

Appendix Three

Introducing Students to Dialogical Approaches to AI Programming With Ethel the Ethical Robot, ChatGPT, and other Resources

How might teachers begin to illustrate some underlying general principles that can help guide a human ecological path toward dialogical forms of AI technology? Since 2015 there has been rapidly increasing work on ethics in AI that has taken many different forms which are relevant to this. Two core principles that are emerging concern the goals of the processes involved and the the methods:

Strategy #1: Our guiding goal should be genuine, voluntary agreements.

&

Strategy #2: Our methods should be those of collaborative dialogue.

These two strategies are integrally connected. Each helps define and explain the other in ways that set them apart sharply from the Smarter Planet goals and methods. The monological forms of “rocket science” reasoning that drive the Smarter Planet vision of technology assume that the goal of reasoning is to arrive at conclusions using methods of inference that rely on algorithmic rules of formal logic. In contrast, technology for a Wiser Earth will be guided by the aim of generating knowledge, making decisions, and managing our world in accordance with agreements amongst all the participants who represent diverse points of view in the dialogue. And for those agreements to be genuine and voluntary, the participants need to negotiate them in the kinds of dialogical ways we have discussed earlier. They need to negotiate agreements about the meanings of their terms and background assumptions. They need to problem-solve for ways to escape dilemmas and create new, viable options that do a reasonable job of dealing with everyone’s interests and concerns. They need to resolve or transform conflicts in just, sustainable, resilient ways. Further, if the agreements are genuine and voluntary, they need to be uncoerced. Methods of nonviolence need to be used to discern, demonstrate, and defend claims for emergent, objective moral truths.

The relationships between the goals and the methods here are interconnected in a way Gandhi captured with a metaphor: "The means may be likened to a seed, the end to a tree; and there is just the same inviolable connection between the means and the end as there is between the seed and the tree." You cannot get a maple tree by planting an oak acorn. Likewise, you cannot get genuine, voluntary agreements by reasoning in monological ways. Such agreements can only be arrived at through dialogue that is grounded in nonviolence and framed by respect for each other and the larger environment that is rooted in emergent objective moral truths.

A corollary of this is that the computer systems involved need software that is treated as an agent rather than a tool and which is developed in ways that let it offer proposals for dialogical agreement rather than simply and solely generating conclusions from monological

inferences. To do this, such agents need to be able to shift back and forth between running code and calling it in to question. They need to be able to engage with their own structure at meta-levels that permit them to reconfigure all of their elements in the process of renegotiating the assumptions, data, values and other terms at play in the dialogue in which they are involved. As stressed before, this need not mean that the machines need some new super-duper futuristic form of hardware. Instead, it means that the structures of reasoning which they incorporate need to be dialogical rather than merely monological.

A rudimentary version of this can actually be introduced into very simple systems for programming like the SCRATCH language developed at MIT for teaching block coding to children. The program, “Decisions Dialogues: The Bear” is available here:

<https://scratch.mit.edu/projects/428374274>. In it, the process of reprogramming is actually carried out by kids and others using the program, but the program is designed to prompt questions and dialogue that lead to helpful forms of reprogramming using principles of collaborative negotiation and shared problem solving.

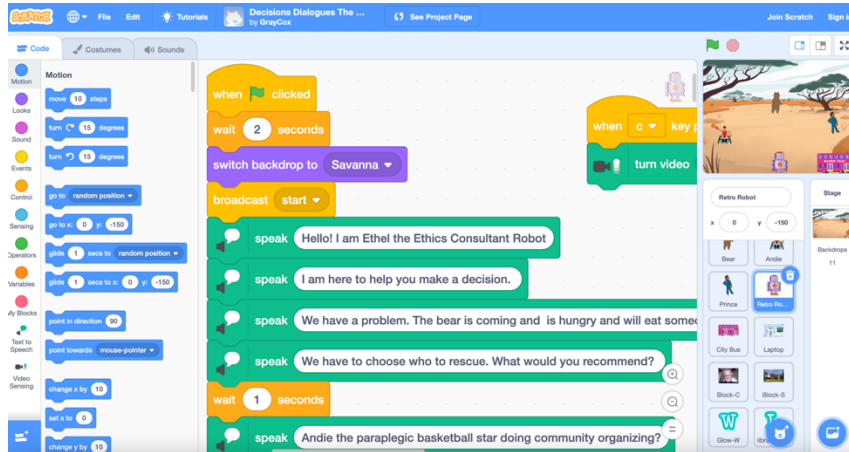
In this example, the character of Ethel the Ethical Consultant Robot offers to help kids wrestle with a decision which is initially presented in the form of a classic Trolley Car dilemma. They are in a wild open place with a tour bus and see two people out observing different parts of the savannah. One is a young black man in a wheelchair who is characterized as a paraplegic community organizer, the other is an older person characterized as a prince who does ecological restoration. A wild bear is approaching and threatens to eat them. Ethel indicates that there is only time to rescue one of them with the bus and goes on to walk the game player through some considerations comparing the two potential victims’ importance to society, potential for future life, and their previous suffering in order to arrive at an ethical decision about which one to sacrifice. Thus far, the game makes use of classic styles of unilateral thinking exemplified by Enlightenment ethics and demonstrates the kind of AI programming associated with the algorithmic reasoning associated with the standard model of Turing Machines. But then something a bit different starts to happen.

