

## Characteristic Distributions in Multi-agent Systems\*

Stefan J. Johansson

Department of Software Engineering and Computer Science, Blekinge Institute of Technology, Sweden  
email: sja@ipd.hk-r.se

### Abstract

In game theory, iterated strategic games are considered harder to analyze than repeated ones. However, iterated games are in many cases more fit to describe the situation of artificial agents than repeated games. The reason being that they take into account previous actions of others, rather than just assigning each possible action a certain probability. We introduce the notion of *Characteristic Distributions* and discuss how they can be used to simplify and structure the analysis of strategies. This do not only provide a good basis for the agents to choose strategies for future interactions, but also helps in the design of environments that support certain types of agent behavior.

**Keywords:** Game Theory, Multi-Agent Systems, Characteristic Distributions, "No Free Lunch" Theorem

### Introduction

In Multi-Agent Systems, (MAS), agents interact with each other and the environment in order to meet their design objectives. From a game theoretic point of view, this process is a choice of strategies for playing games i.e. choosing behavioral patterns in a given environment, a choice that itself is a meta-game. Previous work on meta-games (the game of selecting a strategy for a game) include (Binmore & Samuelson 1992; Abreu & Rubinstein 1988; Rubinstein 1986).

We distinguish iterated games from repeated ones. In repeated games the players have no memory, while in the iterated games, the players remember all the previous actions made by others, i.e. they have a history of the game so far. Agents are in general able to remember previous actions taken by themselves and other agents, thus well-suited to be modeled by means of iterated games (Rosenschein & Zlotkin 1994).

Iterated strategic games are known to be harder to analyze and find equilibria in than repeated ones because of the exponentially increasing number of possible states (and thus taken into consideration when choosing the next move). Recently, promising attempts have been made, especially

\*The author wish to acknowledge the following colleagues for their support and their comments on various drafts of this work: Paul Davidsson, Rune Gustavsson, and Bengt Carlsson at Blekinge Institute of Technology, Magnus Boman at the Swedish Institute of Computer Science, Kristian Lindgren at Chalmers Institute of Technology and Santa Fe Institute, Fredrik Ygge at Enersearch, Johan Schubert at Ericsson and last, but not least, the anonymous reviewers, wherever you are.

Copyright © 2001, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

in the field of evolutionary game theory, to use e.g. adaptive dynamics to describe how equilibria might be reached among simple strategies in iterated games (Hofbauer & Sigmund 1998; Samuelson 1997; Weibull 1996). While these attempts try to answer the question "How are strategies behaving?", we will here try to focus on the question "What is the result of their behavior?" and "How can this result be used at the meta-level?". To help us answer these questions we will use *Characteristic Distributions* or *ChDs*.

The paper is organized as follows. First, we describe the distinction between agents and strategies and cover some formalities. Then some properties of the approach are discussed and exemplified. Finally some conclusions and further work are presented.

### Characteristic Distributions - Definitions

We begin by defining some central concepts:

**Definition 1 (Strategy)** *By a strategy we mean a function that projects the sequence of previous actions (made by itself and others) to the set of possible actions.*

**Definition 2 (Game)** *A game is a function where each combination of the choice of actions of  $n$  strategies is projected to payoffs among the strategies.*

The definition states that for each move, each participating strategy is assigned a certain value (the payoff of the move). From here on, only strategic, symmetric, two-player games are considered, unless explicitly said otherwise. The ideas may easily be extended to other types of games.

**Definition 3 (Size of a Game)** *The size  $d_g$  of a game  $g$  is the number of possible combinations of actions in each iteration.*

**Lemma 1** *The size of a game with  $n$  players where each strategy have  $k$  actions to choose from is  $k^n$ .*

**Definition 4 (Agent)** *The behavior of an agent corresponds to a meta-strategy that choose strategies for playing games.*

Other definitions, e.g. the one in (Wooldridge 1997)<sup>1</sup> are more explicit in the sense that they put the agents in a context by describing their properties and capabilities to interact with the environment. Our definition does not contradict

<sup>1</sup>"An *agent* is an encapsulated computer system that is situated in some environment, and that is capable of flexible autonomous action in that environment in order to meet its design objectives."

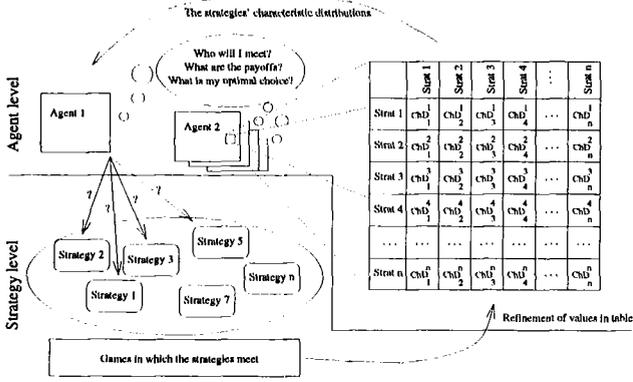


Figure 1: The two-layered approach to games. At the strategy level, games are played according to conditions given by the environment, whereas at the agent level the decision of what strategy to use include modeling other agent's choice of strategies, a process that is facilitated by the *ChDs*.

that, but raise the level of abstraction to strategies and games, instead of actions in specific environments.

