Studying human problem solving using citizen science games.

Jacob F. Sherson and www.scienceathome.org team

Center for Community Driven Research (CODER) at department of Physics and Astronomy, Aarhus University, Denmark

Abstract:

In the following we explain initial steps in turning the www.scienceathome.org (S@H) citizen science platform into a massive online laboratory for the systematic study of human cognition and behavior with the intention of better understanding the general principles underlying human problem solving (HPS). First, we point out that in our opinion natural science research games such as our Quantum Moves game constitute a so-far unexplored resource for the investigation of HPS. Next, we explain our ongoing efforts to implement reinforcement machine learning methods in order to explicitly highlight the differences between human and state-of-the-art machine problem solving strategies. Next, we describe first work on a cognitive science inspired Quantum Minds game based on the dynamics of Quantum Moves but carefully designed in order to be able to extract information about the human learning process while searching for optimal strategies. Finally, we describe the work to create a palette of HPS games with similar mathematical framework but gradually increasing complexity with the aim to understand how HPS changes with game dimensionality.

Background:

The disciplines focusing on the study of processes of problem solving have traditionally influenced each other and e.g. have shared assumptions about the mind as an information processor [1] or symbolic AI [2].

Machine learning and AI strategies utilizing big data have recently drawn attention with remarkable achievements in e.g. Go [3], Arcade games [4], as well as outstanding progress regarding speech or image recognition [5]. These successes have been won by brute force machine learning and *specialized* learning mechanisms [6]. But in order to develop more *general intelligence* based on reinforcement learning, leading AI scholars demand a “return to the roots”, i.e. finding further inspiration in cognitive sciences and work on human problem solving (HPS) processes [7].

In the last few decades research on HPS has identified error-prone heuristics [8,9] and simple adaptive heuristic algorithms [10]. Recent developments in this field offer a new theoretical foundation for understanding how humans select between two distinct modes of action: goal-directed (model-based) and habitual (model-free) learning [11].This perspective addresses how humans develop and select alternative problem representations [12] and explains how humans are able to abstract from past knowledge [13] to extrapolate to novel domains [14] as well as what is the role of perceptual learning (i.e. long-term performance increase due to perceptual experience) in constructing these models [15]. Despite these advances many fundamental questions remain:

a) Can human behavior in artificial, simplified laboratory tasks be generalized to more complex and natural settings? [16]
b) Do people recruit categorically different decision making mechanisms for perceptual/motoric tasks versus cognitive tasks and is the resulting performance systematically different? [17,18]
c) What is the relative contribution to complex problem solving skills of crystalline and fluid intelligence [19] and model-based vs model-free search strategies [11]?
d) Are there different basic decision processes in different domains of behavior/culture, or is there one evolutionarily selected fundamental decision module? [20]
e) How do cognitive constraints on an individual level hamper or aid the choice of problem solving strategies?

The current fault line between artificial and HPS is continually probed by the emerging field of citizen science [21]. It actively seeks out hard scientific problems and utilizes the process of gamification - the use of game design elements in non-game contexts [22] – to build games enabling ordinary citizens to contribute to the solution. This has led to scientific contributions in e.g. the study of protein [23] and RNA [24] folding, neuron mapping [25], cancer research [26], astronomy [27], and quantum physics [28]. The massive user engagement provides an excellent but so far largely unexploited platform for the systematic study of HPS processes.

The scienceathome.org HPS project

In the following we explain initial steps in turning the www.scienceathome.org (S@H) citizen science platform into a massive online laboratory for the systematic study of human cognition and behavior with the intention of better understanding the general principles underlying HPS. First, we point out that in our opinion natural science research games such as our Quantum Moves game constitute a so-far unexplored resource for the investigation of HPS. Next, we explain our ongoing efforts to implement reinforcement machine learning methods in order to explicitly highlight the differences between human and state-of-the-art machine problem solving strategies. Next, we describe first work on a cognitive science inspired Quantum Minds game based on the dynamics of Quantum Moves but carefully designed in order to be able to extract information about the human learning process while searching for optimal strategies. Finally, we describe the work to create a palette of HPS games with similar mathematical framework but gradually increasing complexity with the aim to understand how HPS changes with game dimensionality.

