Live and Learn, Ask and Tell: Agents Over Tasks

Don Perlis, Clifford Bakalian, Justin Brody, Timothy Clausner, Matthew D. Goldberg, Adam Hamlin, Vincent Hsiao, Darsana Josyula, Chris Maxey, Seth Rabin, David Sekora, Jared Shamwell, and Jesse Silverberg

Steven Spielberg's "A.I." tells the story of two artificial agents: David and Teddy. While David resembles a human child, his companion Teddy is much simpler. Its behavior, however, still suggests a crucial mix of capabilities that stretch the state of the art in AI today. We argue that, unlike most contemporary AI, Teddy qualifies as a bona fide *agent*, and that implementing such a system would represent a valuable advance in our understanding of agency. We then describe a project to integrate our existing work to create a simple agent with Teddy-like capabilities.

Teddy learns by dialog as well as observation, and makes helpful comments based on its assessment of other agents. A real-life Teddy would seem to involve a synergistic integration of multiple areas including robotics, vision, learning, NLP, knowledge representation, and commonsense and multi-agent reasoning. Teddy doesn't need to *do* a whole lot in the ordinary sense. It is there learning/knowing more and more all the time, is available for simple interactions (such as asking and telling), and can offer suggestions and remember events as they transpire over time (e.g., who else was present on a particular occasion).

By contrast, a great deal of AI is task-based: write a program that accepts input type I and produces output type O. The term "agent" is sometimes used loosely for almost any AI program. But, as we intend the term here, an agent (e.g., Teddy) is an entity that knows what tasks it is attempting, and when and why, and can reason and make decisions about them while doing them, as well as before and after. What is of interest here is not that Teddy succeeds at a particular task X, but that Teddy knows that it its attempting X, and can reason about it and remember it later.

Here we summarize several relevant advances which we are now engaged in integrating into such a (robotic-dialogic) agent:

All except below
University of Maryland, College Park, MD 20742

Justin Brody
Goucher College, Baltimore, MD 21204

Darsana Josyula
Bowie State University, Bowie, MD 20715

Anomaly-handling: We designed and implemented the Metacognitive Loop (MCL) with the purpose of allowing a system to handle anomalies [2, 4] by i) *noting* when the system's actions do not line up with its expectations, ii) *assessing* any such anomaly and examining the available routes for resolving it, and iii) *guiding* the system toward implementing the chosen strategy. This was built using our Active Logic Machine [3], which is a paraconsistent and temporal reasoning engine. [13] discusses how agent-centric logics, such as active logic, are well suited for resolving contradictions through introspection due to the temporal stratification of the inference process. MCL in a way embodies an implementation of the appearance-reality distinction [9]: things may be different from what the agent has believed them to be. This leads to a need for representation of one's past beliefs. Active logic is used for this purpose, for instance to create expressions for "the thing I formerly took to be X" in the case of misidentification [12]. **Dialog and intention:** [15] describes ways that active logic can be used to prevent implications that could misleadingly impact the intended meaning of input phrases. In [11], the TRAINS system [1] is extended using MCL to determine if it understood a speech request and then takes actions to correct errors. **Dialog and robotic self-knowledge:** We sketched a robotic activity: on request, a robot goes to another room to find a particular book [5]. This may sound simple enough, but it is surprisingly complex when designed in a general way. Two robots (Alice and Julia) explored this and unexpectedly brought us face-to-face with new issues. In particular, when Alice was programmed to point at Julia while saying that it was doing so, it would enter an unintended loop by responding to its own utterance of "Julia". Using the neuroscientific idea of efference copy, we enabled Alice to know when she was engaged in talking and thus not respond [7]; this differs markedly from the approach of [6] who programmed a robot to recognize the sound of its own voice, and then decide on that basis that it must be speaking. **Dialog and self/other knowledge:** We recently examined self/other knowledge complexities arising even in a simple dialogue [10]. If asked, "Is there milk in the fridge?" Teddy should: (i) realize another agent is addressing it, (ii) consider if it has appropriate knowledge of the fridge and milk, or (iii) if it realizes it doesn't know, then possibly (iv) infer that it can find out by looking in the refrigerator, and so on. Finally, after responding, it should understand that the other agent now knows the information in the response. **Navigation:** We have developed a voice-controlled navigation interface for our robot agent. Current work includes: (1) expanding the command vocabulary to include more complex movements and objectives; (2) using computational linguistics to attempt to resolve unknown/misheard commands. **Object recognition:** We are using object-detection software from state of the art deep neural networks to locate and classify objects within an image. Our network can detect bottles, chairs, people, and other common indoor objects. This can be made more complex via different neural network architectures, increased network object classes, and eventually specific agent identification (or misidentification leading to anomaly detection and response). *We expect that a Teddy-bot endowed with a variety of abilities, particularly learning, asking, and telling, operating over its lifetime, will be able to assess and adjust its actions, and, in so doing, will yield new discoveries about how dialog and lifelong learning enable task-general behaviors.*