**Definition 5 (Population)** A population (of strategies),  $P$  is the union of all strategies that the agents in a particular game consider.

There are two things worth noting here. Firstly, the agents may or may not consider the same set of strategies. Secondly, we distinguish between the agent level and the strategy level, where the strategies are purely projective, while the agents may have capabilities to reason about other agents choices of strategies, to analyze what game is the most suitable for describing the present situation in the environment, etc. For the agent to be able to compare different strategies in order to choose an appropriate one, it may use *ChDs* (see Fig. 1).

**Definition 6 (Characteristic Distribution)** The Characteristic Distribution (ChD) of a strategy  $s$  when meeting a strategy  $t$  in a game of size  $d$  is defined to be the  $d$ -entry matrix that describes the distribution of outcomes (distribution of combinations of moves made by the strategies). We denote this *ChD* by  $ChD_t^s$  and let  $ChD_t^s(i)$  be the  $i$ :th entry in this matrix.

The enumeration of the entries is reduced to one index  $i$ , although the *ChD*-matrix in the two-player case would require two index variables. This simplification is valid, as long as the enumeration of the entries is unambiguous. Also note that Fig. 1 has a matrix of *ChDs*, i.e. a matrix of matrices, since each *ChD*-entry in itself is a distribution of outcomes stored in a matrix.

**Lemma 2** Since all possible outcomes are considered, the sum of the entries, i.e.  $\sum_{i=1}^{d_g} ChD_t^s(i) = 1$ .

**Definition 7 (Population Distribution)** The population distribution  $P_d$  of a population  $P$  is the function  $P_d : P \rightarrow [0, 1]$  that tells the estimated probability of meeting each of the strategies in the population; especially, let  $P_d^a$  denote the population distribution of agent  $a$ .

**Lemma 3** Since  $P_d$  is a probability distribution,  $\sum_{t \in P} P_d(t) = 1$ , all considered strategies are in  $P$ .

**Definition 8 (Weighted Characteristic Distribution)**

We let  $\hat{ChD}_{P_d}^s$  denote the weighted *ChD*, i.e. the sum  $\sum_{t \in P} P_d^s(t) \cdot ChD_t^s$ .

**Lemma 4**  $\sum_{i=1}^{d_g} \hat{ChD}_{P_d}^s(i) = 1$

$\hat{ChD}_{P_d}^s$  is in itself a *ChD*, since  $P_d$  can be regarded as a mixed strategy. Thus it sums up to 1 according to Lemma 3.

**Definition 9 (Payoff Matrix)** The payoff matrix of a game  $g$  of size  $d$ , denoted  $\pi_g$  is a matrix of size  $d$ . Let  $\pi_g(i)$  be the  $i$ :th entry of the matrix.

The payoff matrix is of the same dimensions as the *ChDs* and the same enumeration is used for the entries.

**Definition 10 (Payoff)** Let  $s$  be a strategy. Its expected payoff  $\pi_g(ChD_t^s)$  in a game  $g$  when meeting an opponent strategy  $t$  is defined by

$$\pi_g(ChD_t^s) = \sum_{i=1}^{d_g} \pi_g(i) \cdot ChD_t^s(i) \quad (1)$$

Since the payoff simply is a linear function of the *ChDs*, it is easy to determine what strategy is the most successful one in a certain environment of other strategies. It is also easy to take a subset of the entries in a *ChD* and make comparisons between them. An example of such a comparison is the one done in previous work where two of these derived properties, generosity and greediness, have been studied (Carlsson, Johansson, & Boman 1998). For instance, generous strategies have shown to get higher payoffs in noisy chicken games than the greedy ones.

## On optimal strategies and games

Based on the definitions above, we will now discuss some properties that strategies and games can be shown to have.

**Theorem 2 (Existence of Nash equilibria in meta-game)** Given a population of strategies  $P$  (able of playing an arbitrary strategic game  $g$ ), the meta-game of choosing a mix of strategies that play  $g$  has a Nash equilibrium.

Sketch of proof: (a full proof is given in (Johansson 1999)): We conclude that the choice of strategies given the *ChD* matrix, is itself a repeated game. Repeated games have Nash equilibria, and thus, the choice of strategies for playing an iterated game must have Nash equilibria. The result mean that, regardless of what game we play, we know that there

is a distribution of strategies that is optimal, even if the underlying game for which we choose the strategies, may lack such equilibria.

