunted—the robots that can teach are unable to learn, and the robots that can learn are unable to teach.

Due to the robot's inability to act as both teacher and student, humans miss out on the benefits that teaching provides to the learning process. Within educational theory, having a student teach others has long been accepted as a valuable exercise for improving that student's knowledge [28, 29, 32, 63]. This is often credited to Jean Pol Martin who developed a practice called "Lernen durch Lehren" (LdL), German for "Learning by Teaching" [9] through which pupils were asked to teach relevant concepts to their classmates. Several studies of LdL have confirmed its value in improving learning gains [1, 26, 61, 66]. The method used to train surgeons, described by the adage "See One, Do One, Teach One," also requires medical students to learn, perform, and teach a procedure to be considered proficient [39]. Its benefits have been similarly validated [69].

We believe that by combining robot learning, doing, and teaching into a single process, we can improve learning for both humans and robots in multiple ways. In our methodology, the robot begins as a novice. An expert teaches the robot a skill by providing a set of demonstrations, tests the robot, and provides feedback on its attempts. Once the robot has learned the skill, the robot can then perform the task on its own, or collaboratively with another agent. The robot then teaches a novice the skill through explanation and examples. It then tests the novice, correcting them when a mistake is made. The interaction can even continue on to have the novice teach the robot the recently acquired skill. This allows the novice to improve their understanding through teaching, and the robot to better assess how well the novice has learned. This methodology could potentially be applied to any knowledge that can be transmitted from an expert to a robot, and from a robot to a novice.

There are understandable drawbacks to this methodology. It requires significant time and investment, which may be unnecessary if the primary goal is only for the robot to complete a simple task. Therefore, there are likely circumstances in which this system does not prevail over traditional learning and teaching systems. Nevertheless, if we want robots whose knowledge is well-evaluated and who can transmit the knowledge they have learned to others, we must build robots that can both teach and learn.

To study this methodology, we chose playing three-note harmonies on the glockenspiel as the simplistic proof-of-concept skill for the robot to learn, perform, and teach. We carried out a case study in which the robot learned the skill of three-note harmonies through observing an expert, collaborated with the expert on playing three-note harmonies, and then taught the skill to novices, who in turn, taught the robot what they had learned (Figure 1). We use this demonstration as both a feasibility proof and to highlight some of the benefits and drawbacks of this approach.

2 RELATED WORK

Because our work involves the process of a robot learning from humans, completing a task independently or through collaboration with humans, and teaching humans, our research was informed by previous work in each of these areas.

2.1 Robot as Student

There are many ways of providing a robot with the ability to complete a task. Traditionally, this required having a roboticist program the robot exactly how to complete the task. But in the past 20 years, there has been a growing interest in providing robots with models that allow them to learn from humans without being explicitly programmed what to do. This area of research generally falls within the category of learning from demonstration.

Learning from demonstration occurs when a robot observes a set of demonstrations from a human accomplishing a task and infers an action sequence based on those demonstrations to accomplish the task itself. This can happen at a variety of abstraction levels. At the action level, the robot learns the physical motions of the demonstrator [5, 6, 18, 21]. At the program level, the robot considers the higher-level structure of more complex actions [5, 22, 38, 48, 59]. Demonstrations can be combined with verbal communication [55] and dialog and gestures [15] to improve learning capabilities. Some research has also focused on enabling the robot to ask questions [20] or to request more examples when those provided are not deemed sufficient [19, 24, 43].

Our work expands on these findings by extending the learning process, using collaboration and teaching to better evaluate and improve the robot's understanding.

2.2 Robot as Collaborator

Much of human-robot collaboration involves establishing a shared task model. Having a shared task model has been shown to increase efficiency of a human in a human-robot collaboration setting [58]. Though task models can be hard-coded, it is generally preferred that they be learned during the interaction [23]. This learning can be done through various means, such as questions asked by the robot in natural language [51]. Some prior work has focused on learning both the high-level structure of a task and the low-level action sequences [16, 34, 59]. Other work has studied how to use robust task models to adapt a robot's plan to accommodate human actions [8, 35]. In general, learning a well-structured task model that is robust to human intervention dramatically improves the human-robot collaboration experience. Our work requires the learning of task models not just for collaboration, but also for teaching, ensuring that learned models lend well to introspection and explanation.

2.3 Robot as Teacher

Robot tutoring systems are designed to conduct the learning experience via social interaction between the robot and a human learner [10]. It is well established that personalized tutoring that adapts to student needs improves learning outcomes [11, 33, 52, 64].

Robots have been placed in the roles of peer and novice to aid human learning as well. As a peer, the robot might act as a learning companion, providing motivational support for the student without passing on any knowledge [44]. As a novice, the robot allows the

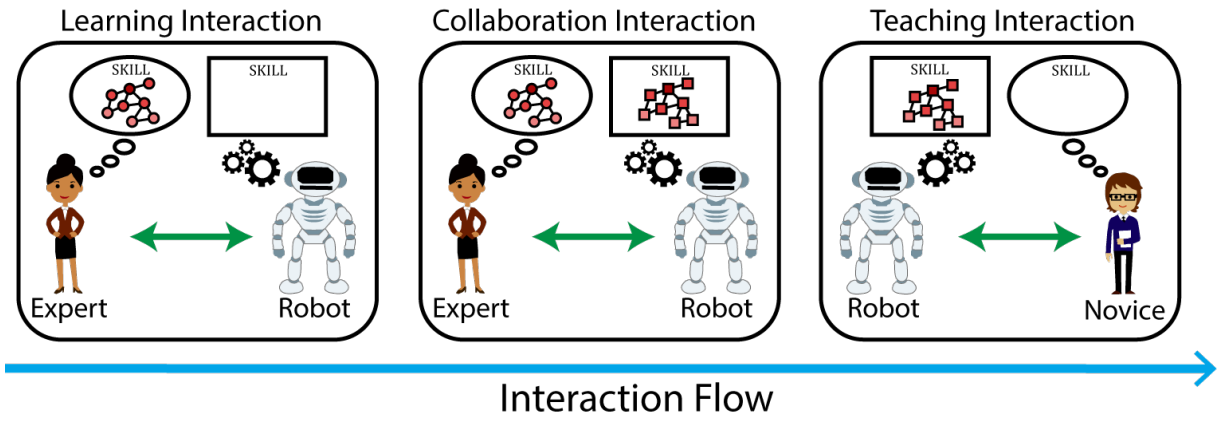

Figure 2: High-level overview of the methodology: the robot learns a skill from an expert, collaborates with the expert on that skill, and then teaches a novice the skill, such that the task model becomes shared between the expert, robot, and novice. The bi-directional arrows in the figure highlight how learning, collaboration and teaching are bi-directional processes.

