# Novice Human Teachers of a Virtual Toddler: A Case Study

SUB #558


## Abstract

While home robotic agents have been successful in learning or adapting to simple tasks (such as vacuuming), often they can perform only preordained tasks and cannot be easily trained by a non-expert. This study is part of a larger project that with the goal of addressing this challenge by creating an interface and feature set for a child-like virtual robot automated learner teachable by human teachers (HTs) with no prior training. The current study contributes to this project by exploring the ways in which the HTs interact with our virtual robot automated learner (VRAL). Through the study, we found that humans tended to make a repeated and reliable set of assumptions about the learner, and that these assumptions affected how the HTs attempted to teach VRAL. Past research has shown that humans teaching robots (or agents) have three reliable tendencies: they attempt to infer a model of learner; they exhibit bias towards positive reinforcement; and they direct the agents attention. This study suggests that there are two more tendencies: humans do not ask the robot to explore and humans assume much about the nature of a robot before evaluating it. We explore the evidence of and reasons for these tendencies.


## 1 Introduction

There are myriad benefits to teaching robots or agents using human teachers (HTs), and we see them every day. For example, reliable, logical virtual learners have been attempted by a variety of websites (such as spiders, automated shopping agents, etc), but the virtual agents have often been too simplistic for untrained humans to meaningfully communicate complex tasks to them. While home robotic agents (such as the Roomba [iRobot Corp., 2008] or the AIBO [Sony Corp., 2008]) have seen success in learning or adapting to simple tasks (such as vacuuming), they can perform only preordained tasks and cannot be easily trained by a non-expert. This study is part of a larger project that with the goal of addressing this challenge by creating an interface and feature set for a child-like virtual robot automated learner (hereafter called "VRAL" for "Virtual Robot Automated Learner") teachable by human teachers (HTs) who have no prior training.

The current study contributes to this project by exploring the ways in which the untrained HTs interact with VRAL, the virtual robot automated learner. The HTs were asked to produce a lesson plan for the automated learner prior to teaching it and then to enact that lesson with VRAL. Through the study, we found that humans tended to make a repeated and reliable set of assumptions about the learner, and that these assumptions affected how the HTs taught the automated learner.

The term "virtual robot" has had many definitions and correspondingly mixed success as a tool for designing AI learners [Maes, 1990]. In this paper, we define "virtual robot" to mean an agent that acts in a simulation of the physical world. In other words, a virtual robot is an imperfect analogue of a physical robot but is designed to appear, act, and have similar capabilities to existing physical robots. Virtual robots have the broad benefit of being easily testable since the simulation is virtual, it is necessarily portable and repeatable. Due to the vagaries of physicality, a physical robot cannot exactly repeat actions, as action on the physical world necessarily changes the physical world. Virtual robots also have the benefit of being able to perform a variety of virtual tasks, and algorithms designed for physical robots can often be ported relatively easily to a virtual robot, assuming the simulations model of the world is sufficiently complex. Furthermore, the virtual robot is a teachable concept for a novice (e.g. [Berland and Wilensky, 2005]).

Instances of novices teaching virtual robots are relatively rare, but we found that our participants (the HTs) had little trouble making sense of this task. In many representations of robots in culture (such as C3PO, the Terminator, Data, and Robby the Robot), humans interact with and teach the robots without any specialized language or training. Furthermore, the training of robots by humans in natural settings has long been a "gold standard" in robotics research (since, at least, [Weiner, 1965]), though our study is relatively limited. That said, training with novice human teachers introduces a variety of problems: interface problems are pervasive; their language is usually imprecise; and they tend not to selectively introduce only relevant data, as expert trainers do. The novice teachers also introduce their own idiosyncrasies and biases into the project. For instance, some teachers tried to speak to VRAL using something akin to baby talk, in part, which obfuscated important data.

