Some Notes on Anticipation in AI and Robotics

Alessandro Saffiotti
AASS Cognitive Robotic Systems Lab
Örebro University, 70182 Örebro, Sweden
asaffio@aass.oru.se

The concept of anticipation is found across several fields, and hence it is a topic that can, and should, be studied in a cross-disciplinary way. These notes are meant to contribute to this study, by providing an accessible account of the many ways in which anticipation is used by intelligent robots. Although we always refer to robots in our narration, all the considerations below apply to AI agents in general.

1 Robots that live in the present

Back in 1948, in Bristol, the cybernetician William Grey Walter built two of the first electronic autonomous robots, named Elmer and Elsie [1]. These robots coupled photoreceptors to the motors controlling their wheels, so that their motion would depend on the direction and intesity of a light source. Elmer and Elsie are early examples of what we now call *reactive robots*: robots whose actions only depend on their immediate perceptions. These robots take decisions based on the current state of their environment, limited to those parts of the environment that their sensors can observe.

Reactive robots are very simple but they may be surprisingly powerful: Braitenberg [2] presents a systematic study of simple robots like Elmer and Elsie, and Brooks [3] energetically argued that we can go very far in realizing intelligent robots with a purely reactive approach, avoiding any internal mental representation. Reactive robots, however, are limited by their temporal and sensory horizon: they can reach a light source if this in view of their sensors, but they cannot do so if it becomes occluded by a moving object or if it is behind a wall, even the same source had been seen before. To do so, robots would need to have some memory of the past.

2 Robots that remember the past

In most of today's robots, behavior is influenced not only by the current perceptions but also by the memory of past perceptions. A robot with memory could, for instance, reach a light source that has been seen before but it is currently switched off or occluded; and it could decide to open a box that contains the key to unlock a door, if it has previously seen someone putting the key in that box.

Memories of the past can be kept as individual data points, or they can be processed and compiled into a more abstract *model* of the world. Both approaches are common in robotics, and they are often combined. Consider the basic "self-localization" problem, that is, the problem for a robot to know where it is in the environment. The robot can memorize images taken at different places, so as to recognize these places when it is back there [4]; or it can use all its perceptual experiences to build a global, structured map of the environment [5]. In cognitive sciences, these ways to store past experiences are referred to as *episodic memory* and *semantic memory* [6], respectively.

3 Robots that care for the future

Action and sensing happen locally, in the here and now of the robot, but the goals they are aimed at, or the undesired consequences they can bring about, may lie far away in time and space. An intelligent robot needs to connect its local actions to its global goals and desires [7]. World models, like maps, allow robots to reason about information beyond the spatial horizon of their sensors. To reason beyond their temporal horizon, robots use *predictive models*. Such models allow a robot to forecast the possible evolutions of the world and how its own actions may affect those evolutions, and to use this information to decide its actions. Robot that do so are anticipatory systems in the sense of Rosen [8].

Robots can use predictive models in the context of several different functionalities.

- *Anticipatory control.* When you want to catch a ball thrown to you, you start moving your arm before the ball is in reach, so your hand will be at the right position by the time ball arrives there. Similarly, when a robot moves in a dynamic world, it must take into account not how the world is now, but how the world will be by the time the motion is completed. Techniques like feedforward control and model-predictive control [9] use a predictive model of how movements are performed, and how the world evolves, to realize this type of anticipatory movements.

- *Goal achievement.* When you invite your friends for dinner, you carefully plan what to buy and how to organize your cooking activities, depending on the preferences of your friends and on what you already have in your fridge. Similarly, robots that are given new goals on the fly, rather than having a goal built-in, must plan a course of actions to achieve those goals given the current context. Automatic planning techniques [10] use predictive and causal models to explore possible courses of actions, to understand how they will affect the future state of the world, and to eventually chose one course of action that will achieve the intended goals.

- *Proactive behavior.* When you see someone in front of a door who is carrying an heavy weight, you will probably open the door without being asked. Warneken and Tomasello [11] note that such proactive behavior is already present in 18 month old infants, and claim that proativity has given humans an evolutionary advantage compared to other species. Most current robots wait for requests by humans, but a few recent works study robots that proactively generate their own goals [12]. To do so, they use predictive models to forecast the possible future evolutions and how the robot's actions may affect them, and to proactively decide to initiate actions that lead to desirable evolutions or that avoid undesirable ones.

4 Discussion

The above notes shows how intelligent robots use anticipation to connect the actions they performs now to their final goals in a dynamically changing world. Despite the central role played by anticipation in AI and robotics, anticipation is rarely studied as a topic *per se*. Anticipation is usually included as one of the necessary ingredients of other, well defined areas of investigation in robotics, like model predictive control or task planning, but it is rarely considered an independent area of investigation, and no study in the robotic literature analyzes this problem in a systematic way. We hope that these notes will be instumental in initiating such a general study of anticipation in AI and robotics.

Acknowledgemnts

These notes grew from a dialogue with Luis de Miranda, who continuously challenged me from a philosopical point of view, in the context of the CREA initiative (`crea.aass.oru.se`). The parts on proactive behavior reflect joint work with Jasmin Grosinger.

References

[1] W G Walter. An imitation of life. Scientific American 182(5), pp.42-45, 1950

[2] V Braitenberg. Vehicles: Experiments in synthetic psychology. MIT press, 1986.

[3] R A Brooks. Intelligence without representation. Artificial intelligence 47(1-3) pp.139-159, 1991.

[4] I Ulrich and I Nourbakhsh. Appearance-based place recognition for topological localization. Proc of the IEEE Int Conf on Robotics and Automation (ICRA) pp.1023-1029, 2000.

[5] B Kuipers. The spatial semantic hierarchy. Artificial intelligence 119(1-2) pp.191-233, 2000.

[6] L R Squire. Memory systems of the brain: a brief history and current perspective. Neurobiology of Learning and Memory 82(3) pp.171-177, 2004.

[7] A Saffiotti, K Konolige, EH Ruspini. A multivalued logic approach to integrating planning and control. Artificial intelligence 76(1-2) pp.481-526, 1995.

[8] R Rosen. Anticipatory Systems: Philosophical, Mathematical & Methodological Foundations. Pergamon Press, 1985.

[9] B Siciliano and O Khatib, eds. Springer handbook of robotics. Springer, 2016.

[10] M Ghallab, D Nau and P Traverso. Automated Planning: theory and practice. Elsevier, 2004.

[11] F Warneken and M Tomasello. The roots of human altruism. British Journal of Psychology 100(3) pp.455-471, 2009

[12] J Grosinger, F Pecora and A Saffiotti. Robots that maintain equilibrium: Proactivity by reasoning about user intentions and preferences. Pattern Recognition Letters 118, pp.85-93, 2019.