student to act as instructor, leading to improvements in engagement and meta-cognitive skills [62]. Research into these roles has demonstrated that both can increase student learning [36, 71]. A recent study found that an adaptive robot, capable of dynamically exhibiting roles of both tutor and novice while teaching a student, can even further improve learning outcomes [25].

Our methodology builds on previous efforts by combining the roles of novice, collaborator, and teacher, and granting the robot the ability to acquire new skills. Through this system, both the robot and the novice have the potential to benefit from each role.

3 METHODOLOGY

This section describes the different components of our methodology and analyzes the advantages and disadvantages of creating robots that learn, collaborate, and teach.

3.1 Description of Components

Combining the three components of our methodology—learning, collaboration, and teaching—provides a system with a variety of advantages over systems that exhibit only one or two components. Figure 2 shows a high-level visual overview of our approach.

In the first component, the *Learning Interaction*, the robot learns a skill from a human. A vital property of the learning model is that it needs to be interpretable to facilitate the transfer of knowledge to humans. For our purposes of having a shared knowledge representation for learning, collaboration, and teaching, we only need the robot to explain and reason about what decisions it makes for a given scenario using a particular model. We do not necessarily need the robot to understand how the model works, in the sense that each input, parameter, and calculation would admit an intuitive explanation [42]. Hence, this component could leverage a wide variety of learning techniques, including those that are easily understandable as well as those that are more opaque. The more opaque learning techniques like machine learning, deep learning, and reinforcement learning can be used as long as the robot can explain the rationale behind its decisions during the collaboration and teaching processes. Methods that add this level of explainability to these opaque learning techniques are an active research area and require significant time and effort to create [4, 45, 46, 65].

The *Collaboration Interaction* allows the human and the robot to collaborate on a task in a shared environment. For a robot to become a skilled collaborator, both the robot and the human should have a shared model of the task. Establishing a shared model can be challenging. Previous work used high-level task models combined with POMDPs [54] and gaze cues [49] to help develop shared task models. Our methodology supports shared task models in two ways. The first is the robot's ability to learn from its collaborator. Learning from the collaborator allows the robot to adjust its task model to better align with the human's task model. The second is the robot's ability to teach its collaborator. The robot can utilize its internal model's explainability to share with the collaborator its internal model of the task, leading to better alignment of the two models.

Finally, in the *Teaching Interaction*, the robot aims to transfer its task model to a novice. For a robot to be an effective teacher, the robot should periodically test the human on the concepts taught. To do that, a robot needs to predict and contextualize a student's performance throughout an interaction to track the student's progress. This motivates the need for knowledge tracing. Existing literature shows that common knowledge tracing models such as Bayesian Knowledge Tracing (BKT) offer adequate performance when task success is binary [27, 70]. Other works explore general-purpose models such as Deep Knowledge Tracing (DKT) frameworks that utilize deep learning models that can account for nuances in student learning that are difficult to quantify in other structured models [37, 50]. Depending on the the task model, our methodology can leverage any knowledge tracing model in complex task settings, allowing the robot to modify its teaching strategy as necessary and ensure successful transfer of knowledge to the student.

3.2 Advantages

Our methodology provides many advantages over the current standard for robot learning [53]. These include more comprehensive evaluation metrics, flexible role selection for the robot, the ability

to transfer knowledge across temporal and geographic separation, and the ability to robustly teach people.

3.2.1 More comprehensive evaluation metrics. Our methodology supports current evaluation metrics for assessing the robot's performance on a learned task which typically compare a robot's ability to perform a task in relation to a human's. It also proposes new metrics to evaluate how quickly a robot learns, how well it collaborates, and how successfully it teaches. These metrics, used alongside the traditional metrics, provide a more comprehensive assessment of the robot's understanding of a task. We propose the following qualitative and quantitative measures of evaluating the process, including, but not limited to:

- how long it took before the expert deemed the robot to be competent
- how good of a collaborator the expert rates the robot to be
- how many times the expert had to correct the robot while collaborating
- how many demonstrations and explanations the robot gave before the novice gained competence at the task, and
- how good of a teacher the novice rated the robot to be

It should be noted that many of the proposed evaluation metrics likely have greater variability than standard task evaluation metrics because the proposed metrics involve a human-in-the-loop [12].

Additionally, our methodology allows us to evaluate the novice's proficiency at a task based on the robot's trained models. This is possible by having the novice re-train the robot after they have learned from it. The two task models, one taught by the expert and another taught by the novice, can be compared using an appropriate difference metric. If the robot internalized and relayed the knowledge it learned, this difference is expected to be a small value. This provides a strong quantitative assessment of how well the novice has learned the skill because one can now directly compare the novice's task model with the expert's task model and identify which parts of the task the novice has and has not learned.

3.2.2 Flexible role selection. There are benefits to enabling a shared knowledge representation between our methodology's learning, collaboration, and teaching components. Notably, this shared representation allows the robot to switch between components seamlessly, based on its and the human partner's needs. For example, if the robot identifies that it doesn't understand a skill, it could take a student's role and ask for help from its partner. The robot could also take a teacher's role if the robot infers that its partner is facing difficulties performing the assigned task. These interactions are only possible because the robot uses the same task models for learning, collaboration, and teaching. Therefore, it can infer when a partner's task model does not align with its own and adapt accordingly.

