Anticipation as a Mechanism for Complex Behavior in Artificial Agents

Heinrich Mellmann, Benjamin Schlotter, Lea Musiolek and Verena V. Hafner

Adaptive Systems Group, Department of Computer Science, Humboldt-Universität zu Berlin, Berlin, Germany
mellmann@informatik.hu-berlin.de

Abstract

Anticipation is a skill that enables complex decision making in humans and other biological agents. We review different implementations of anticipatory behavior in robots and give an overview on anticipation in biological systems. Based on an example of anticipatory behavior in humanoid robots, we discuss decision making and anticipation in artificial agents. We show that anticipation can enable fast decisions in highly dynamic and complex situations. Our findings are supported by experimental results performed in simulation and on real robots in large scale experiments.

Introduction

Humans and many other biological agents are able to anticipate the consequences of their actions to a certain extent and adapt their decisions based on available information on current and future states. This makes anticipation a powerful tool enabling complex decision making and behavior. This should also hold for artificial agents that interact in complex environments.

Winfield and Hafner (2018) consider anticipation in artificial agents through the mechanism of predictive internal models which generate a prediction of a particular state in the future. These models can generate predictions of the self, others, and the environment, and could be either predefined or learned and adapted through experience.

In robotics, the skill to self-model and to anticipate has been demonstrated in several studies and experiments. One of the earliest implemented in real hardware was the starfish robot presented in Bongard et al. (2006). There, a four-legged robot was able to simulate its locomotion and was therefore able to adapt to changes in morphology such as removed leg parts. In Mirza et al. (2008) information-theoretic measures are used to learn a self-model of a robot from experience. In Schillaci et al. (2016), a humanoid robot learned internal models of sensorimotor relationships through an exploratory phase inspired by infants' body babbling. The acquired models were used for decision making in tool-use as presented in Schillaci et al. (2012). In Matsumoto and Tani (2020) predictive coding and active inference is used for goal directed planning with only partial knowledge.

Anticipation has also been studied in human-robot interaction scenarios on an iCub robot, implementing human intention reading (Duarte et al., 2018) and multi-modal models (Dermy et al., 2019).

Other approaches do not learn a self-model of the robot but directly predict the consequences of both its own actions and those of others. Blum et al. (2018) showed that such a mechanism improves the safety for a scenario in which a robot needs to traverse a corridor without bumping into humans and other robots. In Mellmann and Schlotter (2017); Mellmann et al. (2017), forward simulations in soccer playing robots allow for fast decision making.

Anticipation in Biological Systems

Anticipation plays an important role in the cognition of biological agents, and from an evolutionary perspective, it fulfills a number of crucial functions: It provides additional information to help the decision making process, in order to plan next actions or changes of actions. It facilitates sensory attenuation, and it guides attention mechanisms. In this paper, we focus on the first aspect related to decision making.

In recent decades, cognitive neuroscience has turned its attention to the predictive capacities of biological agents. Wolpert and Kawato (1998) proposed the existence of an internal representation of sensory inputs and motor output signals in the brain responsible for sensorimotor prediction processes. Recent research found evidence for neural representations of probabilistic predictions in human brain imaging data (Aitchison and Lengyel (2017), Kopp et al. (2016), Ostwald et al. (2012)) as well as in human nonverbal social behaviour (Candidi et al. (2015); Pezzulo et al. (2017)). In Devaine et al. (2014), an economic gaming experiment with model comparisons showed an advantage for agents with "mentalizing" skills realized as recursive Bayesian predictions. Interestingly, humans were able to recruit these abilities only when they believed they were playing against a human opponent. Recursive Bayesian predictions were also used in Bordallo et al. (2015), where artificial agents observe and predict the movement of other agents by attributing a goal position to each observed agent. One of the most

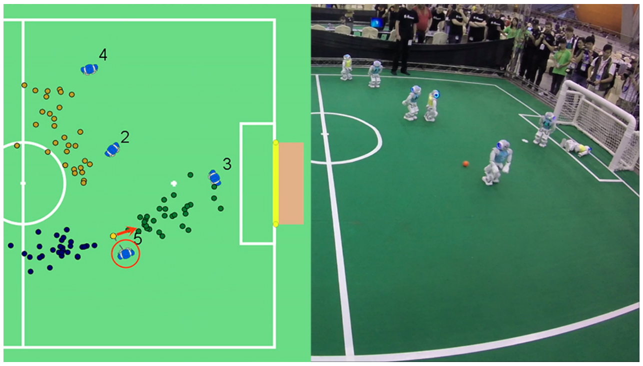

Figure 1: Situation from a robot soccer game. The robot at the ball simulates possible outcomes of three different kick actions illustrated by circles of different colors. The selected action is depicted by the red arrow.

notable attempts to unify predictive approaches is the Free-energy principle (Friston, 2010). Nasuto and Hayashi (2016) discuss anticipation as a bridging concept between philosophy, and biological and cognitive sciences.

