Detect-Interpret-Respond: A Framework to Ground the Design of Student Inquiry into AI Systems

Eric Greenwald[1], Ari Krakowski[1], Timothy Hurt[1], and Ning Wang[2]

[1] Lawrence Hall of Science, University of California, Berkeley, CA, USA
[2] University of Southern California, Los Angeles, CA, USA

Abstract. In this paper, we describe our experience developing a framework for understanding AI systems that we use to drive the design of AI learning experiences for elementary-aged youth in an informal, free-choice environment. This framework—detect/interpret/respond (DIR)—shows promise as a flexible and age-adaptable model for youth to connect across learning experiences and work toward a coherent understanding of AI. As an application of DIR, we describe our experience designing and testing learning experiences related to a cutting-edge AI system (a Virtual Human). Insights from our initial studies suggest that this framework can help unify youth encounters with various AI systems and provide a productive schema that accommodates increasing complexity as youth advance in their understanding. DIR, therefore, offers a heuristic for sense-making and inspection of AI systems that is both accessible to very young children and robust enough to be useful as they mature.

Keywords: Artificial Intelligence, Learning Design

1 Introduction: The Need for Pre-College AI Learning Experiences

AI is a foundational technology that is profoundly reshaping our society, particularly through its power to transform the workforce [1]. AI has emerged as a tool with tremendous potential in STEM as well, as scientists begin to utilize it to make transformative improvements and discoveries in science, engineering, mathematics, and technology itself. It is imperative to develop a thoughtful response to this phenomenally transformative technology and address the need to make substantial investments in developing the AI workforce, not just through higher education, but beginning with childhood learning [2]. While a small percentage of youth will become the future AI developer workforce, a majority of them will utilize AI in their work, and all will become consumers of AI [3]. It is critical therefore to prepare future generations with basic knowledge of what AI is, its capability, and what impact it will have on their lives and career. Yet, in today's society, there are widespread misconceptions about AI. In Elsevier's interview with leading figures in AI, including Stuart Russell who co-authored the most widely-used AI textbook for higher education, experts pointed out that the common misconceptions about AI roughly fall into two categories:

overestimating the capabilities of AI, especially artificial general intelligence [4], and underestimating the prevalence of AI in our daily lives [5]. A perception that AI is a powerful yet complex tool that is beyond the general public's understanding could negatively impact the current generation of students. They may be discouraged from choosing career paths that require basic AI knowledge because of their lack of AI knowledge and low self-efficacy [6]. The lack of awareness of the prevalence of AI in our daily lives can also impact parents' ability to make informed decisions about the ethics issues related to AI, such as data privacy, both for themselves and for their children. More recent studies of how youth build understanding of AI concepts [see, e.g., 7] suggest additional challenges such as youth encountering difficulty parsing a problem space in ways that enable an AI system to operate on it.

The call to help youth learn AI is coming from practitioners as well. The AI4K12 working group, organized by the Association for the Advancement of AI (AAAI) and the Computer Science Teachers Association (CSTA), has documented the demands from practitioners around the world for curricula and guidelines to help their youth learn AI [3]. This need is recognized by the National Science Foundation (NSF), from which there has been a surge of effort to develop research programs to make AI education accessible to the K-12 population. Efforts are underway to develop a national strategy for research and development in AI (NSTC, 2016), as well as to establish guidelines for K-12 AI education [8] from other organizations. This much-needed effort primarily focuses on community building, curriculum development, or teaching specific AI skills. There is yet a lack of research into education for youth on basic knowledge of AI, such as fundamental mechanisms, capabilities, ethics and implications of AI. There is particularly a gap in our understanding of how to help younger children learn about AI, as most of the current K-12 effort focuses on high-school students, where AI education programs have been more well-established. Because students in their early years are at a critical time for developing their perceptions and dispositions toward STEM [see, e.g., 9], creating engaging AI learning experiences for youth is of paramount importance.

