Alexandra Mark (Autor)

Development of an Anticipating Decision Making System Using Evolution and Learning

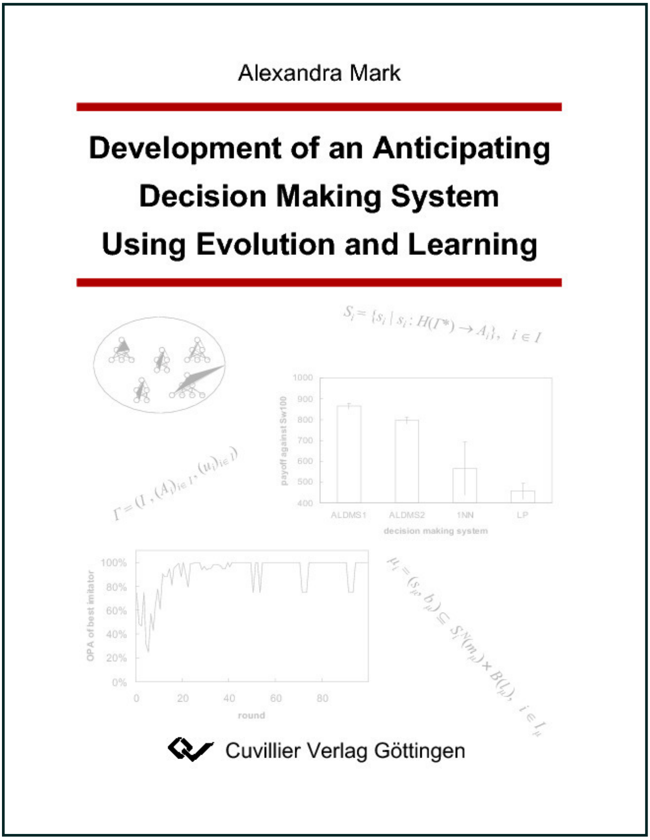

https://cuvillier.de/de/shop/publications/2302

Chapter 1

Introduction

1.1 Motivation

Decision making is of fundamental importance for our daily life. This is tightly connected to the cognitive abilities of thinking, learning, or building internal representations of the environment in which decisions should be taken. Since these concepts are very general, they are – and have already been for a long time – an important research topic in various scientific disciplines. However, the underlying mechanisms in the brain are only fragmentarily understood until today.

decision making: important research topic

The very first approaches to decision making have been of philosophical nature, e.g. by Aristotle. Later, in economics, researchers have tried to mathematically formalise decision making problems in order to be able to compute optimal solutions. Two prominent research fields of this area are decision theory and game theory. The task of game theory is to analyse and understand the behaviour of several rational players in an interaction by solving all players' decision problems simultaneously (Myerson 1997). In contrast to this, decision theory employs a probabilistic approach: the expected utility for an individual decision maker is maximised based on his belief about the future actions of the other decision makers with which he interacts. These approaches require many unrealistic assumptions, however, in order to reduce the complexity of the task. Psychological findings on human decision making have revealed that humans' behaviour significantly departs from these formal theories which assumed decision makers to be fully rational.

early approaches to decision making

It is not by accident that game theory was one of the first formal theories to investigate decision problems. Games provide an ideal test bed for decision problems, because of their diversity and arbitrary complexity – depending on the definition of the behaviour of the participating players. Moreover, games are an integral part of human life, as every child learns to make decisions in a variety of strategic games. More recent research has tried to avoid the problems of game theory by abandoning the rationality assumption and introducing simulation instead of pure mathematical analysis. Human behaviour is considered to a growing amount in order to develop suitable decision making models.

game theory

Humans use predictions in the form of internal reflections and pre-

human decision making

ventive anticipations in order to solve complex decision tasks. For competitive decisions humans employ various decision strategies which are adapted during the decision process in order to reflect the most current situation: new alternatives are generated, and old ones, which do not seem to be very promising, are discarded. Humans learn to identify the strategy of others and develop a suitable counter-strategy (Payne, Bettman, & Johnson 1993). The simulation approach of cognitive function contributes to explain the neuro-biological foundations, which allow the brain-internal simulation and evaluation of hypothetical action sequences before their execution (Hesslow 2002). One important area for prediction is to anticipate the future decisions of other persons, in order to be able to make good own decisions. Ramnani & Miall (2004) indicate, that humans use basically the same cognitive mechanisms for the preparation of own actions and for the prediction of the actions of another person.

definition

Depending on task and research area, different definitions of decision making are available. The following working definition is used within the present thesis.