Shared mathematical framework: The games, or the underlying quantum optimal control problem, and the social science analyses can be formulated using the framework of fitness landscapes [29-31] in which the coordinates of the landscape are formed by the input variables and the height is determined by the quality of the solution for any given set of input variables. This framework quantifying e.g. the choice between "exploitation" of acquired knowledge and "exploration" of new potentials is also common in computer science fields such as numerical optimization and AI [32,33]. This shared mathematical language allows for an interdisciplinary approach where insights from the fields can seamlessly be interchanged.

The Quantum Moves game: The first major citizen science game from S@H is Quantum Moves (available on mobile platforms and for PC download on S@H) optimizing the manipulation of a single ultra-cold atom using a trap realized by an ultra-focused laser beam of relevance for the creation of a quantum computer. In the game, which has so far been played by 200,000 citizen scientists from around the world. The player sees a liquid-like representation of the quantum state of an atom in the dynamically controlled potential landscape. The goal is to quickly move the atom into the target area without having any residual sloshing of the atom. The state of the atom is updated live using the quantum physical equation of motion, the timedependent Schrödinger equation, in response to the actions of the players and at the end a score based on the overlap with the desired state is calculated and shown. Citizen science and social science games exploring HPS are quite often puzzles or involve pattern recognition. Quantum Moves is different, since it is a dynamic game that requires quick responses and proficient hand-eye coordination. In [28] we demonstrated that the players found solutions outcompeting those obtained by standard quantum optimal control algorithms. Moreover, the player solutions were used as seeds for a human-computer hybrid optimization scheme in which the players intuitively explored the vast 1000-dimensional landscape and the computer algorithm performed efficient local optimization. This combined strategy thoroughly outcompeted the previously best results. It is interesting that a very large fraction of the player results were situated in the very small subregion of the vast optimization space leading to optimal solutions. It thus seems that the players were able to combine their normal everyday intuition for manipulating regular liquids with the characteristic quantum features in the game to form a highly efficient, heuristic *quantum intuition*. In general games very often introduce unfamiliar rules of physics into the game play as challenges and thereby represent unique opportunities for the systematic study of the emergence of these domain specific intuitions. In particular, in Quantum Moves the players have to get accustomed to a weird quantum property called quantum tunneling, which is the ability of quantum particles to tunnel through energy barriers despite classically not having sufficient kinetic energy to make it over the top. In the game, this phenomenon manifests itself as the liquid flowing uphill under certain circumstances. The successful player therefore automatically recalls his or her everyday experience of carrying water glasses and coffee cups and intuitively accepts that in this particular game water *can* in fact flow uphill.

The quantum Minds game: in order to systematically investigate the dynamical process of the formation of the choice of strategy we developed the cognitive science game, Quantum Minds (available for PC download on S@H). The game mechanics is very similar except that it has four levels of increasing difficulty and in order to win the game a certain threshold score has to be surpassed three times in a row on any given level. The basic hypothesis is to use this as a proxy for true learning of the level as opposed to simply finding the good solution once by chance. One interesting feature of this game is the introduction in the advanced levels of (slightly unphysical) zones of death in which the players looses the particular gameplay if a small fraction of the atom enters. The advanced level are sufficiently difficult that no player will solve the challenge in the first trial. The fact that a score of zero is output until successful poses an extreme challenge for computer optimization because no gradient information of the fitness function can be formed between iteration. In contrast, in humans we can study the multi-dimensional processing going into formulating a meaning of "nearly making it" or "almost not dying" and how this is translated into game hypothesis or heuristics to be tested in subsequent game play. The plan is then to use machine learning feature extraction in order to i) extract the features characteristic of good solutions and ii) identify the influence of learning style on the speed and likelihood of growth of these good features as the player progresses. In a step towards the benchmarking of online social and cognitive science we aim to study this process in three distinct settings: a) "in-the-wild" by a few thousand users of S@H, b) in a semi-structured class-room setting involving roughly 200 students, and 3) in small-scale controlled laboratory studies.

