From Problem Solving to Problem Posing

Yoonsuck Choe, Timothy A. Mann

***Abstract*—Artificial intelligence and machine learning approaches are both very good at problem solving. However, the various methods accumulated in these fields have not been able to give us truly autonomous agents. The main shortcoming is that the problems themselves are formulated by human designers and subsequently fed to the problem solving or learning algorithms. The algorithms do not question the validity of the problems nor do they formulate new problems. This latter task is called "problem posing", and is in fact an active area in education research. In this article, we will discuss the importance and relevance of problem posing to autonomous intelligence and speculate on key ingredients for effective problem posing in an AI and machine learning context.**

***Index Terms*—Problem posing**

I. The problem in problem solving

ARTIFICIAL intelligence (AI) and machine learning have come a long way in automated problem solving and learning from data. However, in most cases the algorithms we use are geared toward solving problems that are precisely represented and/or defined, whether those are logical inference problems or data classification problems. This is good, but not good enough if we consider our ultimate goal in AI–building truly autonomous intelligent agents. Here we argue that the main shortcoming is that of "problem posing": AI and machine learning algorithms cannot identify or come up with novel yet relevant problems to solve within the broader context of their overall goal. We strongly believe this imbalance between problem solving and problem posing needs to be addressed in order to make progress in autonomous agent research.

II. Problem posing and related approaches

It turns out that problem posing has been intensively investigated, not in AI and machine learning, but in the education research community: Problem posing has been used as both a metric of conceptual understanding and also as a pedagogical tool. For example, Brown and Walter applied problem posing to mathematics education [1]. A notable quote in the book is that once people realized "Can we prove ...." was the right question rather than "How can we prove ...", a seemingly intractable question suddenly became tractable. Mestre on the other hand experimentally investigated the relationship between problem posing ability and the degree of conceptual understanding of physics concepts in high-performing college students [2]. Also, it is well accepted that the process of science heavily depends on posing the right questions [3] (as cited in [2]).

On the AI and machine learning side, there are some preliminary approaches related to problem posing. We are not talking about IBM's Watson [4]: Although the output of Watson to Jeorpardy-like queries is in the form of a question, Watson is fundamentally a question answering system, which is a form of problem solving. A recent work that addresses problem posing more directly is that of Cakmak and Thomaz [5], where they ask how can we design robots that ask good questions. However, in their work, the questions are limited to only three categories (label, demonstration, and feature), so it is more about "when" to ask "which" question, rather than asking more general, novel, task-related questions. An earlier related work by Fong et al. [6] investigated how humans react differently to robots that they can directly control as opposed to robots that ask questions answered by humans for indirect control. As in Cakmak and Thomaz [5], the questions asked by the robots were simple and canned, so this was not problem posing in its true sense. There is some relation between these works above and active learning [7], [8], [9] but active learning is more about actively selecting which data point to include in learning so it is not problem posing in the strictest sense. Finally, the problem of problem posing has in some way been foreseen in the autonomous mental development framework where agents are born and subsequently given novel tasks [10]. If the agents can identify novel tasks by themselves (instead of being given by the designer) through the use of a developmental program, then that is equivalent to problem posing.

III. Automating problem posing

How can we automate problem posing? This is a key open question that could lead to major advances in autonomous agent research. Although the answers are unclear, we can consider several classes or problems that could be posed, and factors that could be important.

First, let us look at different types and methods for problem posing:

- Recognizing an event as a problem when it arises in the agent's environment, e.g., a stove on fire: This could be the easiest to achieve and maybe we should start from here.
- Questioning existing problems: One way to pose a new relevant problem is to start from existing problems and check if they are well-posed. If they are ill-posed, there is an opportunity to reformulate it. In many cases, ill-posed problems are due to invalid or unrealistic assumptions, so checking the assumptions could be a good first step.
- Given an overarching goal, posing problems that could gradually lead to the goal: This may sound similar to subplanning but it is more general than that since unlike

The authors are with the Department of Computer Science and Engineering, Texas A&M University, College Station TX, 77843-3112 USA e-mail: {choe,mann23}@tamu.edu.