Experimental Setup and Results

In this section, we discuss experimental results demonstrating how anticipation and internal simulation can be used to realize complex behavior in artificial agents presented in Mellmann and Schlotter (2017) and Mellmann et al. (2017). The approach is inspired by the decision problem within the RoboCup domain, where humanoid robots are playing soccer autonomously and have to perceive the environment, make decisions and execute actions in real time.

Consider a simple situation in which a robot is approaching a ball on the soccer field and is faced with a decision to select from a number of discrete actions, e.g., front kick, side kick, and a continuous action - adjusting the direction by turning around the ball. Despite the apparent simplicity, this task can present a significant challenge in a complex situation of a soccer game. The outcome of a particular action may depend on a wide variety of environmental factors, such as the robot's position on the field or the location of other players. In addition, the robots' perception of the situation is often uncertain, noisy and incomplete, and execution of the actions is subject to noise and uncertainty as well.

The overall decision process can be split into three distinct phases: *predict*, *evaluate* and *decide*. The robot envisions the possible outcomes of the available actions. Each of the outcomes is evaluated and the option with the best outcome is selected. In this particular scenario the predicted outcome of a kick action is the final position of the ball.

In order to deal with computational complexity of the uncertainty, the simulation of an action is split into a number of simple deterministic simulations (*samples*) in a Monte-Carlo fashion based on the uncertainties in the estimated state and in the model of the action. Stripped of the uncertainty the prediction of the individual possible positions of the ball and their respective evaluation become significantly simplified. The individual evaluations are recombined to represent the uncertainty in the outcome of the action and are compared with those of other actions to inform the overall decision. Figure 1 illustrates the decision process in a game situation.

The decision scheme has been implemented and evaluated in simulation as well as on real robots in a competition setting. Simulated experiments have shown that a robust decision can already be achieved with a small number of samples. To enable evaluation under real world conditions, a data set consisting of videos of the games, perceptual information and behavior states of the robots was collected from a series of games during international competitions. The decisions made by the robots were annotated qualitatively by a human operator. The results presented in Mellmann et al. (2017) have shown that even in its simple form the decision scheme is able to provide fast and robust decisions outperforming a rule based baseline.

Our ongoing work is focused on including simulation of decisions of other players (2nd level anticipation) and analysis of the emergent cooperative and antagonistic behavior.

Discussion

The results of the experiments have shown that anticipation and forward simulation can provide a versatile and extensible yet simple mechanism for inference of decisions.

An open question is still how to evaluate anticipatory behaviour. Certain behavior observed in biology that may appear anticipatory to an outside observer might be purely reactive in nature. More precisely, is it possible to distinguish whether observed behaviour of (biological) agents has been generated by an explicit predictive mechanism or by simple reactive responses? An approach would be the comparison of the performance of reactive and predictive models. If models equipped with predictive capacities outperformed ones with purely reactive skills at explaining observed behavior, this would suggest that the underlying mechanism is anticipatory. Suitable experiments for future studies could involve obstacle avoidance and leader-follower behavior.

It would be also interesting to study under which circumstances implicit anticipatory behaviour can emerge. The evidence from observations in biological systems as well as from experiments with artificial agents underlines the importance of anticipation as a mechanism for behavior. It could be speculated that anticipation necessarily emerges in cognitive systems with enough complexity.

Acknowledgements

This work was supported by the EU H2020 programme under grant agreement No. 824069 (HIVEOPOLIS) and by the DFG under Germany's Excellence Strategy – EXC 2002/1 "Science of Intelligence" – project number 390523135.

References

Aitchison, L. and Lengyel, M. (2017). With or without you: predictive coding and bayesian inference in the brain. *Current Opinion in Neurobiology*, 46:219 – 227. Computational Neuroscience.