Thus, this study explores this challenge: the goal of creating a robot that novice participants can successfully teach requires understanding how these novice users make sense of the task. How do they structure their instructions? What assumptions about the robots knowledge and abilities do they make? What challenges do they have?

## 2 VRAL

There exists a wide range of complex tasks possible for a virtual robotic automated learner. Our work focuses on a small subset: VRAL took the form of single claw arm in a virtual "blocksworld". This subset has the benefit of being restricted enough to be tractable, and the "blocksworld" domain is well established as an AI test-bed [Winograd, 1972]. VRAL is said to be approximately equivalent to an intelligent 2-year-old human child because that is an age at which children can learn new complex tasks relatively quickly but are not verbally or spatially expert.

In this study, we had five humans teach VRAL to create a post-and-lintel ("doorway") structure using randomly placed sets of both large and small blocks. VRAL had five basic actions that could be combined in order to accomplish this goal:

RAISE LOWER GRAB RELEASE MOVE TO

These actions could be combined to make complex actions. For example, the actions: grab, move to, and release are all necessary to move a block from one part of the world to another.

In this paper, we will not describe VRAL itself because the design of VRAL is not inherent to the project of investigating how users interact with VRAL. In this particular study, the natural language and interface aspects of VRAL were handled partly by a novel design after Cyc [Lenat and Guha, 1990] and [Thomaz and Breazeal, 2008] and assisted by a human researcher ("wizard") who would interpret statements that the system had trouble parsing. All statements were communicated via instant message to VRAL to simplify logging.

## 3 Description of Study

In order to explore the ways in which novice users teach a virtual robot, we recruited five human teachers (2 undergraduates, 2 researchers, and a teacher). These HTs were asked to create a lesson plan to teach a virtual robot how to build a doorway out of basic blocks. None of the participants had used the system beforehand however, most of them had prior experience in engineering. They were told that they would be teaching the robot as if it were a 2-year-old child. Each participant emailed a lesson plan to the researchers prior to teaching VRAL.

In the enactment, the teachers were ushered into a room with a researcher, VRALs wizard and a computer that was set up to communicate with VRAL. VRAL was projected onto a screen at the front of the room, visible by the three humans (researcher, teacher, and wizard). The HT was instructed to give his or her instructions directly to VRAL; the wizard sat at the back of the room and helped clarify the HTs instructions, as necessary.

The enactment was roughly one hour per teacher. The teachers attempted to enact their lesson plans, and modified them accordingly with the progress of the lesson. In the end, all teachers were successful in teaching VRAL the process of making a post-and-lintel doorway out of random blocks.

Human interaction data was collected in the form of instant message transcript (teacher interaction with VRAL), actions taken by VRAL, and field-notes by the researcher. The transcripts varied greatly in their tenor and in the ways that the teachers approached the subject. In the section below, we will analyze the ways that the transcripts varied and categorize the patterns in the HTs interactions with VRAL.

## 4 Data

VRAL learned how to make a post-and-lintel doorway from all 5 HTs. For this study, we define "learning" as the ability to perform procedures in a variety of complex and novel contexts. VRAL is said to "learn" if it can reliably reproduce an abstract instruction that has not been explicitly programmed. For instance, if VRAL can produce a doorway in every randomized blocksworld configuration (when asked to do so), it is said to have "learned" to make a doorway. The HTs approaches to teaching VRAL ranged from configuration-specific (i.e., using a specific configuration of blocks) to generic (i.e., would work with any configuration where no blocks were stacked on each other but laid out on the table in random order). In this section, we organize our discussion of the data around four patterns that emerged through these case studies.

### 4.1 HTs used a bottom-up approach to teaching VRAL

Four of the five participating HTs taught VRAL by providing parameterized definitions of basic actions, slowly building up the composite actions necessary for completing the desired task. Table 1 provides an example of how one HT defined one action – placing one block on top of another – necessary for building a doorway. He then asked VRAL to use this action in combination with others – that were defined using a similar process – to build a doorway.