2 Detect-Interpret-Respond: A Framework to Guide Instructional Design for AI Learning

Context: Developing a Virtual Human. As AI is arguably the science of building intelligence that thinks and acts like humans [10], a virtual human provides an ideal vehicle to illustrate many fields of AI, including computer vision, natural language processing, automated reasoning, character animation, and machine learning. A virtual human is an embodied character that can not only see [11], hear [12; 13], and speak [14; 15], but also think [16; 17], feel [18; 19], and move [20; 21; 22] like real humans. The Integrated Virtual Human, developed through decades of collaborative effort from AI researchers at the Institute for Creative Technologies (ICT) at the University of Southern California (USC), is an exemplar of such AI technology [23]. The embodiment of AI through an interactive Virtual Human can make the abstract and mysterious "black-box" nature of AI concrete, relatable, and accessible to the general public – as the public can relate to how AI is used in a virtual and artificial character to fulfill the same function of a real human. This is particularly important for making the concepts of AI accessible to younger populations. Additionally, a Virtual Human can also demonstrate the limitations of AI: while state-of-the-art AI excels in specific tasks, such as filtering spam emails, it pales in comparison with the human-level general intelligence [24]. In all, the Virtual Human technology introduces significant new possibilities for research into engaging the public in AI education in informal settings such as museums, engaging visitor groups across a wide range of backgrounds.

Driven by the demand for AI education for pre-college aged youth, there has been a recent surge of after-school programs to engage K-12 students in hands-on experience to learn AI. For example, the EmpowStudio, in collaboration with the MIT Media Lab, has developed a week-long summer camp to help youth learn AI using personal robots developed by MIT [8]. The AI4All also organizes summer AI learning camps in collaboration with universities across the US to provide underrepresented high school students with hands-on learning experiences, mentored by AI practitioners at the participating universities [25]. In museum settings, AI is commonly used as an advanced technology to enhance visitor experiences, including pioneering work with Virtual Humans as museum guides at the Boston Museum of Science [26]. Virtual humans have been used in museum settings to discuss computer science and engineering topics with middle school students as well [27]. During the exhibit, installed at the Orlando Science Center, visitors were able to customize the appearance and text-to-speech voice of a computer avatar of Alan Turing, a pioneer in AI, to discuss topics such as planets, dinosaurs, natural disasters, and the science center information.

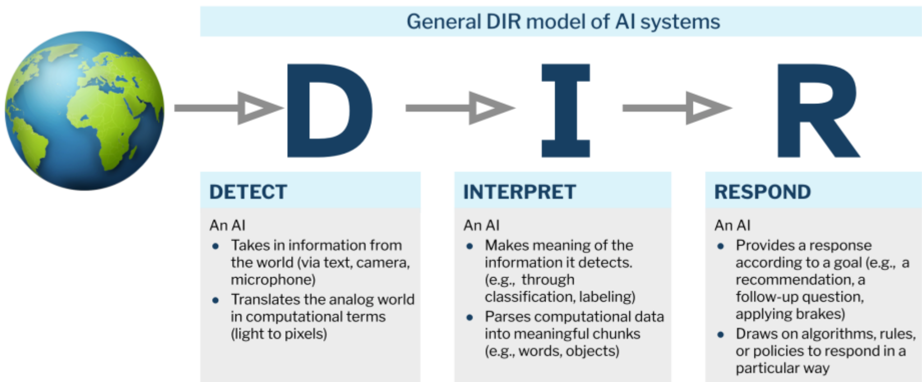

Fig. 1. The Detect-Interpret-Respond Framework for Understanding AI Systems

A Flexible Framework for a Variety of AI Systems. While a Virtual Human offers a rich context for demonstrating cutting-edge AI capabilities, it has yet to be implemented as an object of inquiry to support AI learning for young children. Our work, supported by a grant from the National Science Foundation, has been to do just that: to develop and study learning experiences that can foster youth inquiry into the Virtual Human as an AI system. To ground our work, our team of learning designers, AI developers, and learning science researchers developed and have begun testing an overarching "DIR" framework (detect-interpret-respond) to structure experiences across the portfolio of activities related to the virtual human exhibit (VHX). As illustrated in Figure 1, the intent is for this framework to serve as a working schema for participants to attach their building understanding of AI systems, as embodied in the virtual human. That is to say, participants will come to understand the Virtual Human as a series of interconnected AI subsystems that together enable the VH to **detect** salient objects in its environment (e.g., faces/mouths/eyes), **interpret** those features (e.g., as expressions), and **respond** to that interpretation in human-like ways (e.g., to smile back at a visitor whose detected expression is interpreted to be smiling). While data analysis is ongoing, early review of data from pilot studies have led to initial insights about the DIR framework and learners' interactions along the detect-interpret-respond continuum.

