DE GRUYTER OPEN

REM - Research on Education and Media

Vol. 9, N. 1, Year 2017

ISSN: 2037-0830 – DOI: 10.1515/rem-2017-0002

Theoretical vocabularies and styles of explanation of robot behaviours in children

Edoardo Datteri[a], Luisa Zecca[b]

[a]*Università degli Studi di Milano-Bicocca, Italy, edoardo.datteri@unimib.it*
[b]*Università degli Studi di Milano-Bicocca, Italy, luisa.zecca@unimib.it*

Abstract

How do children describe and explain the behaviour of robotic systems? In this paper, some distinctions between different types of explanations, drawing from the philosophy of science literature, are proposed and exemplified by reference to an activity in which primary school children are asked to describe and explain the behaviour of a pre-programmed Braitenberg-like vehicle. The proposed distinctions are also discussed against other studies drawn from the related scientific literature. A qualitative study has provided insights to further refine the analysis described here, through the introduction of other sub-categories of explanation of robotic behaviours.

Keywords: *educational activities, computer literacy, natural science*

Introduction

Since Seymour Papert's seminal works (Papert, 1980), robots have been increasingly used as teaching tools in a variety of activities aimed at supporting learning of various disciplinary and cross-disciplinary abilities and competences. These activities typically involve constructing and programming robotic systems. For example, Church, Ford, Perova, and Rogers (2010) asked high school students to engineer pendulums using LEGO Mindstorms™ components to foster understanding of basic physical concepts and abilities. Whittier and Robinson (2016) report on a study in which middle school students built so-called Evobots, the learning objective of the teacher being the acquisition of evolution-related concepts such as natural selection, adaptation, and niche specialisation. Note that the two last examples, as well as many others that can be found in previous studies (Benitti, 2012; Mubin, Stevens, Shahid, Mahmud, & Dong, 2013), are focussed on the development and acquisition of disciplinary knowledge and abilities concerning particular branches of science (e.g. physics and evolution). Sullivan (2008), however, describes robot-programming activities aimed at the advancement of cross-disciplinary abilities and competences related to scientific thinking – with a special focus on observation, explanation, and experimentation – in middle school students. They were asked to individually implement a line-following program for a robot and to conduct experiments in succession to "debug" the program. Atmatzidou and Demetriadis (2016) report on a robot programming activity aimed at advancing various abilities related to abstraction and generalisation in adolescents.

The acquisition and development of basic cross-disciplinary scientific reasoning abilities and competences is typically promoted in primary school children through observation, explanation, and experimentation on natural systems, such as plants, simple mechanical objects (e.g. pendulums or levers), and insects, in particularly controlled settings. In this paper, a "robo-ethology" activity is described and discussed, which draws inspiration from – and extends – previous studies reported in literature (Datteri & Zecca, 2016; Levy & Mioduser, 2008, 2010; Mioduser & Levy, 2010). In the activity described herein, a robot was presented to children whose behaviour had been previously programmed by the teacher and the researchers to react in peculiar ways to a variety of sensory stimuli, and the children were instructed to treat it as if it were a little robotic animal. They were asked to observe the behaviour of the robot and to propose explanations (as will be clarified later, they had been left free to choose their preferred theoretical vocabulary). Contrary to the examples reviewed earlier, no programming activity was proposed to them. Based on a pilot study, and taking inspiration from Braitenberg's "Vehicles" (Braitenberg, 1986), Datteri and Zecca (2016) suggested that such a robo-ethology game, called the "Game of Science" (GoS), can in principle – i.e. under a methodologically appropriate teacher guidance and in a suitable setting – support the development of cross-disciplinary abilities connected to scientific reasoning, including the ability to observe, to describe one's observations, to propose explanatory hypotheses, to make predictions based on them, to devise experiments, and to reason on the appropriateness of the planned experiments with respect to the hypothesis to be tested.

The educational goal of the experience discussed here was to foster development of these cross-disciplinary abilities in the students. The goal of the game – in other words, the goal that the children were trying to achieve while playing the GoS – was to discover the mechanism actually governing the behaviour of the robot. Note that students' ability to identify the "right" mechanism was not specifically praised by the teachers and the researchers involved in the research. Their learning objectives included the development of scientific thinking abilities, and one may become able to observe, to propose explanations, to devise and conduct experiments, even though s/he is working on a "wrong" hypothesis. Developing scientific thinking abilities is different from developing "good" explanations of a particular phenomenon – and the learning objectives of the GoS may be achieved, in principle, even though no one wins the game itself, i.e. even though no "good" explanation of the behaviour of the robot has been found. This is why little importance was given by the teachers and the researchers guiding the activity to whether the children's hypotheses were on the right track or not. Specific motivations to use an already programmed robot for this kind of observation/explanation activity, in the place of more traditional systems, as a teaching tool for achieving these learning objectives are given in the next section.