Machine learning on the Quantum Moves game: as more and more human behavioral data is collected, it becomes increasingly important to benchmark the human performance with the newest reinforcement learning technique in order to be able to realistically assess any potential for the useful generalization of the identified HPS strategies. Currently, we are investigating Trust Region Policy Optimization (TRPO) for the Quantum Moves game. For state parametrization, we consider the surface and the liquid on top as two functions, resolved into 128 linearly spaced points along the horizontal axis. Concatenating the two lists leaves us with 256 points for a frame. To account for the sloshing of the liquid multiple frames (i.e. 4 frames) are necessary to capture the state, leaving us with a total of 1024 points. Finally, we have represented the target area as a step function which is 0 outside of the area, and 1 inside. By resolving this function into 128 points, and appending it to the rest of the state, we end up with a dimensionality of 1152. The target area is only included in the state once, because it does not change over time. The 1024 points representing the liquid and the surface can take any value from [0;1]. For action characterization, we implement an agent with continuous actions along two dimensions to mimick the freedom of the players to take a wide variety of different step sizes and directions. Initial results of the TRPO algorithm do exhibit learning from initially completely failed attempt to scores representing non-zero transport of the atom. However, the score is still fairly low (~10% of the max score) and in addition the algorithm at present does not succeed in maintaining this score but "unlearns it". Despite not being competitive with the best player generated scores at present, these first exploratory results do reveal the ability of the reinforcement learning algorithm to exhibit some learning from scratch and we expect further work to improve the overall score dramatically.

Creating a ladder of increasingly complex HPS games: Finally, we turn to the ambition of the S@H project to study HPS systematically in a number of mathematically well defined problems with gradually increasing complexity. The first game is the Alien Game (available for PC download at S@H) in which an 8 bit puzzle representing fixed underlying landscapes is presented to the player in the form of a puzzle posed by an alien visitor to our planet [34]. The player has 25 trials and in each trial an arbitrary number of bits can be flipped. The main research question is to investigate the distribution among three canonical move types: random, locally optimizing and distant model-based moves. Then, the likelihood for a given player to change strategy based on game feedback and advice from other players is studied. Two game variants have so far been studied. In the first one, the player could freely flip the bits at each choice. This made the identification of the move type very complicated since in practice it can be very hard to distinguish between random and model-based moves when one does not have any information of the particular model currently being considered by the player. Instead, a second game mode was developed in which the player actively presses a button for one of the three general strategies and thereafter is presented with the possible moves within this category. While this resolves the problem of identifying the type of search mode being pursued it remains to be investigated if this representation invokes a large bias in the search behavior. The second game is the Crystal Crop Fever (CCF), which has been tested in two different design iterations by roughly 100 participants, and will be available for public download within a few months. Whereas the Alien Game had an implicit 8-dimensional landscape, the game play in CCF is constructed directly around the landscape metaphor by having a 2D landscape with 80x80 tiles. It is a multiplayer game accommodating up to 16 simultaneous players. Again there are 25 game rounds and in each round each player selects a previously unexplored tile to be investigated and then the scores of all player selections is revealed. In the Alien Game, the player advice was built in as a static simulated advice, whereas CCF features true collaborative problem solving [35]. In both Alien Game and CCF landscapes of varying complexity can be implemented and the human response studied. We believe that this systematic low, dimensional HPS study will then inform subsequent studies on the much more high dimensional and realistic natural science research challenges such as Quantum Moves. For these studies it will be crucial to convince the same players to all three games in order to be able to cross correlated behavior as efficiently as possible. With a firm understanding of HPS processes across an increasing dimensionality scale, we hope to be able to extrapolate more confidently to truly complex social interaction dynamics than is possible today.

References:

1. Simon, H. A. (1990). Invariants of human behavior. Annual review of psychology, 41(1), 1-20

2. Haugeland, J. Artificial intelligence: The very idea. MIT press, (1989)

3. AlphaGo 2016: https://deepmind.com/alpha-go

4. Mnih, V. et al. Human-level control through deep reinforcement learning. Nature, 518(7540), 529-533 (2015).

5. LeCun, Y. et al. Deep learning. Nature 521.7553: 436-444 (2015)

6. Schölkopf, B. Artificial intelligence: Learning to see and act. Nature, 518(7540), 486-487 (2015)

7. Lake B.M. et al. Building machines that think like people. eprint arXiv:1604.00289 (2016)

8. Tversky, A.; Kahneman, D. Judgment under uncertainty: Heuristics and biases. Science, 185, 1124–1131 (1974)

9. Kahneman, D. Thinking, Fast and Slow, Farrar, Straus and Giroux (2011)

10. Gigerenzer, G. et al. Heuristics: The foundations of adaptive behavior. Oxford: Oxford University Press (2011)

11. Doll, B. et al. The ubiquity of model-based reinforcement learning. Current opinion in neurobiology, 22(6), 1075-1081 (2012)