While this paper focuses on the application of the DIR framework to the design of learning experiences related to the Virtual Human, an important part of what makes the framework useful is that it is applicable to a wide range of AI systems that youth might encounter in their daily lives. For example, a self-driving car needs to detect road and traffic conditions, interpret that information in terms of the goal of getting to a particular place safely, and then respond to that interpretation with appropriate actions (e.g., applying brakes). Or, with an AI personal assistant, it must detect the user's input (voice or text), interpret that as a request for a particular type of information, then respond with an appropriate answer to that inferred request. While we do not expect the DIR framework to perfectly map onto all AI systems, our hope is that it can serve as a flexible heuristic for initial encounters with many forms of AI so that their basic functionality can be explored and brought into coherence as a case of AI, thereby supporting learners to develop a generalizable explanatory account of how such systems work.

An Extensible Model for Increasingly Sophisticated Understanding of AI Systems. In addition to allowing for flexibility in its application across a variety of AI systems, DIR offers promise as an extensible mental model to ground a learner's increasingly sophisticated understanding of AI systems over time. For example, while detection can be reduced to the idea of a camera or microphone for very young children, it can also be expanded to include the computational capacity involved in converting light into pixels, or sound into a waveform. From this perspective, DETECT locates the notion of perception represented in "Big Idea #1" [28] within the context of a functioning AI system, providing a model for AI perception that can be built upon as a learner advances. It also offers a way to integrate broader notions of computational thinking within a learner's emerging understanding of AI: to detect, an AI system needs to render the analog world in computational form. By the same token, INTERPRET can be expanded to include increasingly technical explication of the algorithms that enable an AI system to, for example, parse a waveform into language, or a visual image into recognizable objects. While the technical features of different machine learning algorithms for speech recognition, for example, are beyond the learning goals for early-elementary children, the more basic notion that an AI system needs to solve the problem of recognizing speech—parsing sounds into words—offers a foundation upon which the more technical knowledge can eventually be built. Similarly, RESPONSE can be expanded to include the algorithms or policies that determine whether a particular action should be taken in a given circumstance and/or the mechanics of how that response takes place (e.g., syncing voiced language to animated facial movements).

3 Experience Report: Applying DIR to the Design of a Virtual Human Learning Experience

Methods. This report is based on a series of small-scale pilots of activities related to (and including) the Virtual Human. The pilots support the ongoing iterative design of a museum-based learning experience centered around an interactive Virtual Human as an exemplar of AI technologies. Two critical considerations for the eventual experience are: 1) that it can succeed in a free-choice environment, like a public museum, where learners can choose, or not, to approach the exhibit, and decide how long to dwell within it; and 2) that it can succeed with the typical visitors to the museum: early elementary youth and their adult caregivers.

Insights about the promise of the DIR model are drawn from a series of cognitive interviews conducted with 26 children from ages 6-13. Interviews participants were asked to think aloud as they interact with the activities in order to understand how youth of this age make sense of AI and how these activities can be made engaging and interesting to youth in a free-choice environment. The cognitive interviews employed a semi-structured protocol, with a series of prompts to elicit student thinking as they: a) initially encountered each activity; b) attempted to engage with the activity (e.g., "I see that you [did X], tell me about your thinking; and, how did you decide to do X?"); and c) after they had completed the activity (e.g., "how do you think this AI worked?"). In-the-moment scaffolding was provided during interviews to: a) enable students to reveal thinking across each phase of the activity; b) surface and test emerging ideas about why a student might be stuck or lose interest; and c) disambiguate between superficial challenges, such as unfamiliar vocabulary or confusing UI, and more conceptual difficulties. Some activities were piloted with actual AI systems in their current state and others were conducted as Wizard-of-Oz (WOZ) style studies in which a human worked behind the scenes to approximate desired AI functionality. The WOZ studies were particularly useful to test youth sensemaking around the "Response" side of the DIR framework, as it enabled us to probe for the range of responses that would be likely to arouse youth curiosity prior to the technical development of those features.

