Dorothy DeLay Meets Rosie the Robot

(Legendary violin teacher Dorothy DeLay, of Julliard, meets Rosie, the cartoon robot from the Jetsons…. *sort of*)

Ron Steiner

Learning Technologies

University of North Texas

Maymester 2016

CECS 6300

May 20, 2016

If we imagine artificial intelligence (AI), in the year 2016, as a discipline still in its infancy, having been born in 1956 at the Dartmouth conference where John McCarthy first used the term, then perhaps we should look to an educator of children for insight on how best to teach computers to learn, especially in the realm of common sense. Because until AI develops common sense that is indistinguishable from that of a human being, it will not be thought of as truly intelligent.

In her 1949 book "The Absorbent Mind," of which I obtained a 1967 publication, Maria Montessori refers to a child's "Periods of Growth" (Montessori 1967, chapter 3), saying that development is a series of rebirths. She describes the earliest period, from birth to six years of age, where the child has "a type of mind that the adult cannot approach, that is to say , we cannot exert upon it any direct influence." The child mostly reacts and absorbs. "An unconscious mind can be most intelligent" (Montessori 1967, p. 23). The child mimics and continues to absorb. The child has no common sense in the sense that they must be taught everything. I submit that this is where we are with computers and AI, albeit 60 years after the birth of the term. Computers may appear smart, in the sense that they can compute quickly and can retrieve well-structured data efficiently. But that's just appearance. And they must be taught (read: programmed) to do every single thing that they do. And they must be programmed by people who actually are intelligent and have common sense. The computer can retrieve information when programmed to do so but without additional programming, and hardware, cannot recognize the face of its programmer (mother/father?) . And without still more programming, and more hardware, they cannot hear. So while they can appear intelligent, knowing exactly how many

miles long the Great Wall of China is or what is the precise value of Avogadro's Number, they cannot synthesize or use any of those facts without guidance from a human programmer.

In short, the computer lacks even the most basic level of common sense. It must be told everything. A child is born with an innate fear of loud noises. The computer doesn't even have that, unless it has been programmed for that specific task (and its associated reaction, since all reactions must also be programmed). It's not even intelligent enough to be afraid of loud noises.

At a minimum, for a computer to mature, in the Montessori sense of the term, denoting "a kind of regulating mechanism for growth, which ensures balance in the organism as a whole…" (Montessori 1967, p. 93), the computer will have to learn (i.e. be programmed with) a level of common sense that permits an understanding of or a recognition of itself as an organism that has the capacity and need for additional learning. At present, the computer is static and unless programmed by a human, will remain in one state forever. A child, even a toddler (especially!), will never remain in one place for very long. I suggest that the child's motion, exploration, fear of loud noises, and capacity to recognize its mother are all elements of common sense, none of which the computer approaches without significant effort by a programmer.

So for the computer to have a level of common sense, they must be programmed for it. But that doesn't mean that they cannot exhibit common sensical behaviors, or show the appearance of common sense (or intelligence) without actually having it. Siri and related voice activated technologies have a sense of humor. Ask Siri to tell you a joke and she will. Ask her Android counterpart to make you a sandwich and see what she says. And, for our purposes, the

appearance of common sense, in narrowly defined domains, might be satisfactory, at least until Moore's Law gives us the computing power to program the computer to do more. It is this appearance of common sense that will allow us to use AI in education effectively, especially in predictable circumstances. But it's important to remember that it's just an appearance, an illusion.

> "Almost without exception, current computer programs to carry out language tasks succeed to the extent the tasks can be carried out purely in terms of manipulating individual words or short phrases, without attempting any deeper understanding; commonsense is evaded, in order to focus on short-term results, but it is difficult to see how human-level understanding can be achieved without greater attention to commonsense." (Davis 2015)

For example, if a middle-school student is using a mathematics tutoring system to improve their knowledge of multiplication tables, the designer might reasonable expect that when presented with the problem of 7 times 9, it's possible that the student will get the answer wrong on the first try. The computer tutors, responding with:

> "That's not right. Please try again."

On the second missed attempt the reply might be:

> "I'm sorry but that's still not correct. Please try again."

And on the third incorrect try the reply might be:

> "Let's see if we can find the teacher to get some help on this one."