Decision making is the system-internal process of repeated

- gathering of information about the changing environment,
- building of hypotheses about possible action alternatives and future consequences,
- selecting one of these alternatives according to a (possibly adaptive) selection mechanism considering short- and long-term consequences with regard to one or several final target(s).

Here a system can be both a computer program or a human being.

developed system

In the present thesis an anticipating and learning decision making system is developed, which implements human-inspired approaches to make decisions in strategic competitive games. The focus is on the implemented mechanisms and not on the scenarios. Games have been chosen as field of application, because of their historically importance for decision making and the possibility to model sophisticated strategies. In particular the modelling and predicting of an opponent's strategy and the hypothetical simulation and evaluation of potential counter-strategies play an important role in the developed system. This division into two parts – modelling and predictive reaction finding – makes the decision task easier and faster to solve. In order to assure good performance of the whole decision system, the accuracy of the opponent model has to be evaluated. Intuitively it is assumed that a better ability to predict the opponent actions (prediction accuracy) leads to a better result, provided that the prediction is used in an optimal way for further processes, which lead to the end result. However, we will see that this straightforward line of thought can be deceptive and misleading. Indeed, a careful revision of measures of prediction accuracy is needed.

1.2 Problem Formulation

research target

The target of this research has been the development of an adaptive decision making system. Since decision making is a very broad concept, it was decided to concentrate on strategic decision problems. "Strategic" is meant in the sense of making decisions, which need some beforehand planning and consideration of their future consequences. From a historical point of view, decision making problems are often solved by rule based methods. Contrary to this in the present thesis a simulation approach is pursued, which is more powerful towards complex decision problems in changing environments with possibly noisy or missing information. This research has been motivated by the concepts employed by humans to make decisions. A focus was set to the internal simulation of possible future events in the relevant environment in order to find a "well considered" decision. It was opted to employ evolutionary and learning techniques for the implementation of the developed concepts, since these methods are very robust and provide a natural way of adaptation.

1.3 Overview

related work

In Chapter 2 different aspects of decision making – how decision problems are regarded or solved in various scientific disciplines – are discussed. Doing this, a focus is put on the modelling of and the adaptation to an opponent in competitive interactions. Also research on human decision making – a rich source of inspiration for the development of artificial decision systems – is considered.

environment

Since each system needs a suitable environment in order to show its functionality and performance, in Chapter 3 the two relevant scenarios for the present thesis are described in detail: the prisoner's dilemma and rock-paper-scissors. Both are well known in the area of game theory, but within this thesis no game theoretic methods are employed.

developed system

The developed anticipating and learning decision making system is presented in Chapter 4. The different components, which contribute to the overall performance of the system, are motivated and explained in detail. Moreover three other decision making systems are introduced, which serve as benchmarks for the new system.

results

The results of this performance comparison are provided in Chapter 5. One important aspect here is the ability of the developed decision making system to adapt online to a changing opponent strategy.

During the analysis of the performance results an unexpected observation was made: the best predicting strategies did not reach the best payoff in the end. This is analysed in Chapter 6. Within this chapter also the development of a new accuracy measure is presented, which is able to measure the accuracy more objectively and can thus avoid the mentioned observation. The achieved results are summarised and discussed in Chapter 7.