Helping Youth Get Purchase on Understanding AI as a System. The learning design task entailed taking an existing AI technology (the Virtual Human) and positioning it as an artifact that learners, ages 5-13, could engage with in an informal science museum setting in order to build their understanding of AI. Before putting this technology in front of visitors, we did an audit to better understand what AI technologies were at work within the system so that we could create surrounds that may better position learners to build their understanding of those specific technologies.

Through identifying how the subsystems worked in a coordinated manner to enable the full AI system, we recognized that some of the subcomponents hung together in terms of their function relative to the larger AI system. Thus, we derived three categories: *Detection*, *Interpretation*, and *Response* (See Figure 2). *Detection* components function to detect the human who was interacting with the AI system. Specifically, the VH uses cameras and microphones to take in visual and audio data. From this data, the AI system can do a number of detections: your face, your body, whether you are speaking, and more. Each of these Detection components feeds into the Interpretation components. The *Interpretation* components, in turn, function to "make sense" of what has been detected. While an AI does not make sense of human actions in the same ways that humans do, they can be designed to functionally appear to do so. Notably, the AI system we were using has the ability to categorize a facial expression as: Angry, Sad, Happy, Disgusted, Surprised, or Neutral. Similarly, body posture in conjunction with head position can be used to infer whether the learner is paying attention or not, and so on. The interpretation components then feed into the Response components of the AI system. The *Response* components function to provide feedback to the learner based on the outcome of the Interpretation components. Our AI system was designed as a Virtual Human, that is, an AI controlled avatar that is supposed to be human-like in its responses to humans: this includes text-to-speech behaviors as well as human body animations and gestures. The Virtual Human we were working with had a number of responses it could make based on the outcome of the Interpretation components; as one example, if the Virtual Human interpreted your facial expression as a smile, and interpreted that you were paying attention with your body position, it would smile back at you.

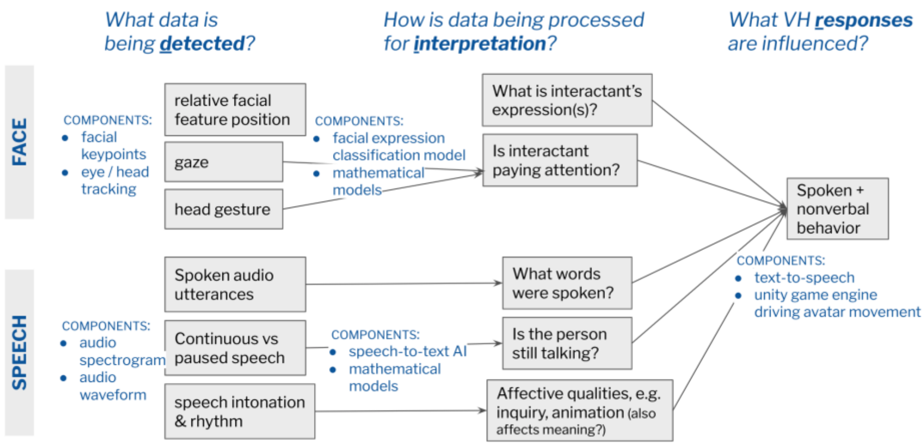

Fig. 2. VH represents an AI system built from integrated subsystems.