But if on the third try the student got the right answer, a common sense reply might very well be:

"You got it! I knew you could do it!"

The computer didn't "know" anything, of course, but by exhibiting this sort of common sense response, the computer endears itself to the child, having expressed confidence in the child's ability, and praise for the child's effort. It's what any tutor with common sense would do and the computer (was programmed….) did it. A context-sensitive and appropriate sense of humor will help AI in educational settings.

Another example of common sense is using more than one sense to gather data about the student's performance of a task. What would it take for a computer to teach a child to play the violin? Both of my daughters began violin lessons at age three. This means a couple of things…. 1) I listened to a LOT of bad violin playing (think cats fighting in a closed burlap sack) and 2) I spent many hours watching teachers teach violin and watching children take those lessons and improve. For a computer (robot?) to teach violin, the computer would need multiple forms of input, because there are multiple factors that affect the violin playing.

The robot would have to watch the child's hands. The grip on the bow, by the bow hand, is crucial to the overall playing and a disproportionate amount of time in early lessons is devoted to the bow hand and the grip. Teachers use a variety of techniques, ranging from putting stickers on the bow to prescribing exercises to strengthen the correct muscles in the bow hand to affect the grip. The robot would have to listen to the tone. I expect that since musical pitch can be measured by a device, indeed my daughters have mechanical tuners with digital readouts of the precise MHz of their note, that the robot would excel in this respect. However it is likely to be a

formulaic or mechanical assessment only, void of any emotion or understanding of the piece being played.

When my older daughter, then age 13, played a violin solo in front of 700 people at a commemorative dinner, I watched grown men in the audience start to cry. I suspect a robot instructor would not have cried, so even with proper programming regarding the expected pitch and duration of every note, the robot instructor would still be missing "something."

The teacher must also pay attention to the ergonomics of the player. Often this is done by the teacher putting a hand on the student's shoulder to feel whether or not there was too much muscle tension. So we've examined the violin teacher's use of vision, hearing, and touch and all the student has done so far is play a single note. Ideally the teacher would communicate their findings, suggestions, compliments, criticisms, and observations verbally. Less frequently they may demonstrate the desired outcome, necessitating that the robot be programmed to play the violin correctly, or retrieve digitally stored representations of the correct outcome, the latter of which seems more likely to me.

It would be reasonable to ask, if AI is in its infancy, is there any indication that a robotic violin tutor such as I have described is even possible? There is some evidence, albeit fragmented, that the answer is yes. For example, a 2016 study by McLeod et al describes the use of wearable sensors on upper extremity prosthetics (sensors worn on the wrist) to measure the use that amputees make of their prosthetics. The key problem the study focused on was differentiating between functional use and extraneous use of the limb. They found that a properly programmed (and properly worn) device could differentiate between functional use vs

extraneous use. This suggests that a properly programmed (and properly worn) device could (eventually) differentiate between a proper and an improper bow movement. This coupled with a properly programmed robotic eye, such as the one described by Carnegie Mellon University (CMU) researchers at the May 17, 2016 IEEE International Conference on Robotics and Automation in Stockholm, Sweden, provides two senses of input to help the teacher correct the bow technique of the student. The ability to incorporate more than one sense (source of input) to craft corrections to student behavior is another positive outcome of common sense in education settings. It is the integration of these multiple senses that will bring about paradigmatic shifts in the use of AI in education settings. It's not enough to have the intelligent tutoring system monotonously reply that an answer is right or wrong, or the pitch of the violin is correct or off by 10 MHz.

For AI to be useful in education it must integrate multiple sources of input and simultaneously be able to suggest multiple steps of corrective action because in the violin example, it may be that the bow hand is precisely correct but the "fiddle hand," as our beloved violin teacher Mrs. Aten called it, is too loose, or too tight, or in the wrong place. So for the automated teacher to replace, or even augment, Mrs. Aten, they must synthesize all of the input data, formulate appropriate corrective action steps, and present the suggestions in a constructive manner, with appropriate scaffolding, encouragement, and discipline

And it would help if the AI knew that it is entirely appropriate to cry upon hearing the theme from Schindler's List played by a 13 year old granddaughter of Holocaust survivors at a Hope for Humanity dinner. (5:10): http://bit.ly/1RceQKA