3.2.3 Knowledge transfer across temporal and geographic separation. Our methodology allows experts and novices to teach and learn without being present in the same geographic location at the same time. An expert can train the robot on any given day, and a novice can learn the same task from the robot on a different day or from a similar robot at a different location. This flexibility can be advantageous in situations where it is logistically infeasible or unsafe for people to learn the task from each other while being present in the same place simultaneously, for example, during quarantine in a pandemic [7, 47, 57].

A robot as a part of this proposed system can also facilitate the transfer of rare skills or cultural heritage that needs to be preserved between generations of people when there is no direct contact between the people who possess that skill and the people who wish to learn it. An embodied robot would be better than virtual sources of learning for this transfer of knowledge. This is shown by existing work [67, 68], which demonstrates that physical embodiment has a measurable positive impact on the performance and perception of social interactions like teaching.

Another potential application of this methodology can be improving the current state of hands-on skill transfer in distance learning. Instructors can now teach skills that can only be acquired by performing the task under expert supervision via a robot, which is not feasible with traditional distance learning.

Finally, this methodology allows for many novices to be trained in an individualized learning environment, with minimal time and effort required of the expert. Once the expert has taught the robot a task, the robot can then perform teaching interactions with many novices. The robot can even customize its instruction for each learner, taking as much time as needed with each novice with no intervention from the human expert.

3.2.4 A robust method for teaching people. This methodology allows for developing more robust learning systems for people by enabling them to teach the robot what they have learned. This provides significant learning gains to the person because they can take advantage of the benefits of "learning by teaching" discussed in Section 1. The robot can also now provide them with an additional, objective metric to evaluate their understanding of a skill, which goes beyond merely their ability to perform the skill. This can come in the form of a comparison between the model taught to the robot by the expert and the model taught by the novice. Additionally, based on the evaluation metrics described in Section 3.2.1, the robot can then tailor its instruction program and its pace of teaching according to what seems to be working best for any particular learner. People also have the opportunity to learn interactively by physically collaborating with the robot.

3.3 Challenges

There are challenges to our methodology that we acknowledge in detail below. The significance of these challenges will vary depending on the application of this methodology.

3.3.1 More robust models take longer to design. Because our methodology requires the robot to learn, collaborate, and then teach, the robot's knowledge representation must be more complex and generalizable than it would be otherwise. Designing explainable models generally requires more effort than creating ones that are opaque and difficult for humans to understand. This cost of time must be borne by those responsible for creating the models. Opaque models can be used in our methodology only if there are additional components that interface with the opaque models to add explainability in the form of visual saliency [13, 17, 31, 56] or natural language explanations [30, 40], which are under active investigation.

Table 1: This table represents the note intervallic arrangement for each harmony type. The first note, represented by *n*, can be any bar on the glockenspiel as long as the third note exists on the instrument. The next two notes are described in reference to this root note, with semitone steps.

Harmony Type	Note 1	Note 2	Note 3
Broadhurst	*n*	*n+3*	*n+5*
Foxwoods	*n*	*n+1*	*n+4*
Belasco	*n*	*n+6*	*n+9*

3.3.2 Complex system requirements. Creating a robot that can not only learn but also collaborate and teach comes with corresponding system requirements. For the robot to teach the human, it must explain its understanding of the task to the human. In addition, the robot needs to semantically disambiguate what the human says for successful bi-directional communication between the human and the robot, which is not trivial. Further, because the robot and human will also likely be interacting in a shared space, the robot's hardware must be designed for safe interactions. Different skills might also demand additional software and hardware components, such as sophisticated perception systems, manipulation capabilities, and language models.

3.3.3 Problems of a generalist approach. This methodology requires a generalist approach to learning. Because the robot must learn from expert demonstration, we cannot make assumptions about the problem space or solution for any of the three interaction types (learning, collaborating, and teaching). The robot must be able to adapt to whatever knowledge is taught by the expert. This leads to more complex, higher-dimensional learning problems. It also makes designing the right collaboration and teaching interactions more difficult because, unlike in traditional robot collaboration and tutoring systems, the robot begins with an internal model that has not been trained.

4 CASE STUDY

In this section we present a proof-of-concept application for our methodology: playing three-note harmonies on the glockenspiel. We outline the interaction flow unique to our case study and present an evaluation of our case study with a small group of users.

4.1 Task Selection

We selected the skill of identifying and playing three-note harmonies, a relevant sub-skill for playing the glockenspiel, as a case study for our methodology. Three-note-harmonies are the composite product of three different notes played simultaneously in a specific intervallic arrangement. In Western music, harmonies are generally categorized as either major, minor, or diminished. However, to make sure that all novices have the same baseline knowledge at the start of the interaction, we've created three of our own harmonies, named Broadhurst, Foxwoods and Belasco. The intervals that denote three-note harmonies in each harmony type are shown in Table 1.

We decided to use the glockenspiel for our case study because this domain worked well with our available hardware, and the glockenspiel's independent bars provided a discretized state space to use for our knowledge representation. A sequence of notes can

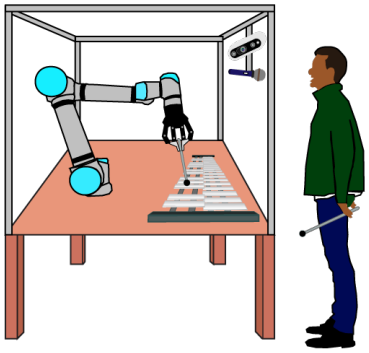

Figure 3: Experimental setup for the case study.

easily be represented by the glockenspiel bars which corresponds to those notes, and the duration of time between the striking of each subsequent bar. The three notes played sequentially either form a correct harmony, or they don't. Although having a discretized state space for both the input to the learning model and the task evaluation can make the model design and development easier, they are in no way requirements of our methodology.

Since our methodology focuses on the learning and teaching of one particular skill, it's reasonable to assume that a robot must have already learned any necessary sub-skills. In the case of learning three-note harmonies, a robot must already know how to hold a mallet, hit notes on the glockenspiel, and determine which bar(s) it's partner struck. The fact that an average human novice has these prerequisite skills in their existing skill set makes the task of learning three-note harmonies ideal for our methodology, adding to the motivation behind choosing this skill as our system's proof of concept.

