

Handout 5: Background of a Game-Theoretic View to MAS

September 5, 2017

(Based on Shoham and Leyton-Brown 2011)

Introduction

1. In single-agent decision theory the key notion is that of an *optimal strategy*.
2. In noncooperative game theory, the basic modeling unit is the *individual* (including its beliefs, preferences, and possible actions) while in coalitional game theory, it is the *group*.
3. Agents are *self-interested*. It means that each agent has its own description of which states of the world it likes—and that it acts in an attempt to bring about these states of the world.
4. The dominant approach to modeling an agent's interests is *utility theory*. The theory also aims to understand how these preferences change when an agent faces uncertainty about which alternative it will receive.
5. The idea of (*expected*) utility can be grounded in a more basic concept of preferences (and *rationality*): *von Neumann and Morgenstern*. And the utility functions are aka von Neumann-Morgenstern utility functions.

Utility Theory

Let O denote a finite set of outcomes. For any pair $o_1, o_2 \in O$, let $o_1 \succcurlyeq o_2$ denote the proposition that the agent weakly prefers o_1 to o_2 . Let $o_1 \sim o_2$ denote the proposition that the agent is indifferent between o_1 and o_2 . Finally, by $o_1 \succ o_2$, denote the proposition that the agent strictly prefers o_1 to o_2 .

Now, a *lottery* is the random selection of one of a set of outcomes according to specified probabilities. Formally, a lottery is a probability distribution over outcomes written $[p_1: o_1, \dots, p_k: o_k]$, where each $o_i \in O$, each $p_i \geq 0$ and $\sum_{i=1}^k p_i = 1$. Let L denote the set of all lotteries. We will extend the \succcurlyeq relation to apply to the elements of L as well as to the elements of O , **effectively considering lotteries over outcomes to be outcomes themselves**.

Axiom 3.1.1 (Completeness). $\forall o_1, o_2 \ o_1 \succ o_2 \text{ or } o_2 \succ o_1 \text{ or } o_1 \sim o_2$.

Axiom 3.1.2 (Transitivity). *If $o_1 \succcurlyeq o_2$ and $o_2 \succcurlyeq o_3$, then $o_1 \succcurlyeq o_3$.*

Axiom 3.1.3 (Substitutability). *If $o_1 \sim o_2$, then for all sequences of one or more outcomes o_3, \dots, o_k and sets of probabilities p, p_3, \dots, p_k for which $p + \sum_{i=3}^k p_i = 1$,*

$$[p : o_1, p_3 : o_3, \dots, p_k : o_k] \sim [p : o_2, p_3 : o_3, \dots, p_k : o_k].$$

Substitutability states that if an agent is indifferent between two outcomes, it is also indifferent between two lotteries that differ only in which of these outcomes is offered.

Let $P_\ell(o_i)$ denote the probability that outcome o_i is selected by lottery ℓ .

Axiom 3.1.4 (Decomposability) If $\forall o_i \in O, P_{\ell_1}(o_i) = P_{\ell_2}(o_i)$, then $\ell_1 \sim \ell_2$.

Decomposability states that an agent is always indifferent between lotteries that induce the same probabilities over outcomes, no matter whether these probabilities are expressed through a single lottery or nested in a lottery over lotteries.

Axiom 3.1.5 (Monotonicity). If $o_1 \succ o_2$ and $p > q$, then $[p: o_1, 1 - p: o_2] \succ [q: o_1, 1 - q: o_2]$.

Lemma 3.1.6 If a preference relation \succcurlyeq satisfies the axioms completeness, transitivity, decomposability, and monotonicity, and if $o_1 \succ o_2$ and $o_2 \succ o_3$, then there exists some probability p such that for all $p' < p$, $o_2 \succ [p': o_1, (1 - p'): o_3]$, and for all $p'' > p$, $[p'': o_1, (1 - p''): o_3] \succ o_2$.

Axiom 3.1.7 (Continuity) If $o_1 \succ o_2$ and $o_2 \succ o_3$, then $\exists p \in [0, 1]$ such that $o_2 \sim [p: o_1, 1 - p: o_3]$.

If we accept Axioms 3.1.1, 3.1.2, 3.1.4, 3.1.5, and 3.1.7, it turns out that we have no choice but to accept the **existence of single-dimensional utility functions** whose expected values agents want to maximize. (And if we do *not* want to reach this conclusion, we must therefore give up at least one of the axioms.)