Reconciling a System of Systems for Learners. The DIR framework positions the different AI components in relation to each other and to the overarching VH capabilities of taking in data from the environment and processing that data to generate appropriate responses. As illustrated in Figure 2, the VH can be conceived of as an interconnected set of AI components that work in concert to activate different functionalities. Some of these components serve to detect specific types of data from the surrounding environment; some components are focused on processing the data for interpretation; and some components on determining the appropriate VH response based on the interpreted data. This framework can help demystify AI systems for learners by decomposing the complex set of functionalities that comprise the VH into component parts focused on detection, interpretation, and response determination & execution. These capabilities share some commonalities with the ways in which humans (often unconsciously) make sense of their environment in their interactions with other humans. In this way, the DIR model offers an avenue for learners to relate to VH as just that - a virtual human programmed to interact with other humans in naturalistic ways.

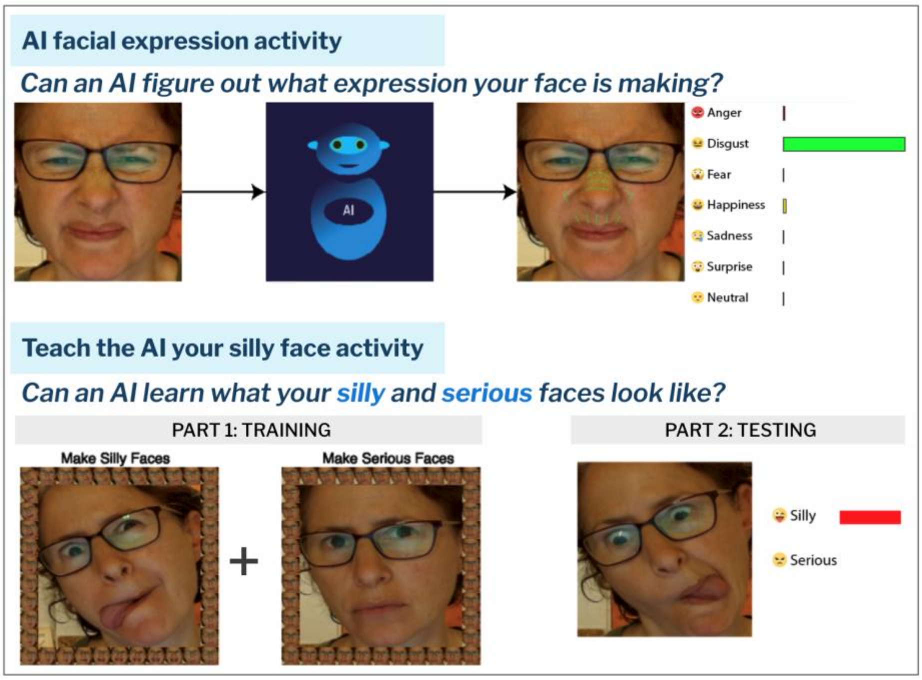

Fig. 3. Examples of activities designed to break down the complexity of the VH.

Building Coherence Across Activities. The development team first implemented a Behind-the-Scenes UI design, which includes a comprehensive display of head-tracking, facial action units tracking, facial expression recognition, speech recognition, etc. The learning design team reviewed the UI and pointed out that such design may be too complex for young learners in a museum setting, presenting cognitive load challenges that could hinder engagement in a free-choice museum environment and thus limit learning about the different AI components that comprise the VH. Instead, the team has begun creating a series of breakout activities, each focusing on a component of the AI technologies embodied in the main VH exhibit. The conceptual design of the exhibit is to have a central VH exhibit (imagined as a pair of visitors, one interacting with VH and a second viewing the comprehensive Behind-the-Scenes UI), and a number of "stations" surrounding the central VH, that visitors can interact with. Figure 3 provides an illustra- tion of two of these activities: the first activity, *Expression Detection,* activity supports the connection between DETECT and INTERPRET and enables the interactant to play with how an AI system can recognize salient facial features (technically, action units) to classify a particular combination of features as happy, sad, surprised, etc.; the second activity, *Teach Me Silly,* is designed to suport deeper engage- ment with the INTERPRET functionality of AI and enables the interactant to use an AI classifier to teach the AI to recognize the interactant's expression of a "silly face." This revised exhibit design can better engage visitors in active learning (hands-on with breakout learning activities at the surrounding stations instead of passively viewing the Behind-the-Scenes UI on a big screen). It also greatly increases the number of visitors the exhibit can engage at the same time, addressing throughput concerns. Important for its implementation, the revised design provides a greater variety of activities (beyond just conversing with the VH) to engage visitors in hands-on and in-depth learning about AI embodied through a VH.