4.2 Implementation

We use a UR5e robot to implement our case study. As illustrated in Figure 3, the robot is placed in front of the glockenspiel such that it can hit every bar with a mallet. Because the UR5e only has one arm, the three notes of the harmony must be played sequentially, instead of following the more common practice of playing them simultaneously. Our robot uses an Intel RealSense camera to observe the glockenspiel bars. It identifies voice commands through a SE VR-2 microphone and Google Cloud's speech-to-text module, and communicates using Google Cloud's text-to-speech module.

The robot detects which bar was hit by locating the mallet's position relative to a reference image. The reference image was annotated to indicate the bounding-box coordinates of each bar on the glockenspiel. Then, for every frame obtained by the camera in real-time, the algorithm looks for the circular mallet's presence by comparing the current video frame to the reference image, using the OpenCV toolbox [14]. If the robot detects a mallet, it matches the position of the circular mallet's centroid to the position of each

pre-annotated bounding-box in the reference image, returning the unique ID of the associated bar.

We developed the robot's internal task model using a deterministic, rule-based heuristic learner. The model is represented using a dictionary whose formulation is shown in equation 1.

$$(ID, bar_2 - bar_1) : [bar_3 - bar_2{}^1, ..., bar_3 - bar_2{}^k], \quad (1)$$

The dictionary's keys are two-valued tuples, each of which maps to a list of integers. Each tuple consists of two integers: the unique ID associated with each subskill, and the difference between the first two bars in the harmony. The list of integers represents all observed differences between the second and third bars in the harmony for k octaves. Storing these as a list prepares the robot for a scenario where a three-note harmony is constructed over multiple octaves. Once the model is trained, the robot can create valid harmonies on its own, complete harmonies when provided one or two bars, and evaluate the validity of a proposed 3-note harmony.

4.2.1 Learning Interaction. During the learning interaction, the robot queries an expert on the name of the skill it is learning and whether there are any sub-skills. The expert demonstrates the skill and the robot is required to repeat the demonstration to show knowledge of the skill. Specifically, for each 3-note harmony skill, the expert plays three notes. The robot's model uses the demonstrations provided to learn the rule. For every example that the expert provides, the robot's model checks whether the example conflicts with previously learned mappings. If so, the robot asks whether the expert can provide a new example that is compatible with the robot's existing model, or alternatively whether the robot should remove the incompatible data points from the model.

4.2.2 Collaboration Interaction. The collaboration interaction allows a human partner to play harmonies in collaboration with the robot, so that the robot can apply its acquired knowledge and fix any errors in its learned model. These are the four possible techniques we designed for the partner to collaborate with the robot:

- The human plays a single note and says the type of harmony. The robot must then fill in the following two notes based on the intervals associated with that harmony.
- The human plays two notes and then says the type of harmony. The robot must then fill in the remaining note based on the second interval associated with that harmony.
- The human states the type of harmony. The robot must build all three notes of the harmony on an arbitrary starting note.
- The human plays all three notes that build a harmony. The robot identifies the type of harmony associated with the three notes, or declares that no match exists.

4.2.3 Teaching Interaction. The robot conveys its internal knowledge representation through demonstrations and explanations to teach the novice three-note harmonies. The robot shows examples until the human announces they feel comfortable with that particular sub-skill. After the demonstrations, the robot tests the human by asking them to play one, two, or three notes needed to complete a harmony.

The robot's internal knowledge representation cannot make any assessment of the novice's current understanding of the skill. Therefore, to ensure that the knowledge has been successfully transferred from the robot to the novice, the robot must assess the the novice's skill comprehension throughout the interaction. This motivates the implementation of knowledge tracing. Specifically, we chose to use Bayesian Knowledge Tracing (BKT) [27, 70], an algorithm that predicts the level of skill mastery of a learner based on the learner's performance history on evaluative tests. Although we can use any form of knowledge tracing during the teaching phase, we chose BKT because of the inherent simplicity of our task model. BKT allows the robot to update the sequencing of administered tests and ultimately conclude the teaching phase. BKT models assume that student knowledge is represented as a set of binary variables, one per subskill, which denote whether the student either has or hasn't mastered the skill. We assign separate BKT models to each variation of the test and independently update the probability that the skill being tested was learned in that step. The parameters used for this calculation are:

- p_{slip}: the probability that the novice knew a harmony but made a mistake in the testing phase
- p_{guess}: the probability that out of all possible bar combinations the novice guessed the correct arrangement
- $p_{transit}$: the probability that skill mastery will be reached in the next testing step
- p_{learn}: the conditional probability that given the novice's input, the novice has attained mastery of the subskill

The calculation of p_{learn} is updated at every testing step and is conditional on the correctness of the novice's test response. An incorrect response will decrease the value of p_{learn} while a correct response will increase the value towards a desired threshold. All previously described model parameters are factored into the calculation of p_{learn} in a set of standard equations [27]. We empirically determined the initialization of p_{learn} to be 0.05 (i.e. there is a 0.05 probability that the novice knows the skill beforehand) and determined the p_{learn} threshold to be 0.95. Once p_{learn} reaches this threshold, the robot tutor assumes that skill mastery has been achieved, and testing is concluded. We initialize separate p_{learn} variables for each variation of the test to allow the robot to assess the novice's mastery of sub-skills separately.

4.3 Evaluation

To evaluate our case study, we had 6 participants (3 males, 3 females) interact with the robot. One participant, who is also an author of this paper, was the expert on playing the three-note harmonies. The other 5 participants were students un-affiliated with this research who were invited to learn three-note harmonies from the robot. Though 4 of the 5 novice participants had musical instrument experience, they were novices to the harmonies they were learning from the robot. The expert interacted with the robot for about 30 minutes, and the novices interacted with the robot for about 1 hour. We evaluated the whole methodology by following these steps:

(1) The expert taught the robot the three harmonies described in Table 1 using the *Learning Interaction*.
(2) The same expert and robot used the *Collaboration Interaction*.
(3) The robot taught the novice, using the *Teaching Interaction*.
(4) The novice and robot used the *Collaboration Interaction*.
(5) The novice taught the robot a new internal model using the *Learning Interaction*.