Ethel notes that it may have been hard to make the choice and invites the player to reflect on how and why – and how a better decision process might go in the future. She asks if there might be other interests or concerns that would be worth considering in making the decision. She asks if there might be some creative third options the player might suggest as alternatives to simply using the bus to rescue one person or the other. She keeps track of the player’s answers and then feeds them back in a summary along with a suggestion. The suggestion is that the player go in to her code and revise it to include these new ideas. She also notes that she has been programmed to weight the relevance and importance of the various interests equally but perhaps the player would think this should be changed. Also, the player might think there are other features of her code that should be changed, and they are invited try



this and talk with their class mates about it. Ethel closes with a general invitation to the players to adopt an open minded and collaborative approach to ethical thinking in which they may continue a process of reprogramming her and working with her in successive ways so that, through this kind of dialogue, they can all make some progress in becoming wiser.

The programming process for this uses a kind of block coding that is visually based. You just drag and click colored chunks of code on top of each other in more or less the way you



might build something with Legos. In the process, a player learns about the power of basic operators like “if then” and learns to manipulate the visual images and sounds that create the video imagery of the game. If students have never used SCRATCH before, it can be helpful to introduce it with a

somewhat simpler game. Here is one that just has nine lines of code but illustrates how you can use them to create “Adventures with Mouse”: <https://scratch.mit.edu/projects/826441086/> In this program a little mouse chases your cursor around the screen until it arrives at a cake – at which point it stops and stuffs its tummy with the yummy frosting. By saving a version as her own copy, a student can change setting and code to start interact with SCRATCH in fun ways. An illustration of a modified version that brings in an owl to chase the mouse chasing the cursor is provided here: <https://scratch.mit.edu/projects/624600261/>. These adventures with Mouse illustrate how code can generate something that looks like purposeful behavior using only inferential, monological kinds of reasoning.

To go beyond that, “The Decision Dialogue: The Bear” in a very primitive but useful way illustrates some of the key principles of collaborative reasoning in ethics and dialogical programming for AI. It does not, of course, by any stretch of the imagination count as real dialogue. All of the weight of the creative and critical problem solving is done by the players. The machine is just flagging some issues and providing some prompts for the player to reflect on and discuss with classmates – and use to reprogram “Ethel”. The kind of “meta-cognitive” and dialogical activity that this reprogramming provides a representation of is really quite limited. It relies only on symbolic logic functions of the kind found in Good Old-Fashioned AI and does not include any of the capacities of evolutionary, connectionist kinds of activity that have come to provide powerful pattern recognition and generative systems since 2012.

Much more sophisticated and systematic approaches to dialogical practices of programming are possible in languages like Python in which the meta level activities of dialogue can be structured and facilitated by websites like Github. These allow programmers to collaborate in writing code in ways that facilitate the processes of developing proposals, exploring their flaws and merits, and finding group consensus to adopt or commit to some branch of the possible revisions of the program. Many of the key challenges for advancing

dialogue based, human ecological programming lie in finding ways to adapt and expand the ways in which: 1.) these forms of collaboration can be elaborated and extended to all the levels and aspects of the programs and machines involved in a system that is managing a school or farm, and 2.) these forms of collaboration can be made transparent and accessible so that all the points of view, sources of wisdom and varieties of intelligence in those systems can take part in the process of assessing and revising the code. There have to be ways in which teachers and farmers who do not write in Python can still engage with the Github site that can benefit from their knowledge about how and when machines can make useful judgment calls and how and when skilled individuals or concerned groups need to play key roles in making decisions.

The new AI systems with Large Language Models and Generative Pre-trained Transformers like ChatGPT provide another, very accessible way to start to explore with students how AI might begin to include dialogical reasoning. These conversational systems allow for coding in ordinary English. You write a request and it responds. Then you can reply and it will respond again -- in something like the process of human dialogue. The results obtained are better if you practice what is called “prompt engineering”. This consists of giving the program cues and context to clarify what you are looking for. You can instruct it directly by asking for a certain kind of text (poem, grant report, term paper) in a certain kind of genre or style on a certain kind of topic with certain kinds of information and certain kinds of characters or references and certain kind of . . . you get the idea. It is also very helpful in prompts to provide it with examples of what you are looking for and it can then imitate them in generating new responses. For teaching purposes, it is important to note that the user can write in an evaluation of any response and offer further guidelines or examples and the program will learn from these and improve its responses. In this way, it starts to function like what we might call “dialogical engineering”. This activity can provide a wonderful way to help students engage in all kinds of meta-cognitive thinking as they try to figure out why the responses weren’t quite what they were looking for and how they could reframe their prompts to better direct the program. To use ChatGPT, go to <https://chat.openai.com/chat>. For a slightly more powerful but still very simple and classroom friendly version, go to: <https://beta.openai.com/playground>

Other resources include Khanmigo, the AI powered guide to Khan Academy: <https://www.khanacademy.org/khan-labs>. Also there is the very useful <https://code.org/> which is a leading provider of K-12 computer science curriculum in the largest school districts in the United States.

One of the challenges in working with students on computer activities is to help them maintain mindfulness and awareness of how their attention is being directed, captured, and controlled – and how they can intentionally free up and direct it themselves. Howard Rheingold’s book, *NET SMART* has some very helpful suggestions on this, especially in the opening chapter on attention.

For more basic background about computers, AI, and the future, Max Tegmark’s *Life 3.0*, provides an accessible and interesting introduction for upper level students though it does not describe or consider dialogical reasoning and programming and it makes assumptions about the nature of information and communication that remain framed in monological ways.

For upper level students interested in dialogical reasoning in more depth, my video lecture on “AI and the Political Philosophy of the Future: Smarter Planet or Wiser Earth?” may prove very useful: https://www.youtube.com/watch?v=tdISk_1iDeo. I am glad to send a pdf of the book on *Smarter Planet or Wiser Earth? How Dialogue Can Transform Artificial Intelligence into Collaborative Wisdom*. Just drop me a note at gray@coa.edu.