The organisation of the activity described in this paper had a research goal too. Datteri and Zecca (2016) identified a number of scientific reasoning abilities exhibited by children during the GoS, which notably include the ability to describe the behaviour of the robot and to propose explanatory hypotheses. The present study aimed at investigating the types of descriptions and explanations of robotic behaviours formulated by the students. As discussed later, a taxonomy of the description and explanation categories had been formulated before the experiment, which extends and refines some distinctions proposed in other studies (Levy & Mioduser, 2008). This taxonomy was used to code the transcriptions of the audio-recorded verbal interactions among children and between the children, the teacher, and the researchers involved in the game. In addition to providing some preliminary figures on the relative predominance of particular types of explanations at different steps of the experiment, the qualitative analysis made in this study provides insights to refine the taxonomy by adding various interesting sub-types of explanations made by the children. The taxonomy is discussed and exemplified in Section 2.

The research activity described here has some educational implications. Better understanding of the types of descriptions and explanations made by children during the observation of a moving robot, in addition to shedding further light on what it means for them to "explain" something, may help teachers plan and conduct robot-assisted, GoS-like activities. For example, knowing that, in particular conditions, children prefer mechanistic to finalistic explanations of robot behaviour, or that mechanistic explanations are typically couched in a psychological versus a physical theoretical vocabulary (as explained later), may help teachers choose appropriate questions for inviting children to reflect on the difference between the two styles of explanation or on their reasons to attribute mental mechanisms to the system. Moreover, more generally, the taxonomy sketched here may be useful to understand children's perception of (educational) robots and to inquire into their styles of explanation of robots' behaviour.

1. Why robots?

As mentioned before, the GoS is similar to the more-traditional observation and explanation activities carried out in primary schools. No construction/programming activity is involved: children are asked to describe and explain the behaviour of an already programmed robot, more or less in the same way as when children are asked to describe and explain the behaviour of an ant or the growth process of a plant. Why, then, should one select a robot as the "thing" to be analysed in an activity of this sort? Robots are often perceived by children and adults as greatly engaging; this cannot be the only reason, however. This perception of (educational) robots is in many cases quite short-lived and may gradually vanish during interaction. Robots are increasingly used in schools nowadays, and one may reasonably expect that their aura of "novelty" will end up being less and less intense in the future. Are there, then, more stable and strong reasons to choose pre-programmed robots as objects of inquiry in activities such as the GoS? Consistently with Datteri and Zecca (2016), it is proposed here that these reasons are concerned with safety, manipulability, speed, repeatability, and predictability.

- The first two reasons are strongly connected with one another. Observing and explaining the behaviour of chemical reactions may be unsafe for children, if they are allowed to manipulate the chemicals involved. On the contrary, manipulating insects may be unsafe for the insects themselves. Therefore, these materials can be safely used in the framework of observation and explanation activities only under severe restrictions on the possibility of interacting with them. Quite on the contrary, the robots typically used in schools can be manipulated without particularly serious risks for the children and for the robots themselves.
- Checking if a plant dies if not watered may take days. After a chemical reaction has occurred, a significant amount of time (and resources) may be needed to reproduce it again for a second round of observation and explanation. On the contrary, it takes relatively little time to put a pre-programmed robot in the starting position again, to check if it is true that it seeks light. This increases the number of experiments that can be conducted in the same amount of time.