Because each "station," and the VH experience as a whole is grounded in the DIR framework, the suite of activities work to complement each other. The hope is that this will enable the visitor to attach understanding constructed at the individual activity level to their working DIR model of the Virtual Human as an AI system, thus building coherency across activities.

DIR as a Heuristic for Inspection of AI Systems. While DIR has provided an important touchstone to support coherent learning design across a series of activities and to help unpack a complex system-of-systems like the Virtual Human, early piloting suggests it may also provide a useful heuristic for learners as they encounter new AI systems. In effect, DIR represents a working, and workable, model of (artificial) cognition [29] *for* AI systems. For example, in a round of piloting in which youth moved between 2 different AI systems (each a subsystem of the VH), we saw some evidence that youth were able to leverage encounters with one AI system to inform their experience of the other by focusing on what each system was noticing about them (DETECT) and how each system was using what it noticed (INTERPRET) to categorize their behaviors (facial expressions, in this case) in meaningful ways (as expressions). Research is ongoing to determine how best to support youth to recognize and apply DIR as a heuristic for these and similar activities. Research will also investigate the extent to which learners are able to apply DIR to encounters with novel AI systems, beyond intentionally designed activities, and explore how applying the heuristic mediates knowledge construction among learners.

4 Discussion

Guiding ongoing design iterations. The interdisciplinary team of AI developers, researchers, and learning designers continues to make use of DIR to guide ongoing design iterations on the VHX and its suite of activities. In particular, we are exploring ways that UI design elements (such as visual coherency in representation of each of D, I, and R components) and activity features (such as prompts to inspect an AI system for DIR) can better arouse curiosity about the VHX as a *case of AI*—as a system that can detect salient features of its environment, interpret those features in meaningful ways, and respond to the user based on those interpretations.

Initial pilot testing and collaborative sense-making of pilot data has led to a number of insights about the application of DIR to support youth understanding of AI, and has surfaced awareness of tensions between the DIR framework and certain kinds of AI systems. First, the design team noted that the that explainability of AI components decreases as one moves from left to right in the DIR framework. *Detection* is wholly explainable for learners. For example, it is relatively straightforward to describe the mechanism underlying the machine learning model that detects facial keypoints from the webcam-generated image: the UI enables youth to see their face mapped with keypoints that roughly correspond to "high interest" areas such as the corners of the mouth, which rise and fall with expressions. The explainability of *Interpretation* varies across different VH model components. For example, we have explored different models for interpreting visitors' facial expressions ranging from rule-based (i.e., based on calculations of facial action units derived from facial keypoints) to black-boxed neural net models. While rule-based systems are easier to explain, they no longer reflect the cutting-edge of AI in many cases. The *response* dimension of the framework is most opaque. The nonverbal VH responses — arguably the ones that make the VH so innovative with respect to naturalistic human interaction — are backchannel behaviors generated by neural networks. The behaviors are apparent from the software configuration file (e.g., head nod, smile), but it is far from clear what determines the backchannel behaviors. These observations bring up interesting connections to explainable AI, and highlight a challenge of "getting under the hood" of advanced AI systems: when an AI model is more explainable, it is also easier for a learner to make sense of the AI capabilities driving the model. At the same time, many state-of-the-art AI models involve opaque AI functionality, and therefore may require additional scaffolding to make them more explainable and understandable to learners. With this challenge in mind, future piloting will explore ways to support visitors in developing a working "theory of cognition" for the VH through the DIR schema that can account for these more opaque systems.