Tables of detailed results are presented in Appendix A, followed by a list of game strategies in Appendix B, and abbreviations in Appendix

C. At the end an index of relevant keywords and the references are provided.

Chapter 2

Different Aspects of Decision Making

Scope

The biological and psychological aspects of human decision making are presented – with a focus on how internal simulations and predictions of future events can be achieved. It is also reviewed how researchers have transferred these concepts into the development of algorithms, methods, and applications. Moreover, classical decision making approaches are presented.

deciding between alternatives

Decision making being a very general concept, there is one task that reoccurs, namely to decide between several alternatives to solve a specific problem. Sometimes these alternatives are explicitly given, sometimes there may be an indefinite number of possible actions, but always some task specific conditions have to be met. The "Tower of London" task illustrates this type of decision problem very well. As described in Figure 2.1 the task is to bring several coloured balls from a start configuration into a given goal configuration with the minimal number of moves.

decision mechanisms

If there is no possibility to solve this task physically, e.g. because no balls are available, the most plausible way for a human to solve this problem, is to mentally simulate several possible movements of the balls. On that occasion the person anticipates possible future constellations and evaluates them in order to choose the most suitable one. Thus, internal simulation and anticipation are important problem solving methods for which relevant literature is reviewed in Sections 2.1 and 2.3.

In more complex decision problems several persons are involved. Therefore, the participating decision makers have to simulate and anticipate not only the own future actions, but also those of the other participants. Hence it should be useful to observe the others' behaviour and to learn a model of it in order to find the best own actions. In Sections 2.2 and 2.4 the importance of opponent modelling and learning is shown.

Since humans memorise not only one, but several possible action alternatives or strategies about future actions, a modularised storage of several strategic action patterns is plausible. Relevant research is presented in Section 2.5. Psychological experiments have revealed that hu-

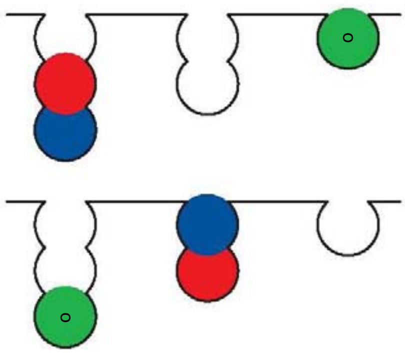

Figure 2.1: *Tower of London* task. The goal is to bring the three differently coloured balls from the bottom configuration with a minimal number of moves to the state shown on top. Only one ball may be moved at a time within the bottom panel, if no other ball blocks it from above. The ball can be put into any of the three columns (left, middle, right), if this is not already filled to its maximal capacity, and "falls" down to the lowest free position. On the left, there are three available positions, in the middle, there are two, and on the right side only one ball can be hold. Figure slightly adapted from Hesslow (2002).

mans do not decide purely rationally in general. Thus, extensions of older more formal theories, such as game or decision theory, have been suggested. This is described in Sections 2.6 and 2.7.

Last but not least, there is the question left, how humans can make decisions at all, i.e. which regions of the brain are involved and which functionality is observable in this complex process. Up to now no clear answer can be given to this question. See Section 2.8 for further information on this subject.

2.1 Anticipation

Que chacun examine ses pensées, il les trouvera toutes occupées au passé ou à l'avenir. Nous ne pensons presque point au présent, et si nous y pensons, ce n'est que pour en prendre la lumière pour disposer de l'avenir.[1]

Blaise Pascal, 1670 (in Pascal (1963), page 506)

As Pascal already stated in the 17th century, human thoughts hardly ever deal with the present, but more frequently plans for the future are done, possibly using experience from the past. This behaviour seems to be a fundamental characteristic of human thinking and decision making. Evidence for this has been put forward by different researchers and is reviewed in this section.

One might ask the question why prediction or anticipation are helpful to solve a certain task. According to Butz, Sigaud, & Gérard (2003a) anticipation can lead to guidance and stabilisation of behaviours, more

[1] English translation: If everybody examines his thoughts, he will find them all occupied with the past or future. We hardly ever think in the present, and if we do it, it is only for taking advise for planning ahead the future.

efficient reactivity, advantages in competitive scenarios, improved cooperative behaviour, faster adaptivity in dynamic environments by internal reflection, and behaviour optimisation by preventive anticipation of possible consequences.