- Predictability and repeatability are connected to each other. Most robots used in education are composed by an algorithmic control system (such as the main brick of the LEGO Mindstorms kit, or the Raspberry Pi used in the robot described later) and a number of physical components such as sensors, motors, and the chassis. On the one hand, the presence of an algorithmic part enables the teacher – who knows how the robot is programmed – to formulate reasonably reliable predictions of how the robot will behave in particular conditions. This may be of great help for the teacher. Suppose, for instance, that three different experiments are proposed by children to test a specific working hypothesis – stating, say, that the robot seeks light – differing from one another in the relative position of a lamp. Based on his/her knowledge of the program implemented in the machine, the teacher can make reasonably accurate predictions of what the robot would do in each of the three different experimental settings: e.g. based on his/her knowledge of the fact that the robot does *not* react to light, (s)he can mentally predict that the robot's behaviour will be nearly the same in all the three conditions. It will be therefore immediately clear to him/her that the three proposed experiments are irrelevant with respect to the testing of the hypothesis at hand: this may help him/her, e.g. to decide how to guide the discussion that will flow from the experiments or how to elicit a reflection on the appropriateness of the proposed experiments before the execution. In addition, the relative predictability of the robot's behaviour, secured by its algorithmic part, guarantees a certain amount of repeatability (if the same experiment is repeated many times, provided that the external conditions are the same, one can be relatively sure that the robot will generate nearly the same behaviour). On the other hand, the presence of the physical – non-algorithmic – components limits the teacher's ability to predict what the robot will do. Yet, in case of mismatch between the teacher's predictions and the observed behaviour, knowledge of the algorithm implemented in the machine enables the teacher to understand whether the prediction failure is due to some environmental or internal physical perturbation or not, thereby helping to guide the ensuing discussion in an appropriate way.

To sum up, there are some interesting reasons – in addition to their ephemeral attractiveness – to choose pre-programmed robots as objects of inquiry in a GoS-like activity.

2. A robo-ethology activity with first year primary school children

2.1. *The robot and the activity*

The GoS experiment described here was proposed to a first year primary school class (23 students, 16 of whom were males and six were females) in Milan (Italy). It was articulated in three weekly sessions lasting 90 minutes each and was guided by the class teacher, one university researcher, and two undergraduate students (they will be referred to as "conductors" from now on).

Two Coderbots (Datteri & Zecca, 2016) were used in the activity. The Coderbot (www.coderbot.org) is a small vehicle based on the Raspberry Pi computer, moving on wheels, embedded with three sonars (mounted on the front, left, and right sides) and a front camera. When turned on, it generates a WiFi network and activates a web server. By connecting to the network and pointing a browser to the appropriate hypertext markup language (HTML) page on the server, one can access the control panel of the robot, which includes buttons for direct tele-operation and a Blockly-based environment for building sensory–motor programs. In this particular edition of the GoS, the Coderbot was directly programmed in Python by the conductors to implement subsumption architecture (Arkin, 1998). Basically, the robot issued avoidance and approaching motor reactions to close and far obstacles, respectively, perceived with the lateral sonars. Detection of close obstacles with the frontal sonar triggered a sequence of apparently bizarre movements. After detecting a certain number of obstacles, the speed of the robot was temporarily halved, simulating a sort of "tiredness". In the absence of obstacles, the robot moved forward.

Each Coderbot was left free to wander in one of two separate experimental arenas, situated outside the classroom. Four groups of pupils were formed by the teacher. In each of the "observation" phases, each one lasting approximately 10 minutes, the groups could observe the robots in the experimental arenas. Occasionally, some of them were allowed to enter the arena and interact with the robot. After each observation phase, the groups went back into the classroom and, under the guidance of the conductors, they were asked to reflect on what they had observed. In particular, they were invited to answer the two following questions: 1) what are the robots doing? 2) Why are they doing that? The conductors invited them to provide answers to these questions, as well as to clarify and justify them. Occasionally, the pupils proposed experiments to test their tentative explanations; these experiments were performed in the subsequent observation phase.

2.2. *Children's perspective on robots: some distinctions*

As mentioned in the Introduction, the GoS activity described here was part of a research study aimed at analysing the styles of description and explanation adopted by primary school children in the analysis of robot behaviours. To this end, the entire activity was audio-recorded. The transcripts were then subjected to a coding process, performed by one of the two undergraduate students participating in the research, based on a taxonomy of categories formulated by the authors before the experiments.

Other authors have investigated the way in which children perceive, describe, and interpret the behaviour of (educational) robots. Based on the study by Ackermann (1991), Levy and Mioduser (2008) distinguish between a *psychological* and a *technological* perspective in the analysis of robot behaviours: "While the first view attributes a robot's behaviours to higher purposes, framed as animate intentions and emotions, personality and volition, the second assigns causality to the inanimate material and informational building blocks which build up the mechanism of the system (i.e., physical parts such as motors and sensors, and the control program, governing the system's interactions). The two approaches are at times distinct, and in other cases entwined and related through growing experience with such artifacts" (p. 338).