Manuscript received May 1, 2012; revised June 3, 2012.

subplanning where the decomposed tasks are the same kind as the original, the approach outlined above will pose problems of various different kinds.

Next, there are several important factors that could contribute to problem posing: Grounding, affordances, analogy, what kind of mathematical formalism to use, and how to measure the goodness of the posed problems? Finally, we often find it effective to ask "what is the nature of X?" and "what kind of objective function to use any why?".

IV. Conclusion

In this article, we argued that our lack of attention to problem posing can be a main obstacle in building truly autonomous agents, and provided some initial insights on problem posing strategies. A more rigorous framework will be needed to bring these ideas into the main stream AI and machine learning.

References

[1] S. I. Brown and M. I. Walter, *The Art of Problem Posing*, 3rd ed. Mahwah, NJ: Lawrence Erlbaum Associates, Inc., 2005.

[2] J. P. Mestre, "Probing adults' conceptual understanding and transfer of learning via problem posing," *Journal of Applied Developmental Psychology*, vol. 23, pp. 9–50, 2002.

[3] K. Dunbar, "How scientists think in the real world: Implications for science education," *Journal of Applied Developmental Psychology*, vol. 21, pp. 49–58, 2000.

[4] D. Ferruci, E. Brown, J. Chu-Carroll, J. Fan, D. Gondek, A. A. Kalyanpur, A. Lally, J. W. Murdock, E. Nyberg, J. Prager, N. Schlaefer, and C. Welty, "Building Watson: An overview of the DeepQA project," *AI Magazine*, vol. 31, pp. 59–63, 2010.

[5] M. Cakmak and A. L. Thomaz, "Designing robot learners that ask good questions," in *Proceedings of the 7th ACM/IEEE International Conference on Human-Robot Interaction*, 2012, in press.

[6] T. Fong, C. Thorpe, and C. Baur, "Robot, asker of questions," *Robotics and Autonomous Systems*, vol. 42, pp. 235–243, 2003.

[7] M. Cakmak, C. Chao, and A. L. Thomaz, "Designing interactions for robot active learners," *IEEE Transactions on Autonomous Mental Development*, vol. 2, pp. 108–118, 2010.

[8] D. A. Cohn, Z. Gharamani, and M. I. Jordan, "Active learning with statistical models," *Journal of Artificial Intelligence Research*, vol. 4, pp. 129–145, 1996.

[9] H. S. Seung, M. Opper, and H. Sompolinsky, "Query by committee," in *Proceedings of the fifth annual workshop on Computational learning theory*, ser. COLT '92. New York, NY, USA: ACM, 1992, pp. 287–294. [Online]. Available: http://doi.acm.org/10.1145/130385.130417

[10] J. Weng, "Developmental robotics: Theory and experiments," *International Journal of Humanoid Robotics*, vol. 1, pp. 199–236, 2004.

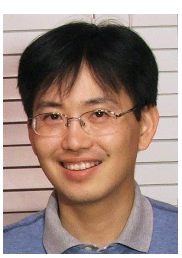

Yoonsuck Choe is an associate professor and director of the Brain Networks Laboratory in the Department of Computer Science and Engineering at Texas A&M University. He received his B.S. degree in Computer Science from Yonsei University (Korea) in 1993, and his M.S. and Ph.D. degrees in Computer Sciences from the University of Texas at Austin in 1995 and 2001. His research interest is broadly in imaging, modeling, and understanding brain function, from the local circuit level up to the whole brain scale, with a focus on the temporal and sensorimotor aspects of brain operation.

Timothy A. Mann received is B.A. in Computer Science and B.A. in Mathematics from the State University of New York at Potsdam College. He is currently a PhD candidate in the Department of Computer Science and Engineering at Texas A&M University. He is a fellow of the NSF East Asia and Pacific Summer Institutes 2011 program and the Texas A&M University Dissertation Fellowship program. His main research interests include the application of transferred knowledge to scale up reinforcement learning algorithms to complex real-world tasks.