In the last few decades experimental psychology has started to accept the concept of anticipation. Since the eighties of the 20th century, definite evidence of anticipatory behaviour in animals has been published, and it has been found that in natural environments anticipatory behaviour increases the chance of survival. More recent literature from cognitive psychology shows further evidence of distinct anticipatory mechanisms in learning, attentional processing, or object recognition tasks. Recent neuron imaging techniques and single-cell recordings provide further proof of anticipatory cognitive processes (Butz, Sigaud, & Gérard 2003b).

types of anticipatory behaviour

Butz, Sigaud, & Gérard (2003b) distinguish between the following types of anticipatory behaviour[2]:

payoff anticipations: Predictions of possible payoff values of different actions are considered and influence current decision making.

sensorial anticipations: Prediction of future stimuli does not influence current behavioural decision making, but current stimulus processing. For example expected sensory input might be processed faster.

state anticipations: Predictions about future states[3] influence current decision making.

terminology: prediction and anticipation

The term "prediction" comes from the Latin "prae-" (before) plus "dicere" (to say). Thus, in its original sense it refers to a statement that a particular event will happen in the future. More general, predictions also apply to the mental imagination that something will happen, without saying it. "Anticipation" refers to the act of looking forward, often in combination with the visualisation of a future event. Butz, Sigaud, & Gérard (2003b) use the term prediction in order to define different types of anticipations (see above). In the present thesis, however, both terms are used as synonyms with the following meaning: the mental (or system internal) imagination that an event will happen in the future.

examples for different types of anticipations

A typical payoff anticipatory method is model-free reinforcement learning, where the learned reinforcement values estimate action-payoffs. This topic is discussed in more detail in Section 2.4.

A typical example for sensorial anticipations is the approach by Gross, Stephan, & Seiler (1998), Gross, Heinze, Seiler, & Stephan (1999), and Stephan & Gross (2001). They present a sensorimotor approach in which the optical flow is predicted by neural networks in order to facilitate a robot navigation task. The importance of the prediction of sensorial consequences has already been published in the 19th century: "An anticipatory

[2]The fourth type (implicit anticipations) is not mentioned here, because no relevant predictions of the future are made.

[3]The term "state" represents a snapshot of the environment, in which the simulation takes place.

image, then, of the sensorial consequences of a movement, plus (on certain occasions) the fiat that these consequences shall become actual, is the only psychic state which introspection lets us discern as the forerunner of our voluntary acts." (James 1890).

An example for the state anticipation method are anticipatory learning classifier systems which contain an explicit prediction component, i.e. a set of rules, possibly combined with evolutionary methods (Butz, Sigaud, & Gérard 2003b). In the latter the knowledge is represented in a population of classifiers representing condition-action-anticipation rules, and conditions are generalised by a genetic algorithm. Another classifier system of this type is presented by Meyer, Ganascia, & Zucker (1997), who apply it to a strategic two-person game, called Alesia game. Within their system first all possible actions for one player are evaluated using game-theoretic methods. If no clear winning strategy can be found, then a model of the opponent player is built and used to predict the future opponent behaviour in order to find a suitable action for the predicting player.

Laird (2001) employs anticipatory mechanisms for the implementation of synthetic characters for computer games, which should show human-level playing behaviour in order to be challenging opponents. The reasoning algorithms of the characters are written with Newell's SOAR architecture. Such a character creates an internal representation of its assumption about the opponent's internal state and then uses its own tactics to predict the opponent's behaviour. A problem of this technique is that it is too slow to be applied extensively in online computer games.

Dubois (2003) describes the mathematical modelling of anticipatory capabilities in discrete and continuous systems. He defines a special type of anticipation systems, which shows multiple potential future states and a decision function to select one of these as actual state. His theory is tightly related to dynamical systems, but has not yet been applied to real world systems.

Within the present thesis anticipations, which influence and even facilitate current decision making, are of particular interest. This does not necessarily have to be the prediction of a state, but might also be the prediction of a behaviour. The idea is to anticipate the behaviour of other actors in an environment, in order to be able to react in an optimal way. To make such an anticipation, it is useful to first observe the behaviour of the other actors and then to learn to predict it. Therefore, it is useful to build models of these other actors, which can fulfil this task.

2.2 Opponent Modelling

Especially in the combination of decision making with game playing, opponent modelling is an important research area (Fürnkranz 2001). Psychological research has revealed that humans do not behave purely rationally in general (see Section 2.7). Not being rational means to make mistakes from time to time. Thus, an intelligent decision making system should recognise this weakness and be able to exploit it. Game theoretic