In another part of the same article, the authors present the distinction between the psychological and technological perspective as the distinction between two ways of expressing the condition–action rules governing the behaviour of a robot. According to the authors, a technological perspective is adopted whenever a robot's behaviour is analysed in terms of sets of condition–action rules connecting the *inputs* (conditions) and *outputs* (actions) of the device. The psychological perspective is instead characterised by the formulation of condition–action rules linking *contextual features* and *behaviour*. This formulation of the distinction by the authors does not fully clarify it, and the examples they provide do not help one to understand. For example, a kindergarten child's statement that the robot "sensed the light" is classified as technological by the authors, even though no input–output rule is described; and the claim that the robot "doesn't know what the white is" is categorised as psychological even though no connection between context and behaviour is made.

In any case, the authors follow Ackermann (1991) in believing that "integrating the two kinds of explanations—synthesis of the behavioural and the psychological—are the core of a whole explanation. She argues that the ability to animate or give life to objects is a crucial step toward the construction of mature cybernetic theories".

Their work provides interesting insights on children's perception of robots. However, the distinction between the two perspectives identified by the authors deserves further analysis. First, *on what* exactly may one adopt a psychological or a technological perspective? The focus of the authors is on children's analysis and understanding of robot behaviours. Here, it is proposed that vague reference to children's "understanding" of robot behaviours may be usefully replaced with a more specific reference to the two main processes that make such understanding possible, namely *description* and *explanation* of robot behaviour. One way to reformulate the claim by Levy and Mioduser is that children may produce psychological and technological descriptions and explanations of the behaviour of a robot.

Second, in what exactly do the two perspectives differ from one another? As mentioned earlier, according to the first definition proposed by the authors, the psychological perspective is characterised by the attribution of higher purposes, intentions and emotions, personality and volition to the robot, whereas the technological perspective is characterised by the identification of causal processes in the inanimate and informational material. This way of couching the distinction is vulnerable to a number of conceptual objections.

First, it is not clear why attribution of "higher purposes", whatever this means, should be distinctively connected to the adoption of a psychological stance towards robots. Pioneers of the cybernetic movement (Rosenblueth, Wiener, & Bigelow, 1943) have argued that technological artefacts can be sensibly regarded as purposeful if they implement a feedback-based mechanism, independently of whether that mechanism is described in psychological terms or not. Accordingly, one may attribute purposes to a physical system within the framework of a physical, non-psychological account of its internal mechanism. Second, reference to the *causal processes* in the *informational* processing mechanism realised by the system distinctively features, according to the authors, in technological (non-psychological) analyses of robot behaviours. However, it is by no means obvious that information-processing mechanisms can be regarded as causal in the same sense in which physical processes are ordinarily thought of. Moreover, it is not clear why the description of a system as an information processing mechanism should be distinctively connected to the adoption of a non-psychological stance, as many psychological systems are said to process information (Piccinini & Scarantino, 2010).

Considering these obscurities, the distinction between the psychological and the technological perspectives is re-formulated here in terms of two orthogonal distinctions. The first one concerns the adoption of a particular theoretical vocabulary in the description and explanation of robot behaviours. One may use a psychological vocabulary, attributing mental states, internal representations, free will, and emotions to a system ("The robot wants to go out of the arena"); or, one may use a non-psychological vocabulary such as that of physics ("The robot moves using electricity") or of biology ("The robot has a brain"). The second distinction specifically concerns *explanation*: one may formulate a *teleological*, finalistic explanation of the behaviour of a robot ("The robot has turned right *to reach the light*") or a *mechanistic* one ("The robot has turned right *because it has sensed an obstacle on the left*"). In the first case, one explains a particular behaviour by identifying its purpose, whereas in the second case, the behaviour is explained by reference to its causes. The distinction between teleological and mechanistic explanations has been extensively discussed in the philosophy of science literature (Psillos, 2003).

The distinction between the psychological and technological perspectives made in the study by Levy and Mioduser (2008) conflates these two orthogonal distinctions. Indeed, in the analysis made here, one may well formulate a psychological, mechanistic explanation that does not attribute "higher purposes" to the system ("The robot has turned right *because it thought that there was an obstacle on the left*"). On the other hand, one may formulate a finalistic, non-psychological explanation. The notional explanation mentioned earlier – "The robot has turned right *to reach the light*" – is clearly finalistic, but neither mental states nor properties are attributed to the system (unless one believes that having a purpose implies having a mind, in contrast with cyberneticians' analysis). Table 1 provides a summary of the proposed distinctions, which, taken together, provide a tentative and preliminary taxonomy for investigating the styles of explanation of robot behaviours in GoSs and in other kinds of scientific reasoning activities.