Blum, C., Winfield, A. F. T., and Hafner, V. V. (2018). Simulation-based internal models for safer robots. *Frontiers in Robotics and AI*, 4:74.

Bongard, J., Zykov, V., and Lipson, H. (2006). Resilient machines through continuous self-modeling. *Science*, 314(5802):1118–1121.

Bordallo, A., Previtali, F., Nardelli, N., and Ramamoorthy, S. (2015). Counterfactual reasoning about intent for interactive navigation in dynamic environments. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pages 2943–2950.

Candidi, M., Curioni, A., Donnarumma, F., Sacheli, L. M., and Pezzulo, G. (2015). Interactional leader-follower sensorimotor communication strategies during repetitive joint actions. *Journal of The Royal Society Interface*, 12(110):20150644.

Dermy, O., Charpillet, F., and Ivaldi, S. (2019). Multi-modal intention prediction with probabilistic movement primitives. In Ficuciello, F., Ruggiero, F., and Finzi, A., editors, *Human Friendly Robotics*, pages 181–196, Cham. Springer International Publishing.

Devaine, M., Hollard, G., and Daunizeau, J. (2014). The social bayesian brain: Does mentalizing make a difference when we learn? *PLOS Computational Biology*, 10(12):1–14.

Duarte, N. F., Tasevski, J., Coco, M. I., Rakovic, M., and Santos-Victor, J. (2018). Action anticipation: Reading the intentions of humans and robots. *CoRR*, abs/1802.02788.

Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138.

Kopp, B., Seer, C., Lange, F., Kluytmans, A., Kolossa, A., Fingscheidt, T., and Hoijtink, H. (2016). P300 amplitude variations, prior probabilities and likelihoods: A bayesian erp study. *Cognitive, Affective, & Behavioral Neuroscience*, 16.

Matsumoto, T. and Tani, J. (2020). Goal-directed planning for habituated agents by active inference using a variational recurrent neural network. *Entropy*, 22(5):564.

Mellmann, H. and Schlotter, B. (2017). Advances on simulation based selection of actions for a humanoid soccer-robot. In *Proceedings of the 12th Workshop on Humanoid Soccer Robots, 17th IEEE-RAS International Conference on Humanoid Robots (Humanoids), Madrid, Spain.*

Mellmann, H., Schlotter, B., and Blum, C. (2017). Simulation based selection of actions for a humanoid soccer-robot. In Behnke, S., Sheh, R., Sarıel, S., and Lee, D. D., editors, *RoboCup 2016: Robot World Cup XX*, pages 193–205, Cham. Springer International Publishing.

Mirza, N. A., Nehaniv, C. L., Dautenhahn, K., and te Boekhorst, I. R. J. A. (2008). Anticipating future experience using grounded sensorimotor informational relationships. In *ALIFE2008*, pages 412–419.

Nasuto, S. J. and Hayashi, Y. (2016). Anticipation: Beyond synthetic biology and cognitive robotics. *Biosystems*, 148:22 – 31. What Synthetic Biology can offer to Artificial Intelligence.

Ostwald, D., Spitzer, B., Guggenmos, M., Schmidt, T. T., Kiebel, S. J., and Blankenburg, F. (2012). Evidence for neural encoding of bayesian surprise in human somatosensation. *NeuroImage*, 62(1):177 – 188.

Pezzulo, G., Iodice, P., Donnarumma, F., Dindo, H., and Knoblich, G. (2017). Avoiding accidents at the champagne reception: A study of joint lifting and balancing. *Psychological Science*, 28(3):338–345. PMID: 28103140.

Schillaci, G., Hafner, V. V., and Lara, B. (2012). Coupled inverse-forward models for action execution leading to tool-use in a humanoid robot. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction*, HRI '12, page 231–232, New York, NY, USA. Association for Computing Machinery.

Schillaci, G., Hafner, V. V., and Lara, B. (2016). Exploration behaviors, body representations, and simulation processes for the development of cognition in artificial agents. *Frontiers in Robotics and AI*, 3:39.

Winfield, A. F. T. and Hafner, V. V. (2018). Anticipation in robotics. In Poli, R., editor, *Handbook of Anticipation: Theoretical and Applied Aspects of the Use of Future in Decision Making*, pages 1–30, Cham. Springer International Publishing.

Wolpert, D. and Kawato, M. (1998). Multiple paired forward and inverse models for motor control. *Neural Networks*, 11(7):1317 – 1329.