Limitations. This paper reflects emerging insights from initial testing with a series of AI learning activities related to a Virtual Human system. Our analysis of these data has been exploratory, and our findings are tentative. As we continue to develop learning experiences and conduct research to study them, we are attuned to possible limitations of DIR as a framework. One concern we are examining is that while the DIR model can be applied to support a basic understanding of a wide variety of AI systems, it may risk introducing an artificial step for some AI approaches, such as neural nets. Specifically, as alluded to above, these more opaque systems, in essence, collapse the “Interpretation” and “Response” steps and apply algorithms to input data and provide a response directly. In such systems, there is no system-specific need to parse input data into meaningful (to a human) chunks prior to determining a response. In such cases, DIR may risk introducing misconceptions, such as a problematic anthropomorphizing of AI as something that needs to make human-like sense of the world in order to operate. At the same time, this potential challenge intersects in important ways with recent calls to improve the explainability of AI, work that can involve reverse engineering opaque algorithms to determine a set of human-comprehensible rules that tie inputs to outputs. It may be that the explainability of an AI system is a requirement for DIR to be a useful framework to support learning.

Directions for future research. As a project in its early phases, we continue to iterate on designed experiences and conduct research to examine the usefulness of DIR as a framework to ground learning design and support learner consolidation of ideas. Thus far, we have found DIR to be a useful framework to ground our learning design work, but considerable work is need to empirically evaluate its affordances and constraints for fostering understanding of AI systems among young children. At the same time, we seek to investigate the applicability of DIR to AI systems beyond the Virtual Human: is it broad enough to support the design of learning experiences for a range of AI systems and across a range of experience levels, while being specific enough to provide traction for advancing understanding of each one. These dual attributes of DIR—applicability across a range of AI systems and extensibility as expertise develops—also create opportunities to explore its usefulness in the design of interdisciplinary learning experiences. One could imagine, for example, how a team of students in computer science, robotics, engineering, science, and psychology could apply their developing expertise to the design and development of a complex AI system like the VH. Future research, then, could explore how a framework like DIR might serve as a boundary object to facilitate communication and coordinate work across disciplines.

5 Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. 2116109. 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 National Science Foundation.

6 References

1. McKinsey Global Institute (2017). A Future That Works: Automation, Employment, and Productivity.

2. Office of Science and Technology Policy, White House (2018) Summary of the 2018 White House Summit on Artificial Intelligence for American Industry.
3. Gardner-McCune, C., Touretzky, D., Martin, F., & Seehorn, D. (2019). AI for K-12: Making Room for AI in K-12 CS Curricula. In Proceedings of the 50th ACM Technical Symposium on Computer Science Education (pp. 1244-1244). ACM.
4. Goertzel, B. (2007). Artificial general intelligence (Vol. 2). C. Pennachin (Ed.). New York: Springer.
5. Roffel, S., & Evans, I. (2018) The biggest misconceptions about AI: the experts' view. Retrieve from https://www.elsevier.com/connect/the-biggest-misconceptions-about-ai-the-experts-view
6. Ritter, Caroline. "User-based barriers to the adoption of artificial intelligence in healthcare." PhD diss., Capella University, 2019.
7. Greenwald, Eric, Maxyn Leitner, and Ning Wang. "Learning Artificial Intelligence: Insights into How Youth Encounter and Build Understanding of AI Concepts." In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 17, pp. 15526-15533. 2021.
8. EmpowStudios (2019). Learn and Design Artificial Intelligence, with Ethics in Mind: One-week workshop taught in collaboration with MIT Media Lab. Retrieved from: https://empow.me/ai-ethics-summer-program-mit
9. Bybee, Rodger W., and Bruce Fuchs. "Preparing the 21st century workforce: A new reform in science and technology education." (2006).
10. Russell, S. J., & Norvig, P. (2016). Artificial intelligence: a modern approach. Malaysia; Pearson Education Limited,.
11. Stratou, G., Morency, L. P., DeVault, D., Hartholt, A., Fast, E., Lhommet, M., Lucas, G., Morbini, F., Georgilia, K., Scherer, S. Gratch, J., Marsella, S., Traum, D. & Rizzo, A. (2015, September). A demonstration of the perception system in SimSensei, a virtual human application for healthcare interviews. In 2015 international conference on affective computing and intelligent interaction (ACII) (pp. 787-789). IEEE.
12. Cummins, N., Scherer, S., Krajewski, J., Schnieder, S., Epps, J., & Quatieri, T. F. (2015). A review of depression and suicide risk assessment using speech analysis. Speech Communication, 71, 10-49.
13. Degottex, G., Kane, J., Drugman, T., Raitio, T., & Scherer, S. (2014, May). COVAREP—A collaborative voice analysis repository for speech technologies. In 2014 ieee international conference on acoustics, speech and signal processing (icassp) (pp. 960-964). IEEE.
14. Govind, D., & Prasanna, S. M. (2013). Expressive speech synthesis: a review. International Journal of Speech Technology, 16(2), 237-260.
15. Leuski, A., & Traum, D. (2011). NPCEditor: Creating virtual human dialogue using information retrieval techniques. Ai Magazine, 32(2), 42-56.
16. Pynadath, D. V., Si, M., & Marsella, S. C. (2011). Modeling theory of mind and cognitive appraisal with decision-theoretic agents. Social emotions in nature and artifact: emotions in human and human-computer interaction, 70-87.
17. Ustun, V., & Rosenbloom, P. S. (2016). Towards Truly Autonomous Synthetic Characters with the Sigma Cognitive Architecture. In Integrating Cognitive Architectures into Virtual Character Design (pp. 213-237). IGI Global.
18. Gratch, J., & Marsella, S. (2005). Evaluating a computational model of emotion. Autonomous Agents and Multi-Agent Systems, 11(1), 23-43.
19. Morency, L. P. (2014). The role of context in affective behavior understanding. Social Emotions in Nature and Artifact, 2.
20. Office of Science and Technology Policy, White House (2018) Summary of the 2018 White House Summit on Artificial Intelligence for American Industry.