Table 1. Summary of the main types of explanation discussed in the paper, with examples

Explanation type	Psychological vocabulary	Non-psychological vocabulary
Mechanistic	The robot turned left because it thought there was something on the right.	The robot turned left because its sensors have revealed an obstacle on the right.
Teleological	The robot turned left to reach the light it wanted to reach.	The robot turned left to reach the light.

3. Varieties of explanation

During the GoS described here, many explanations have been formulated, which exemplify the categories and the distinctions discussed earlier. In particular, the following excerpts are examples of *teleological*, *non-psychological* explanation:

Teacher: Ok, do you all agree? Now, in your opinion, why is the robot digging?

Child: It's digging to find the things it has found.

Teacher: So, children, have you understood anything on what is the robot doing, or why is the robot doing that?

Child: I have understood that it … is making a hole to pass through, it goes inside the other arena.

These explanations are clearly teleological. However, no attribution of a mind is made there: under the assumption that having a purpose does not imply having a mind, these explanations must be sensibly classified as non-psychological. An example of *teleological*, *psychological* explanation is the following.

Teacher: What is the robot doing now?

Child: It's moving!

Teacher: And why is it moving this way?

Child: Because it wants to visit its friend!

Child: It want to go where it was born.

Note that the teleological character of this explanation is controversial. On the one hand, the child identifies a purpose in the robot (i.e. visiting a friend). However, the fact that really explains its behaviour – in the child's opinion – is that the robot *wants* to visit a friend, and the "act" of wanting something may well be regarded as one of the causal motives of the movement, thus making this explanation falling in the mechanistic category. This argument would imply that nearly every psychological explanation in which the robot is said to do something because it wants, believes, desires, or intends to do something should be classified as psychological and mechanistic, in clear contrast with Ackermann's (1991) claims stating that psychological explanations are teleological. In this paper, tentatively – and waiting for a more specific analysis of this conceptual problem, this explanation is classified as teleological as the child is, indeed, pointing to a purpose the robot wants to achieve.

Some *mechanistic* explanations formulated by the children were couched in *non-psychological* terms, as in the following examples.

Child 1: … because the other two wheels are on the other side!

Child 2: Because it has two USB keys inside…

Child 1: But I have understood that when the light on the key is blue… that when it finishes, it stops, then starts moving again.

These are clearly sketches of mechanistic explanations: children do not produce full-fledged descriptions of the mechanism producing the phenomenon of interest. However, what matters to the present discussion is that these tentative explanations point to the causes of that phenomenon and not to the purpose of the robot. Notably, children have formulated *psychological mechanistic* explanations as well:

Child 1: It stopped because it was thinking on where to go.

Child 2: Yes, it realises that we are moving, then it decides to wait.

Child 3: … probably because it's happy.

Child 4: Because it's happy or perhaps because it's angry, maybe it's sad or it's hungry.

These explanations are couched in psychological terms: they attribute mental states and emotions to the robot. However, they are regarded as causal factors producing the phenomenon of interest. Psychological explanations need not be finalistic.

A close analysis of the transcriptions has offered insights to refine the taxonomy of explanation types illustrated so far. For example, it has been found that mechanistic psychological explanations may differ from one another in whether a mind is attributed to the whole robot ("… because the robot wants to go out of the arena") or to a part of it ("The robot goes there because the wheels want it to go there"). They may also differ from one another in whether *cognitive states or processes* ("… because the robot is thinking on where to go") or *emotional states or processes* ("… because it is sad") are attributed to the robot. The emotional states can be permanent ("… because it is shy") or transitory ("… because it is angry with its friend"). The cognitive states or processes may be of various sorts as well. In some cases, mere intentionality and free will are attributed to the robot. Consider the following example again:

Teacher: And why is it moving this way?

Child: Because it wants to visit its friend!

In other cases, the robot is considered as being able not only to want things but also to *decide* how to do what it wants to do, as in the following example.

Teacher: She says that the robot is looking for an exit.

Child: Yes, because it understands that we were moving, then it decides to wait.