12. FitzGerald, T. H., et al. Model averaging, optimal inference, and habit formation.

13. Ullman, T. et al. Theory acquisition as stochastic search Proceedings of the 32nd Annual Meeting of the Cognitive Science Society, COGSCI 2010, Portland, Oregon, August 11-14(2010)

14. Kemp, C. and Tenenbaum J.B. The discovery of structural form. Proceedings of the National Academy of Sciences 105.31: 10687-10692 (2008)

15. Watanabe T. et al. Perceptual learning: toward a comprehensive theory. Annual review of psychology 66: 197 (2015)

16. Smaldino P.E. & Richerdsen P.J. The origins of options. Frontiers in neuroscience 6: 50-50.(2011)

17. Trommershäuser J. et al. Decision making, movement planning and statistical decision theory. Trends in cognitive sciences, 12(8), 291-297 (2008)

18. Glaser C. et al. Comparison of the distortion of probability information in decision under risk and an equivalent visual task. Psychological science, 23(4), 419-426 (2012)

19. Engle, R. W. et al. Working memory, short-term memory, and general fluid intelligence: a latent-variable approach. Journal of experimental psychology: General, 128(3), 309 (1999)

20. Hastie, R. Problems for Judgment and Decision Making, Annual Review of Psychology, p. 672 (2001)

21. Kullenberg, C., & Kasperowski, D.. What Is Citizen Science?–A Scientometric Meta-Analysis. *PloS one*, *11*(1), e0147152 (2016).

22. Deterding, S, et al. From game design elements to gamefulness: defining gamification. Proc 15th Int Acad MindTrek Conf: Envisioning Future Media Environments, 9-15. (2011)

23. Cooper, S. et al. Predictingprotein structures with a multiplayer online game. Nature 466, 756-760 (2010), https://fold.it/portal

24. Lee, J. et al. RNA design rules from a massive open laboratory. Proc. Natl Acad.Sci. USA 111, 2122–2127 (2014).

25. Kim, J. S. et al. Space-time wiring specificity supports direction selectivity in the retina. Nature 509, 331–336 (2014)

26. Curtis, C. et al. The genomic and transcriptomic architecture of 2,000 breast tumours reveals novel subgroups. Nature 486(7403):346-52 (2012)

27. Lintott, C. et al. Galaxy Zoo 1: data release of morphological classifications for nearly 900,000 galaxies. Mon. Not. R. Astron. Soc. 410, 166–178 (2011)

28. Sørensen, J.J. et al. Exploring the Quantum Speed Limit with Computer Games. Nature 532, 210-213 (2016)

29. Sewall Wright. The roles of mutation, inbreeding, crossbreeding, and selection in evolution. *Proceedings of the Sixth International Congress on Genetics*, 1:355–366, 1932.

30. Herschel Rabitz, Re-Bing Wu, Tak-San Ho, Katharine Moore Tibbetts, and Xiaojiang Feng. Fundamental principles of control landscapes with applications to quantum mechanics, chemistry and evolution. In *Recent Advances in the Theory and Application of Fitness Landscapes*, pages 33–70. Springer Science + Business Media, 2014. doi: 10.1007/978-3-642-41888-4_2.

31. Daniel A. Levinthal. Adaptation on rugged landscapes. *Manage. Sci.*, 43(7): 934–950, 1997. ISSN 00251909, 15265501. URL http://www.jstor.org/stable/2634336.

32. Katherine M. Malan and Andries P. Engelbrecht. Fitness landscape analysis for metaheuristic performance prediction. In *Recent Advances in the Theory and Application of Fitness Landscapes*, pages 103–132. Springer Science + Business Media, 2014. doi: 10.1007/978-3-642-41888-4_4.

33. Mersmann O, et al. (2011) Exploratory landscape analysis in *Proceedings of the 13th annual conference on Genetic and evolutionary computation*. (ACM), pp. 829–836.

34. Vuculescu O., Bergenholtz C., Kock M. and Sherson J. Fitness landscapes in organizational theory: research challenges and future directions. Druid 20th Conference (2016)

35. Mason, W. & Duncan J. Watts Collaborative learning in networks. Proceedings of the National Academy of Sciences 109.3: 764-769 (2012)