That four of the five participants used these parameterized procedures is striking when compared to teaching strategies with young children. That is, young children are often taught through a more top-down approach in which their instructor identifies goals and then asks the students to explore the available resources in order to achieve those goals. In this case, an instructor might draw a doorway and ask the young children to build something like that with blocks. However, the HTs were unable to teach VRAL in this manner and used a bottom-up approach instead. That is, the HTs provided VRAL with detailed procedures to follow.

In fact, even those HTs that began their teaching session with a more top-down approach, generally abandoned it in favor providing direct instructions for performing specific procedures (e.g., a bottom-up approach). For example, HT 4 took the most aggressive approach in "teaching" VRAL treating it

Table 1: Transcript Excerpt from HT 5

| | |
|---|---|
| HT > VRAL | first we will learn how to “plop” |
| HT > VRAL | grasp a |
| VRAL > HT | Done |
| HT > VRAL | raise a |
| VRAL > HT | Done |
| HT > VRAL | release a |
| VRAL > HT | I can’t do that. |
| HT > VRAL | lower a |
| VRAL > HT | Done |
| HT > VRAL | release a |
| VRAL > HT | Done |
| HT > VRAL | this is called plop(a),<br>plop(b) would do the same thing<br>as was done to block a to block b |
| HT > VRAL | the definition of plop(block) takes<br>any block as an argument<br>plop(block) is defined as:<br>grasp block<br>raise block<br>lower block<br>release block<br>end<br>plop(h) |
| VRAL > HT | Done |

Table 2: Transcript Excerpt from HT 2

<table>
<tr><td rowspan="5">HT > VRAL</td><td>define action MoveToRightOf</td></tr>
<tr><td>arguments (block, support)</td></tr>
<tr><td>steps<br>MoveClaw(“over”, block)<br>Grasp(block)<br>Raise(block)<br>MoveClaw(“right”, support)<br>Lower(block)<br>Release(block)<br>Done</td></tr>
<tr><td>define predicate<br>isRightOf(support, support)</td></tr>
<tr><td>define function<br>support getRightOf(support)</td></tr>
</table>

as a planner and working to help VRAL determine the steps necessary for accomplishing the goal. That is, at the beginning of the session, HT 4 was using a top-down approach in which he explained the general overall goals and then broke them down into component pieces and pushing down to the detailed level. However, when encountering a challenge with teaching VRAL to “choose” an arbitrary block he resorted to providing VRAL with parameterized procedures to followsimilar to other teachers in this study.

### 4.2 All the HTs overestimated what the student knew or could do.

Obviously, no current system has the ability to learn like a child would learn. Human intelligence, even at a young age, is formidable. As a result, even though the wizard was translating instructions for the automated learner, there are many complex statements that were too difficult for VRAL. For example, after finding out that VRAL had no notion of “choose”, HT 3 defined an algorithm for choosing a block (in this case beginning at the left and searching to the right for a block). Moreover, the mechanism for choosing had to be refined to take into considerations specific characteristics of the blocks (size, whether it has a block on top of it, etc). Thus, HT 3 had to repeatedly simplify his instructions to VRAL in order to accommodate the systems familiarity with these concepts.

In addition, it is very hard to assess learning and achievement in any human, especially a young child; a human teacher teaching a human child will probably not know whether the child has reliably learned the lesson. This is much more easily evaluated with VRAL by simply asking it to perform an explicit task and evaluating the result in blocksworld.

### 4.3 Many of the HTs employed repetition and mnemonics when teaching VRAL

Where repetition and reinforcement would be a standard teaching technique with a human student, it was not clear to the HTs whether such repetition is useful at all with VRAL. Due to the structure of VRAL, memory was not an issue. Repetition and mnemonic techniques (such as funny names) are very common with young children but not at all helpful with VRAL. Many of the lessons employed repetition and mnemonics. HT5, for example, used humorous mnemonic names for blocks, though that was not relevant to VRAL.