Conclusions

Doing science involves the ability to describe and explain phenomena. Pre-programmed, animal-like robots can be used as didactic tools to advance these and other key scientific thinking abilities in children. In this paper, a "robo-ethological" activity called the "Game of Science" has been described and exemplified. A taxonomy of categories for analysing the types of explanations produced by the children, which draws insights from the philosophy of science literature and refines those proposed in previous studies, has been provided and discussed. The qualitative analysis presented here offered the first insights on the usefulness of the taxonomy and leads to the identification of sub-categories. In future studies, this analysis will be subjected to further refinements and articulation, in addition to being applied to other GoS activities to obtain quantitative data. The underlying empirical hypothesis is that this taxonomy can be useful for teachers to analyse children's way of understanding – in the proposed account, of describing and explaining – robot behaviours, in GoSs as well as in other kinds of activities. Knowing what styles of description and explanation of robot behaviours children typically adopt may facilitate didactic planning of GoS-like activities. More generally, the distinctions proposed herein between the use of different theoretical vocabularies and different types of scientific explanation can be useful to inquire into children's perceptions of robotic systems.

References

Ackermann, E., (1991). The agency model of transactions: Towards an understanding of children's theory of control. In J. Montangero, & A. Tryphon (Eds.), *Psychologie genetique et sciences cognitives*. Geneve: Fondation Archives Jean Piaget.

Arkin, R. C. (1998). *Behavior-based robotics*. The MIT Press.

Atmatzidou, S., & Demetriadis, S. (2016). Advancing students' computational thinking skills through educational robotics: A study on age and gender relevant differences. *Robotics and Autonomous Systems*, 75, 661–670. http://doi.org/10.1016/j.robot.2015.10.008

Benitti, F. B. V. (2012). Exploring the educational potential of robotics in schools: A systematic review. *Computers and Education*, 58(3), 978–988. http://doi.org/10.1016/j.compedu.2011.10.006

Braitenberg, V. (1986). Vehicles. Experiments in Synthetic Psychology. Cambridge, MA: The MIT Press.

Church, W., Ford, T., Perova, N., & Rogers, C. (2010). Physics With Robotics-Using LEGO MINDSTORMS In High School Education. *Educational Robotics and Beyond*, 47–49. Retrieved from http://www.aaai.org/ocs/index.php/SSS/SSS10/paper/viewPDFInterstitial/1062/1398

Datteri, E., & Zecca, L. (2016). The Game of Science: An Experiment in Synthetic Roboethology with Primary School Children. *IEEE Robotics & Automation Magazine*, 23(2), 24–29. http://doi.org/10.1109/MRA.2016.2533038

Levy, S. T., & Mioduser, D. (2008). Does it "want" or "was it programmed to..."? Kindergarten children's explanations of an autonomous robot's adaptive functioning. *International Journal of Technology and Design Education*, 18(4), 337–359. http://doi.org/10.1007/s10798-007-9032-6

Levy, S. T., & Mioduser, D. (2010). Approaching complexity through planful play: Kindergarten children's strategies in constructing an autonomous robot's behavior. I*nternational Journal of Computers for Mathematical Learning*, 15(1), 21–43. http://doi.org/10.1007/s10758-010-9159-5

Mioduser, D., & Levy, S. T. (2010). Making sense by building sense: Kindergarten children's construction and understanding of adaptive robot behaviors. *International Journal of Computers for Mathematical Learning*, 15(2), 99–127. http://doi.org/10.1007/s10758-010-9163-9

Mubin, O., Stevens, C. J., Shahid, S., Mahmud, A. Al, & Dong, J.-J. (2013). A Review of the Applicability of Robots in Education. *Technology for Education and Learning*, 1–7. http://doi.org/10.2316/Journal.209.2013.1.209-0015

Papert, S. (1980). Mindstorms: Children, Computers, and Powerful Ideas. New York: Basic Books.

Piccinini, G., & Scarantino, A. (2010). Computation vs. information processing: Why their difference matters to cognitive science. *Studies in History and Philosophy of Science Part A*, 41(3), 237–246. http://doi.org/10.1016/j.shpsa.2010.07.012

Psillos, S. (2003). *Causation and Explanation*. Montreal, Canada: McGill-Queen's University Press.

Rosenblueth, A., Wiener, N., & Bigelow, J. (1943). Behavior, Purpose and Teleology. *Philosophy of Science*. http://doi.org/10.1086/286788

Sullivan, F. R. (2008). Robotics and Science Literacy: Thinking Skills, Science Process Skills and Systems Understanding. *Journal of Research in Science Teaching*, 45(3), 373–394. http://doi.org/10.1002/tea

Whittier, L. E., & Robinson, M. (2016). Teaching Evolution to Non-English Proficient Students by Using Lego Robotics. *American Secondary Education*, 35(3), 19–28.