(6) The robot used the new internal model during the *Collaboration Interaction* with the novice.

By allowing the novice to teach the robot a new internal model and collaborate with the robot a final time (Steps 5 and 6), we were able to explore how a different application of the robots' flexible roles can allow a human to practice learning by teaching. We were also able to evaluate how well the robot learned and taught the novice by comparing the models taught by the novice and the expert.

The surveys used throughout this case study had 7-point Likert responses [41] and short-answer questions based on standard course and instructor evaluation surveys, typically administered in educational institutions [2, 3]. On the Likert responses, 1 stood for *"Strongly Disagree"* and 7 stood for *"Strongly Agree."* The first survey, consisting of questions about the novices' experience with music and robots, was administered before Step 3. A second survey, where the novice evaluated the robot's teaching capabilities, was conducted after Step 4. A final survey in which the novice evaluated the methodology as a whole was administered after Step 6.

4.3.1 Initial Analysis. Our case study evaluation showed that all five novices successfully learned how to play the same three harmonies which the expert had taught to the robot. This is validated by the fact that, by the end of each novice's teaching interaction with the robot, the novice could play each of the harmonies without any assistance. On average, each novice user made 0.9 mistakes and 4.4 successful harmony completions for each of the three harmonies during the teaching interaction in Step 3.

4.3.2 Applying novel evaluation metrics. Implementing our methodology allowed us to evaluate our glockenspiel-playing robot with a variety of novel metrics. We describe some of these metrics below, along with the accompanying results from our case study.

(1) Measuring the novices' performance provides insight into the success of the robot's teaching. The novices' average success rate when tested across all harmonies was 62%, which was lower than their success rate when learning the individual harmonies. This can be explained by the novice playing the same harmony multiple times in the teaching interaction rather than switching between all three harmonies in the performance assessment.
(2) Comparing the robot's task performance after learning from the expert and after learning from the recently-taught novice allows for evaluating how well the robot taught and learned. The robot's average success rate at playing the harmonies after learning from the novices was 66%. This shows that most novices were able to teach back to the robot what they had just learned.
(3) Calculating the difference between the model learned from the expert and the model learned from the novice allows for a fine-grained analysis of what knowledge the robot was and was not able to transfer to the novice. Our robot's knowledge model can be represented with a set. This allows us to compare the model learned from the expert and the model learned from the novice using set operations. This comparison reveals which specific subskills the novice did not understand well and can potentially help the robot tailor its future instruction to fill the novice's understanding gaps. Because of a small perception error that occurred when the expert was teaching Belasco harmonies to the robot, the robot learned an extra rule which resulted in the robot's model having three more elements than it should have had. This increased the symmetric distance and the set difference between the two models by three.
    (a) Symmetric distance represents what the novice was supposed to learn, but did not learn, along with any excess information that the novice taught the robot. The symmetric difference of the two models averaged across all the novices for Broadhurst, Foxwoods, and Belasco harmonies was 0, 2.4 and 7.8 respectively. The lower the symmetric distance, the better.
    (b) The set intersection evaluates how well the novice was able to re-teach the robot what it had learned. The intersection of the two models averaged across all novices for Broadhurst, Foxwoods, and Belasco harmonies was 6, 4.8, and 3.6, respectively, with the max intersection being 6. The higher the intersection, the more rules the participants were able to re-teach the robot correctly.
    (c) The set difference accounts for what the robot taught the novice that the novice did not teach back to the robot. The set difference averaged across all the novices for Broadhurst, Foxwoods, and Belasco harmonies was 0, 1.2, 5.4, respectively. The lower the set difference, the better.
(4) Using surveys to assess the robot's teaching and learning abilities provides a human-centered evaluation metric. We administered our surveys to the novices to obtain their subjective assessment of the robot's teaching and learning abilities.
    (a) In evaluating the robot as a teacher, the Likert question *"The instructor knows the subject area very well"* had an average score of 5.4 while the question *"The instructor effectively explained and illustrated the concepts"* had an average score of 6.8.
    (b) In evaluating the robot as a student, the Likert question *"The robot was able to learn quickly"* had an average score of 6 while the question, *"By the end of our interaction, the robot had a strong understanding of the material which I taught it"* had an average score of 5.8.

These evaluation metrics provide additional insight into the robot's knowledge beyond merely assessing the robot's ability to play the harmonies independently.

4.3.3 Demonstrating advantages of the methodology. Our case study also allows a robot to learn and teach across temporal and geographic separation. Evaluating our system would have been difficult without this advantage due to the COVID-19 pandemic. An expert taught the robot on a given day and the robot taught the majority of the novices days after it had learned from the expert, allowing us to follow adequate public health safety protocols.

Implementing and evaluating our case study demonstrated how our methodology allows learning by teaching. Analyzing our qualitative data showed that having the novice teach the robot what they just learned (Steps 5 and 6), allowed them to better understand the task themselves. When learning from the robot, none of the 5

novices realized that the notes in a three-note harmony could be played in any order. However, 3 of the 5 novices made this realization when they were teaching the robot. For example, when P5 was teaching the robot and saw it play the 3 notes in an order he did not expect he said "the order is flipped… so it's right."

5 DISCUSSION

5.1 Case Study Limitations

Our case study is admittedly straightforward. Since three-note harmonies follow a specific set of rules and because we've used a heuristic rule-based learner, the robot understood the rules that govern this skill with few demonstrations. In the future, we would like to implement this methodology for more complex tasks, such as sorting recyclables or assembling furniture. Our simplistic learning model allowed us to focus our efforts on building a "full-cycle" methodology, but using a probabilistic learning model with our methodology would be an excellent direction for future work.

Our case study may have benefited from a comparison to another system, but there does yet exist a system, capable of learning, collaborating, and teaching, to provide a meaningful baseline. Though Chen et al.'s [25] robot can act as both teacher and peer, the robot is unable to gather new knowledge. Comparing our system to one that only teaches, learns or collaborates would distract from our case study's primary purpose of exploring the advantages and disadvantages of a system that incorporates all three components.