Theorem 3.1.8 (von Neumann and Morgenstern, 1944) If a preference relation \succcurlyeq satisfies the axioms completeness, transitivity, substitutability, decomposability, monotonicity, and continuity, then there exists a function $u: O \rightarrow [0, 1]$ with the properties that

1. $u(o_1) \geq u(o_2)$ iff $o_1 \succcurlyeq o_2$, and
2. $u([p_1: o_1, \dots, p_k: o_k]) = \sum_{i=1}^k p_i u(o_i)$

Normal Form Games

We have seen that under reasonable assumptions about preferences, agents will always have utility function whose expected values they want to maximize. *This suggests that acting optimally in an uncertain environment is conceptually straightforward*—at least as long as the outcomes and their probabilities are known to the agent and can be succinctly represented.

Agents simply need to choose the course of action that maximizes expected utility!

But in real life, that's often too good to be true.

Prisoner's Dilemma!

	Player 2 No Betray	Player 2 Betray
Player 1 No Betray	1,1	-4,3
Player 1 Betray	3,-4	-3,-3

Game theory gives answers to many of these questions. It tells us that any rational user, when presented with this scenario once, will adopt Betray—regardless of what the other user does. It tells us that allowing the users to communicate beforehand will not change the outcome. It tells us that for perfectly rational agents, the decision will remain the same even if they play multiple times; however, if the number of times that the agents will play is infinite, or even uncertain, we may see them adopt No Betray.

	Player 2 No Betray	Player 2 Betray
Player 1 No Betray	a,a	b,c
Player 1 Betray	c,b	d,d

*It works as long as $c > a > d > b$.

Definition of Normal Form (AKA Strategic Form) Games

A game written in this way amounts to a representation of every player’s utility for every state of the world, in the special case where states of the world depend only on the players’ combined actions.

Definition 3.2.1 (Normal-form game) A (finite, n-person) normal-form game is a tuple (N, A, u) , where:

- N is a finite set of n players, indexed by i ;
- $A = A_1 \times \dots \times A_n$, where A_i is a finite set of actions available to player i . Each vector $a = (a_1, \dots, a_n) \in A$ is called an action profile;
- $u = (u_1, \dots, u_n)$ where $u_i: A_i \rightarrow \mathbb{R}$ is a real-valued utility (or payoff) function for player i .

Note that we previously argued that utility functions should map from the set of *outcomes*, not the set of *actions*. Here we make the implicit assumption that $O = A$. A natural way to represent games is via an n -dimensional matrix.

Example 1. Prisoner’s dilemma

Example 2. Common-payoff games

Definition 3.2.2 (Common-payoff game) A common-payoff game is a game in which for all action profiles $a \in A_1 \times \dots \times A_n$ and any pair of agents i, j , it is the case that $u_i(a) = u_j(a)$.

Common-payoff games are also called *pure coordination games* or *team games*. In such games the agents have no conflicting interests; their sole challenge is to *coordinate on an action that is maximally beneficial to all*.

As an example, imagine two drivers driving towards each other in a country having no traffic rules, and who must independently decide whether to drive on the left or on the right. If the drivers choose the same side (left or right) they have some high utility, and otherwise they have a low utility.

	Driver 2 Left	Driver 2 Right
Driver 1 Left	1,1	0,0
Driver 1 Right	0,0	1,1

Example 3. Zero-sum games

At the other end of the spectrum from pure coordination games lie *zero-sum games (constant-sum games)*. Unlike common-payoff games, constant-sum games are meaningful primarily in the context of two-player games.

Definition 3.2.3 (Constant-sum game) A two-player normal-form game is constant sum if there exists a constant c such that for each strategy profile $a \in A_1 \times A_2$, it is the case that $u_1(a) + u_2(a) = c$.

A classical example of a zero-sum game is the game of *Matching Pennies*. In this game, each of the two players has a penny and independently chooses to display either heads or tails. The two players then compare their pennies. If they are the same then player 1 pockets both; otherwise player 2 pockets them.

	Player 2 Heads	Player 2 Tails
Player 1 Heads	1,-1	-1,1
Player 1 Tails	-1,1	1,-1

	Player 2 Paper	Player 2 Scissors	Player 2 Rock
Player 1 Paper	0,0	-1,1	1,-1
Player 1 Scissors	1,-1	0,0	-1,1
Player 1 Rock	-1,1	1,-1	0,0