We may also prove the possibility of finding an optimal *game* for a certain mix of strategies, a property that could be useful in the design of agent environments. Imagine a situation where we have an agent with a certain behavioral pattern and we would like to design an environment that suits that particular agent. We could easily, by extracting the information about (i) the *ChDs* involving the strategy that describe the behavior of our agent, and (ii) the predicted  $P_d$ , derive a payoff matrix that will favour our own agent. Similar lines of reasoning may be used to custom-design environments that (in general) favour the ways in which we would like the agents to behave, i.e. to create a system of rewards and punishments in order to get the agents to follow some norms set up by the designers of the environment.

**Theorem 3** (*Existence of optimal games*) For all strategies  $s$  and a population distribution  $P_d$ ,

1. It is always possible to find a game  $g$  in which  $s$  is optimal.
2. If  $ChD_{P_d}^s$  is a corner of the convex hull of the set of *ChDs*, it is always possible to find a strictly optimal game for  $s$ .

Sketch of proof (full proof in (Johansson 1999)): 1. is easily proven, since all games that give all outcomes an equal payoff satisfy that property. 2. is shown by assigning the combination of outcomes that the corner represents a higher payoff than the other outcomes, thus rewarding  $s$ .

The implication of this result is that we may *not* tell generally that a certain strategy will be unsuitable for all games, i.e. there are generally speaking no “bad strategies”. On the contrary, we may always find games in which every strategy, given an arbitrary, but specifically chosen distribution of opponents, will be among the best.

**Theorem 4** (*NFL theorem for strategies*) Let  $G$  be the set of all possible games. Then, for arbitrary strategies  $s_1$  and  $s_2$  and population distributions  $P_d^1$  and  $P_d^2$ :

$$\sum_{g \in G} \pi_g(ChD_{P_d^1}^{s_1}) = \sum_{g \in G} \pi_g(ChD_{P_d^2}^{s_2}) \quad (2)$$

Sketch of proof (full proof in (Johansson 1999)): Similar to the NFL theorem proof in (Wolpert & Macready 1996). Given every pair of strategies, we show that they are equally good, when compared over all possible payoff matrices.

From an agent perspective, this theorem imply that without any knowledge about the context, an agent is unable to tell whether a certain action will be good or bad, *regardless* of what action it is.

## Conclusions and Future Work

We have introduced the concept of Characteristic Distributions and explained how they can be used to structure knowledge about how different strategies behave when they meet. This knowledge is useful for agents in order to make optimal choices (in a given context).

The *ChDs* combined with the distributions of choices of strategies is a powerful tool for modeling an agents' choice of strategies. It may be used in a variety of other settings such as:

- Various types of simulations of population dynamics.
- Different types of *situated games* where the  $P_d$  is based upon the local environment.
- Prevent invasions of “nasty” strategies in the population through asserting a class of such strategies a small probability in  $P_d$ . This may (hypothetically) prevent genetic drift toward weaker strategies, unable to defend themselves against the invaders.
- Use *ChDs* as a design tool in MAS environment design.

Some questions arise though, e.g. how do changes in the level of noise or the length of the game affect the theories? The theory works as long as these parameters remain the same; however, if they are changed, what the resulting *ChD*'s will look like is an open question.

In all we believe that Characteristic Distributions may be a powerful tool for agents in deciding how to behave, a decision that has to be made based upon the context they are in, according to the NFL theorem for strategies.

## References

- Abreu, D., and Rubinstein, A. 1988. The structure of nash equilibrium in repeated games with finite automata. *Econometrica* 56:1259–1282.
- Binmore, K., and Samuelson, L. 1992. Evolutionary stability in repeated games played by finite automata. *Journal of Economic Theory* 57:278–305.
- Carlsson, B.; Johansson, S.; and Boman, M. 1998. Generous and greedy strategies. In *Proceedings of Complex Systems 98*.
- Hofbauer, J., and Sigmund, K. 1998. *Evolutionary Games and Population Dynamics*. Cambridge: Cambridge University Press.
- Johansson, S. 1999. Game theory and agents. licentiate thesis. Dept. of Software Engineering and Computer Science, University of Karlskrona/Ronneby, Sweden.
- Rosenschein, J. S., and Zlotkin, G. 1994. *Rules of Encounter*. MIT Press.
- Rubinstein, A. 1986. Finite automata play the repeated prisoners dilemma. *Journal of Economic Theory* 39:83–96.
- Samuelson, L. 1997. *Equilibrium Selection and Evolutionary Games*. MIT Press.
- Weibull, J. 1996. *Evolutionary Game Theory*. MIT Press.
- Wolpert, D., and Macready, W. 1996. No free lunch theorems for search. Technical report, Santa Fe Institute. SFI-TR-95-02-010.
- Wooldridge, M. 1997. Agent-based software engineering. *IEE Proceedings on Software Engineering* 144:26–37.