5.2 Flexibility and Configurability

Until now, robots that teach a skill or collaborate by applying a skill have generally been designed with that singular skill in mind. The addition of new skills to the robot's repertoire requires retraining with a new skill set. Within our framework, a robot could instead be taught the skill by an expert and incorporate the skill into its existing task model, provided that the existing hardware supports it and that the robot has the relevant sub-skills to learn the skill. Skills that leverage existing sub-skills would additionally be easier for the robot to both learn and teach. The robot would build up an interconnected model of skills, in which each sub-skill could be evaluated and taught individually by the robot.

We presented the three components of our methodology in the order of learning, collaboration, and teaching. Though all interactions must begin with the learning component so that the robot has an initial task model, the methodology allows for significant flexibility of what follows next. After the robot has learned from an expert, it could then go on to collaborate and teach as many or few times as desired with different people or the same person. The various components can also be chained together multiple times to compute valuable evaluation metrics for the robot or the human learning the task. For example, once the robot has learned a task model from an expert, a peer could teach the robot that same skill. The robot could compare its original task model with that taught by the peer in order to determine how proficient the peer is with the skill.

5.3 Future Directions

We can think of several improvements to specific components of our methodology. The robot could be trained on using a skill for a specific task, and then be presented with a different task that combines aspects of the skill with the robot's existing skillset in a new manner. For example, a robot trained to catch falling objects could be taught how to pitch a fastball and then be tasked with catching a ball thrown by the human. The robot could also be trained to adapt to their human partners' preferences and make its next intended action very clear to the human, to make their collaboration fluent.

Though our case study did not use the *Teaching Interaction* to improve the robot's teaching abilities, we see this as a promising opportunity for future work. The robot could try a variety of teaching strategies for a given task to learn which work best for explaining that specific skill, using student performance and feedback as a reward. The robot could also be designed to identify when its own task knowledge is insufficient to successfully teach the skill. The robot would then know to seek an expert to receive more instruction to fill in the gaps in it's understanding. We see this direction as another promising opportunity for future work.

6 CONCLUSION

In this paper, we argue for the need to create robots capable of teaching and learning. We took inspiration from the "Learning by Teaching" (LdL) educational theory and the medical adage "See one, Do one, Teach one" to create a methodology that incorporates three components: learning, collaborating, and teaching. This required deriving insights from various fields, including learning from demonstration, human-robot collaboration, and intelligent tutoring systems. A robot system that utilizes our methodology gains significant advantages, including new and robust evaluation metrics, and the ability to transfer knowledge representations across temporal and geographic separation.

We implemented a proof-of-concept case study in which a robot learns three-note harmonies from an expert in the skill, collaboratively plays the harmonies with the expert, and finally teaches the same harmonies to a novice user until the robot's internal knowledge tracing models conclude that the novice has a sufficient understanding of the skill. We then had the novice teach what they had learned back to the robot to further improve the novice's learning gains evaluate the robot learner.

The three components of learning, collaboration, and teaching can be combined multiple times and in different arrangements, allowing for more robust robot learners and efficient robot tutoring systems. We recognize that such a system requires an explainable learning model which is challenging to design and implement. However, due to their generalizability to various task models and human-robot settings, we argue that for many applications — particularly for robot tutors and collaborators — the difficulties are outweighed by the benefits of building such systems.

ACKNOWLEDGMENTS

This work was partially funded by the National Science Foundation (NSF) under grants No. 2033413, 1955653, 1928448, 1936970, 1813651. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

REFERENCES

[1] Ahmed Gasm Alseed Ahmed. 2013. The impact of learning by teaching strategy on university students' linguistic competence. *University of Bakht Alruda Scientific Journal* 9 (2013), 185–200.

[2] Lawrence M Aleamoni. 1972. Illinois Course Evaluation Questionnaire (CEQ) Results Interpretation Manual Form 66 and Form 32. (1972).

[3] Lawrence M Aleamoni and Richard E Spencer. 1973. The Illinois course evaluation questionnaire: A description of its development and a report of some of its results. *Educational and Psychological Measurement* 33, 3 (1973), 669–684.

[4] Alnour Alharin, Thanh-Nam Doan, and Mina Sartipi. 2020. Reinforcement Learning Interpretation Methods: A Survey. *IEEE Access* (2020).

[5] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. 2009. A survey of robot learning from demonstration. *Robotics and autonomous systems* 57, 5 (2009), 469–483.

[6] Christopher G Atkeson and Stefan Schaal. [n.d.]. Robot learning from demonstration. Citeseer.

[7] Besart Avdiu and Gaurav Nayyar. 2020. When face-to-face interactions become an occupational hazard: Jobs in the time of COVID-19.

[8] Andrea Bajcsy, Dylan P Losey, Marcia K O'Malley, and Anca D Dragan. 2017. Learning robot objectives from physical human interaction. *Proceedings of Machine Learning Research* 78 (2017), 217–226.

[9] Anne Kathrin Barnbeck and Thomas Neumann. 2006. Lernen durch Lehren (LdL) in theory and practice.

[10] Tony Belpaeme, James Kennedy, Aditi Ramachandran, Brian Scassellati, and Fumihide Tanaka. 2018. Social robots for education: A review. *Science robotics* 3, 21 (2018).

[11] Matthew L Bernacki and Candace Walkington. 2018. The role of situational interest in personalized learning. *Journal of Educational Psychology* 110, 6 (2018), 864.

[12] AC Bittner Jr. 1993. Requirements testing and evaluation: Human-in-theloop systems. *Ergonomics in Design* (1993), 29–32.

[13] Ali Borji, Ming-Ming Cheng, Huaizu Jiang, and Jia Li. 2015. Salient object detection: A benchmark. *IEEE transactions on image processing* 24, 12 (2015), 5706–5722.

[14] Gary Bradski. 2000. The OpenCV Library. *Dr. Dobb's Journal of Software Tools* (2000).