References

1. Allen, J. F., Miller, B. W., Ringger, E. K., Sikorski, T.: A robust system for natural spoken dialogue. In: Proceedings of the 34th annual meeting on Association for Computational Linguistics, pp. 62-70. Association for Computational Linguistics (1996, June)
2. Anderson, M. L., Perlis, D. R.: Logic, Self-awareness and Self-improvement: the Metacognitive Loop and the Problem of Brittleness. Journal of Logic and Computation **15.1**, 21–40 (2005)
3. Anderson, M. L., Gomaa, W., Grant, J., Perlis, D.: Active logic semantics for a single agent in a static world. Artificial Intelligence **172 (8-9)** 1045–1063 (2008)
4. Anderson, M. L., Oates T., Chong W., Perlis, D. R.: The metacognitive loop I: Enhancing reinforcement learning with metacognitive monitoring and control for improved perturbation tolerance. Journal of Experimental and Theoretical Artificial Intelligence **18.3**, 387–411 (2006)
5. Bhargava, P., Cox, M., Oates, T., Oh, U., Paisner, M., Perlis, D., Shamwell, J.: The robot baby and massive metacognition: Future vision. In: Proceedings of the 2012 IEEE International Conference on Development and Learning and Epigenetic Robotics, pp. 1–2 (2012, November)
6. Bringsjord, S., Licato, J., Govindarajulu, N.S., Ghosh, R., Sen, A.: Real robots that pass human tests of self-consciousness. In: Robot and Human Interactive Communication (RO-MAN), 2015 24th IEEE International Symposium on. pp. 498–504. IEEE (2015)
7. Brody, J., Perlis, D., Shamwell, J. (2015, September). Whos TalkingEfference Copy and a Robots Sense of Agency. In: 2015 AAAI Fall Symposium Series.
8. Chong, W., O'Donovan-Anderson, M., Okamoto, Y., Perlis, D. (2002, January). Seven days in the life of a robotic agent. In Workshop on Radical Agent Concepts (pp. 243-253). Springer, Berlin, Heidelberg.
9. Flavell, J. H., Green, F. L., Flavell, E. R., Watson, M. W., Campione, J. C. Development of Knowledge about the Appearance-Reality Distinction. Monographs of the Society for Research in Child Development **51.1**, i-87, 1986.
10. Goldberg, M. D., Brody, J., Clausner, T. C., Perlis, D.: The overlooked role of self-agency in artificial systems. In: Workshop on Shortcomings in Vision and Language at European Conference on Computer Vision (2018)
11. Josyula, D. P., Fults, S., Anderson, M. L., Wilson, S., Perlis, D.: Application of MCL in a dialog agent. In: Third Language and Technology Conference (2007)
12. Miller, M.: A View of One's Past and Other Aspects of Reasoned Change in Belief. University of Maryland at College Park (1993)
13. Perlis, D., Brody, J., Kraus, S., Miller, M.: The internal reasoning of robots. In: Proceedings of the Thirteenth International Symposium on Commonsense Reasoning (2017)
14. Perlis, D., Purang, K.: Conversational adequacy: mistakes are the essence. (1998)
15. Purang, K., Perlis, D., Gurney, J.: Active logic applied to cancellation of gricean implicature. In: AAAI 96 Spring Symposium on Computational Implicature (1996)
16. Sekora D., Barham, S., Brody, J., Perlis, D.: Anatomy of a Task: Towards a Tentative Taxonomy of the Mind. In: AAAI Fall Symposium on a Standard Model of the Mind (2017)