References

Davis, E., & Marcus, G. (2015). Commonsense reasoning and commonsense knowledge in artificial intelligence. *Communications of the ACM Commun. ACM, 58*(9), 92-103. doi:10.1145/2701413

Mcleod, A., Bochniewicz, E. M., Lum, P. S., Holley, R. J., Emmer, G., & Dromerick, A. W. (2016). Using Wearable Sensors and Machine Learning Models to Separate Functional Upper Extremity Use From Walking-Associated Arm Movements. *Archives of Physical Medicine and Rehabilitation, 97*(2), 224-231. doi:10.1016/j.apmr.2015.08.435

Montessori, M., & Claremont, C. A. (1967). *The absorbent mind*. New York, NY: Dell.

Moor, J. H. (2006). The Dartmouth College artificial intelligence conference: The next fifty years. *AI Magazine*, 27(4), 87-91

Robot's In-Hand Eye Maps Surroundings, Determines Hand's Location-CMU News – Carnegie Mellon University. (n.d.). http://www.cmu.edu/news/stories/archives/2016/may/robot-hand-camera.html

Appendix
(excerpt of obituary)

March 26, 2002

Dorothy DeLay, Teacher of Many of the World's Leading Violinists, Dies at 84

By ALLAN KOZINN http://www.nytimes.com/2002/03/26/arts/dorothy-delay-teacher-of-many-of-the-world-s-leading-violinists-dies-at-84.html?pagewanted=all&pagewanted=print

Dorothy DeLay, one of the world's most famous teachers of the violin and a mentor to two generations of players ranging from Itzhak Perlman to Midori and Sarah Chang, died on Sunday. She was 84 and lived in Upper Nyack, N.Y.

Miss DeLay, as her students continued to call her even after they became star performers, began her teaching career almost as an afterthought in 1947, when she was a student of Ivan Galamian at the Juilliard School. Accepting a handful of invitations for part-time teaching and assistantships at the Henry Street Settlement, the Juilliard School and Sarah Lawrence College led her to realize that she enjoyed teaching more than she enjoyed performing.

For more than 20 years she worked largely in Galamian's shadow, but in the 1970's she became a sought-after teacher in her own right, and became the first woman -- and the first American-born violnist -- to be regarded as a master violin teacher in the tradition of Galamian and Leopold Auer.

The list of students who became famous soon after leaving her studio quickly blossomed starting in the 1970's. In addition to Mr. Perlman, Midori and Ms. Chang, they include Nadja Salerno-Sonnenberg, Nigel Kennedy, Gil Shaham, Shlomo Mintz and Cho-Liang Lin. Also among her pupils are members and often founders of some of the world's great chamber groups, among them the Juilliard, the Tokyo, the Cleveland, the Vermeer, the Takacs and the Ying Quartets. Others like Joseph Swenson and Peter Oundjian became conductors and many became concertmasters and section players in orchestras around the world as well as members of conservatory faculties. At the Juilliard School, the faculty includes 14 of her former students.

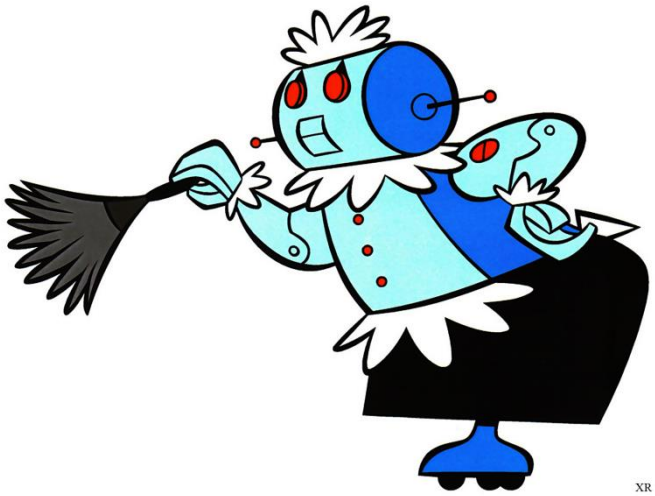

XR