[15] C. Breazeal, A. Brooks, D. Chilongo, J. Gray, G. Hoffman, C. Kidd, H. Lee, J. Lieberman, and A. Lockerd. 2004. Working collaboratively with humanoid robots. In *4th IEEE/RAS International Conference on Humanoid Robots, 2004.*, Vol. 1. 253–272 Vol. 1.

[16] Cynthia Breazeal, Guy Hoffman, and Andrea Lockerd. 2004. Teaching and working with robots as a collaboration. In *AAMAS*, Vol. 4. 1030–1037.

[17] Alessandro Bruno, Francesco Gugliuzza, Roberto Pirrone, and Edoardo Ardizzone. 2020. A Multi-Scale Colour and Keypoint Density-Based Approach for Visual Saliency Detection. *IEEE Access* 8 (2020), 121330–121343.

[18] Richard W Byrne and Anne E Russon. 1998. Learning by imitation: A hierarchical approach. *Behavioral and brain sciences* 21, 5 (1998), 667–684.

[19] Maya Cakmak, Crystal Chao, and A. L. Thomaz. 2010. Designing Interactions for Robot Active Learners. *IEEE Transactions on Autonomous Mental Development* 2, 2 (2010), 108–118.

[20] Maya Cakmak and Andrea L Thomaz. 2012. Designing robot learners that ask good questions. In *2012 7th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 17–24.

[21] Sylvain Calinon and Aude Billard. 2007. Incremental learning of gestures by imitation in a humanoid robot. In *Proceedings of the ACM/IEEE international conference on Human-robot interaction*. 255–262.

[22] Sylvain Calinon, Florent D'halluin, Eric L Sauser, Darwin G Caldwell, and Aude G Billard. 2010. Learning and reproduction of gestures by imitation. *IEEE Robotics & Automation Magazine* 17, 2 (2010), 44–54.

[23] Angelo Cangelosi. 2010. Grounding language in action and perception: From cognitive agents to humanoid robots. *Physics of life reviews* 7, 2 (2010), 139–151.

[24] Crystal Chao, Maya Cakmak, and Andrea L. Thomaz. 2010. Transparent active learning for robots. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 317–324.

[25] Huili Chen, Hae Won Park, and Cynthia Breazeal. 2020. Teaching and learning with children: Impact of reciprocal peer learning with a social robot on children's learning and emotive engagement. *Computers & Education* 150 (2020), 103836.

[26] Jonathan M Coffman, Marina J McConkey, and James Colee. 2020. Effectiveness of video-assisted, self-directed, and peer-guided learning in the acquisition of surgical skills by veterinary students. *Veterinary Surgery* 49, 3 (2020), 582–589.

[27] Albert T Corbett and John R Anderson. 1994. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User modeling and user-adapted interaction* 4, 4 (1994), 253–278.

[28] David Duran. 2017. Learning-by-teaching. Evidence and implications as a pedagogical mechanism. *Innovations in Education and Teaching International* 54, 5 (2017), 476–484.

[29] Logan Fiorella and Richard E Mayer. 2013. The relative benefits of learning by teaching and teaching expectancy. *Contemporary Educational Psychology* 38, 4 (2013), 281–288.

[30] Amirata Ghorbani, James Wexler, James Y Zou, and Been Kim. 2019. Towards automatic concept-based explanations. In *Advances in Neural Information Processing Systems*. 9277–9286.

[31] Debasmita Ghose, Shasvat M Desai, Sneha Bhattacharya, Deep Chakraborty, Madalina Fiterau, and Tauhidur Rahman. 2019. Pedestrian Detection in Thermal Images using Saliency Maps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 0–0.

[32] Barbara Goldschmid and Marcel L Goldschmid. 1976. Peer teaching in higher education: A review. *Higher education* 5, 1 (1976), 9–33.

[33] Goren Gordon, Samuel Spaulding, Jacqueline Kory Westlund, Jin Joo Lee, Luke Plummer, Marayna Martinez, Madhurima Das, and Cynthia Breazeal. 2016. Affective personalization of a social robot tutor for children's second language skills. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*. 3951–3957.

[34] Kelsey P Hawkins, Shray Bansal, Nam N Vo, and Aaron F Bobick. 2014. Anticipating human actions for collaboration in the presence of task and sensor uncertainty. In *2014 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2215–2222.

[35] Guy Hoffman and Cynthia Breazeal. 2004. Collaboration in human-robot teams. In *AIAA 1st Intelligent Systems Technical Conference*. 6434.

[36] Deanna Hood, Séverin Lemaignan, and Pierre Dillenbourg. 2015. When children teach a robot to write: An autonomous teachable humanoid which uses simulated handwriting. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. 83–90.

[37] Mohammad Khajah, Robert V Lindsey, and Michael C Mozer. 2016. How deep is knowledge tracing? *arXiv preprint arXiv:1604.02416* (2016).

[38] George Konidaris, Scott Kuindersma, Roderic Grupen, and Andrew G Barto. 2010. Constructing skill trees for reinforcement learning agents from demonstration trajectories. In *Advances in neural information processing systems*. 1162–1170.

[39] Sandra V Kotsis and Kevin C Chung. 2013. Application of see one, do one, teach one concept in surgical training. *Plastic and reconstructive surgery* 131, 5 (2013), 1194.

[40] Samantha Krening, Brent Harrison, Karen M Feigh, Charles Lee Isbell, Mark Riedl, and Andrea Thomaz. 2016. Learning from explanations using sentiment and advice in RL. *IEEE Transactions on Cognitive and Developmental Systems* 9, 1 (2016), 44–55.

[41] Rensis Likert. 1932. A technique for the measurement of attitudes. *Archives of psychology* (1932).

[42] Zachary C Lipton. 2018. The mythos of model interpretability. *Queue* 16, 3 (2018), 31–57.

[43] Manuel Lopes, Francisco Melo, and Luis Montesano. 2009. Active learning for reward estimation in inverse reinforcement learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 31–46.

[44] Nichola Lubold, Erin Walker, and Heather Pon-Barry. 2016. Effects of voice-adaptation and social dialogue on perceptions of a robotic learning companion. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 255–262.