21. Lee, J., & Marsella, S. (2006, August). Nonverbal behavior generator for embodied conversational agents. In International Workshop on Intelligent Virtual Agents (pp. 243-255). Springer, Berlin, Heidelberg.
22. Lhommet, M., Xu, Y., & Marsella, S. (2015, March). Cerebella: automatic generation of nonverbal behavior for virtual humans. In Twenty-Ninth AAAI Conference on Artificial Intelligence.
23. Swartout, W. R., Gratch, J., Hill Jr, R. W., Hovy, E., Marsella, S., Rickel, J., & Traum, D. (2006). Toward virtual humans. AI Magazine, 27(2), 96.
24. Goertzel, B. (2007). Artificial general intelligence (Vol. 2). C. Pennachin (Ed.). New York: Springer.
25. AI4All (2019). Learn About a Creative New Technology that Can Help the World and Have Your Best Summer Yet. Retrieved from: http://ai-4-all.org/summer-programs/
26. Swartout, W., Traum, D., Artstein, R., Noren, D., Debevec, P., Bronnenkant, K, Williams. J., Leuski, A., Narayanan, S., Piepol, D., Lane, C., Morie, J., Aggarwal, P., Liewer, M, Chiang, J-Y, Gerten. J., Chu. S., & White, K. (2010, September). Ada and Grace: Toward realistic and engaging virtual museum guides. In International Conference on Intelligent Virtual Agents (pp. 286-300). Springer, Berlin, Heidelberg.
27. Hollister, J., Parker, S. T., Gonzalez, A. J., & DeMara, R. (2013, October). An extended Turing test: A context based approach designed to educate youth in computing. In International and Interdisciplinary Conference on Modeling and Using Context (pp. 213-221). Springer, Berlin, Heidelberg.
28. Touretzky, David, Christina Gardner-McCune, Fred Martin, and Deborah Seehorn. "Envisioning AI for K-12: What should every child know about AI?." In *Proceedings of the AAAI conference on artificial intelligence*, vol. 33, no. 01, pp. 9795-9799. 2019.
29. Brown, Nathaniel JS, and Mark Wilson. "A model of cognition: The missing cornerstone of assessment." *Educational Psychology Review* 23, no. 2 (2011): 221-234.