### 4.4 The HTs differed in their assumptions regarding the linguistic capability of VRAL

The language of the teachers ranged from predicate logic to natural language. HTs made markedly different assumptions about the linguistic capability of the learner. In actuality, the wizard could translate any natural language into usable instructions for VRAL. Nonetheless, four out of five of the teachers used a code-like language to teach VRAL. To demonstrate this variation, compare Tables 2 and 3 in which HTs 2 and 3 are using very different language to teach VRAL.

## 5 Conclusions

Thomaz and Breazeal [2008] found that humans teaching robots (or agents) had three reliable tendencies: they attempt to infer a model of learner; they exhibit a bias towards positive reinforcement; and they attempt to direct the agents attention. This study suggests that there are at least two more tendencies: humans do not ask the robot to explore; and humans assume a lot about the nature of a robot before evaluating it. Learning theorists (such as [Piaget, 1928]) argue that humans learn by exploring new situations/problems and connecting these novel contexts to familiar ones. However, no HT made that assumption about VRALs learning in this study. Indeed, no HT created a space for VRAL to do anything other than follow simple instructions. Moreover, that the majority of the HTs used a pseudo-code like language when interacting

Table 3: Transcript Excerpt from HT 3

| | |
|---|---|
| HT > VRAL | let's define what it means to push a block. To push a block y that is in position x one square to the right, moveblocktolocation(y, x+1). |
| | A stack consists of three blocks. One is on the slab. A second one is on top of that first one. The third one is on top of the second one. |
| | To push a block y that is in position x one square to the left, moveblocktolocation(y, x-1). |
| | To build a stack, do three things. First, choose a block x that is sitting on the slab. |
| | Next, choose a block y that is also sitting on the slab. MoveClaw(over, y). Then grasp(y). Then MoveClaw(over, x). Then Release(y, x). |
| | Third, choose a block z that is also sitting on the slab. MoveClaw(over, z). Then grasp(z). Then MoveClaw(over, y). Then Release z, y). |
| | Build a stack. |

with VRAL suggests that they were focused on the computational ability of the program rather than considering it as a learner. The one interesting contrast to these findings were the instances in which the HTs used repetition and mnemonics to teach VRAL – a tendency that is necessary for young children but less so when working with a computer. As the participants were not robotics researchers, this may have to do with some outside knowledge about interactive robotics, or it could be bias. Either way, it seems fairly reliable. If we are to teach automated learners as if they are humans, we need the teachers to believe that they can be taught in a more reliably human way.

## References


[Berland and Wilensky, 2005] M. Berland and U. Wilensky. Complex play systems: Results from a classroom implementation of vbot. In *The annual meeting of the American Educational Research Association*, Montreal, 2005.

[iRobot Corp., 2008] iRobot Corp. iRobot Roomba: Vacuum cleaning robot, 2008.

[Lenat and Guha, 1990] D.B. Lenat and R.V. Guha. *Building Large Knowledge-Based Systems*. Addison-Wesley, Reading, MA, 1990.

[Maes, 1990] P. Maes. *Designing Autonomous Agents*. MIT Press, Cambridge, MA, 1990.

[Piaget, 1928] J. Piaget. *The Child's Conception of the World*. Routledge and Kegan Paul, London, 1928.

[Sony Corp., 2008] Sony Corp. AIBO Entertainment Robot, 2008.

[Thomaz and Breazeal, 2008] L. Thomaz and C. Breazeal. Teachable robots: Understanding human teaching behavior to build more effective robot learners. *Artificial Intelligence*, 172:716–737, 2008.

[Weiner, 1965] N. Weiner. *Cybernetics: or the Control and Communication in the Animal and the Machine*. MIT Press, Cambridge, MA, 1965.

[Winograd, 1972] T. Winograd. *Understanding natural language*. Academic Press, New York, 1972.