[45] Prashan Madumal, Tim Miller, Liz Sonenberg, and Frank Vetere. 2019. Explainable reinforcement learning through a causal lens. *arXiv preprint arXiv:1905.10958* (2019).

[46] Prashan Madumal, Tim Miller, Liz Sonenberg, and Frank Vetere. 2020. Distal Explanations for Explainable Reinforcement Learning Agents. *arXiv preprint arXiv:2001.10284* (2020).

[47] Ali Ahmad Malik, Tariq Masood, and Rehana Kousar. 2020. Repurposing factories with robotics in the face of COVID-19. (2020).

[48] Neville Mehta, Sriraam Natarajan, Prasad Tadepalli, and Alan Fern. 2008. Transfer in variable-reward hierarchical reinforcement learning. *Machine Learning* 73, 3 (2008), 289.

[49] Bilge Mutlu, Allison Terrell, and Chien-Ming Huang. 2013. Coordination mechanisms in human-robot collaboration. In *Proceedings of the Workshop on Collaborative Manipulation, 8th ACM/IEEE International Conference on Human-Robot Interaction*. Citeseer, 1–6.

[50] Chris Piech, Jonathan Bassen, Jonathan Huang, Surya Ganguli, Mehran Sahami, Leonidas J Guibas, and Jascha Sohl-Dickstein. 2015. Deep knowledge tracing. In *Advances in neural information processing systems*. 505–513.

[51] Mattia Racca and Ville Kyrki. 2018. Active robot learning for temporal task models. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. 123–131.

[52] Aditi Ramachandran, Alexandru Litoiu, and Brian Scassellati. 2016. Shaping productive help-seeking behavior during robot-child tutoring interactions. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 247–254.

[53] Harish Ravichandar, Athanasios S Polydoros, Sonia Chernova, and Aude Billard. 2020. Recent advances in robot learning from demonstration. *Annual Review of Control, Robotics, and Autonomous Systems* 3 (2020).

[54] Alessandro Roncone, Olivier Mangin, and Brian Scassellati. 2017. Transparent role assignment and task allocation in human robot collaboration. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 1014–1021.

[55] Paul E Rybski, Kevin Yoon, Jeremy Stolarz, and Manuela M Veloso. 2007. Interactive robot task training through dialog and demonstration. In *Proceedings of the ACM/IEEE international conference on Human-robot interaction*. 49–56.

[56] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*. 618–626.

[57] Shuvro Sen, Neel Antara, Shusmita Sen, and Sunny Chowdhury. 2020. The apparel workers are in the highest vulnerability due to COVID-19: a study on the Bangladesh Apparel Industry. *Sen S, Antara N, Sen S, Chowdhury S. The apparel workers are in the highest vulnerability due to COVID-19: a study on the Bangladesh Apparel Industry. Asia Pacific J Multidiscip Res* 8, 3 (2020).

[58] Julie Shah, James Wiken, Brian Williams, and Cynthia Breazeal. 2011. Improved human-robot team performance using chaski, a human-inspired plan execution system. In *Proceedings of the 6th international conference on Human-robot interaction*. 29–36.

[59] Jaeeun Shim and Andrea L Thomaz. 2011. Human-like action segmentation for option learning. In *2011 RO-MAN*. IEEE, 455–460.

[60] Gustavo Taboada Soldati. 2015. Knowledge transmission: the social origin of information and cultural evolution. In *Evolutionary ethnobiology*. Springer, 89–104.

[61] Martin Stigmar. 2016. Peer-to-peer teaching in higher education: A critical literature review. *Mentoring & Tutoring: partnership in learning* 24, 2 (2016), 124–136.

[62] Fumihide Tanaka and Shizuko Matsuzoe. 2012. Children teach a care-receiving robot to promote their learning: Field experiments in a classroom for vocabulary learning. *Journal of Human-Robot Interaction* 1, 1 (2012), 78–95.

[63] Neil Thomason. 1990. Making Student Groups Work:"To teach is to learn twice". *Teaching Philosophy* 13, 2 (1990), 111–125.

[64] Judy CR Tseng, Hui-Chun Chu, Gwo-Jen Hwang, and Chin-Chung Tsai. 2008. Development of an adaptive learning system with two sources of personalization information. *Computers & Education* 51, 2 (2008), 776–786.

[65] Abhinav Verma, Vijayaraghavan Murali, Rishabh Singh, Pushmeet Kohli, and Swarat Chaudhuri. 2018. Programmatically interpretable reinforcement learning. *arXiv preprint arXiv:1804.02477* (2018).

[66] Ricardo Borges Viana, Mário Hebling Campos, Douglas de Assis Teles Santos, Isabela Cristina Maioni Xavier, Rodrigo Luiz Vancini, Marília Santos Andrade, and Claudio Andre Barbosa de Lira. 2019. Improving academic performance of sport and exercise science undergraduate students in gross anatomy using a near-peer teaching program. *Anatomical sciences education* 12, 1 (2019), 74–81.

[67] Joshua Wainer, David J Feil-Seifer, Dylan A Shell, and Maja J Mataric. 2006. The role of physical embodiment in human-robot interaction. In *ROMAN 2006-The 15th IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, 117–122.

[68] Joshua Wainer, David J Feil-Seifer, Dylan A Shell, and Maja J Mataric. 2007. Embodiment and human-robot interaction: A task-based perspective. In *RO-MAN 2007-The 16th IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, 872–877.

[69] Victoria Weiss and Robert Needlman. 1998. To teach is to learn twice: resident teachers learn more. *Archives of pediatrics & adolescent medicine* 152, 2 (1998), 190–192.

[70] Michael V Yudelson, Kenneth R Koedinger, and Geoffrey J Gordon. 2013. Individualized bayesian knowledge tracing models. In *International conference on artificial intelligence in education*. Springer, 171–180.

[71] Cristina Zaga, Manja Lohse, Khiet P Truong, and Vanessa Evers. 2015. The effect of a robot's social character on children's task engagement: Peer versus tutor. In *International Conference on Social Robotics*. Springer, 704–713.