Cooperative Game Theory

TDDE13 - Multi Agent Systems

Fredrik Präntare

LINKÖPING UNIVERSITY

Coalitional Games

Cooperative game theory offers another perspective on agents and multi-agent systems:

- in *non-cooperative games*, we have modelled the basic unit as the individual agent;
- in *cooperative games* (i.e., *coalitional games*), the main focus is instead on the *coalition*.

Agents may still have independent preferences, but we are more interested in groups of agents—how they interact, and what they can achieve.

Example 1 (Hospital)

- Suppose we aim to organize staff at a hospital by forming healthcare teams aimed at helping a number of patients in the best possible way.
- Ideally, we would like to pair each patient with the most suitable group (coalition) of doctors/nurses.
- We can model this scenario as a cooperative game.

Transferable Utility Assumption

Most work in cooperative game theory assumes *transferable utility* (Shapley, 1953): coalitions' *payoffs* can be distributed/transferred among members.

Coalitional Games with Transferable Utility

Definition 1. A *coalitional game with transferable utility* is a pair $\langle N, v\rangle$ where:

- $N = \{1, ..., |N|\}$ is a finite set of players; and
- $v : 2^N \mapsto \mathbb{R}$ is a function, called the *characteristic function*, that maps a value to each coalition $C \subseteq N$ that its members can distribute freely among themselves. $v(\emptyset) = 0$ is typically assumed.

The coalition $C = N$ is called the *grand coalition*.

A coalition's value is also called its *payoff* (or worth).

Example 2 (Voting Game)

A parliament is made up of four political parties: A, B, C, and D, which have 45, 25, 15, and 15 representatives, respectively.

They are to vote on whether to pass a $100 million spending bill, and how much of this amount should be controlled by each of the parties.

A majority vote, that is, a minimum of 51 votes, is required in order to pass any legislation, and if the bill does not pass, then every party gets zero to spend.

(example from the course's textbook by Shoham and Leyton-Brown)

Example 2 (Voting Game – cont'd)

In this game, there is a set of coalitions $\mathcal{W} \subseteq 2^N$ that are "winning" (i.e., that have a majority vote).

To each coalition $W \in \mathcal{W}$, we assign $v(W) = \$100M$, while for $L \notin \mathcal{W}$, we have $v(L) = 0$.

Coalitional Games with Transferable Utility

Major questions:

1. Which coalitions will/should form?
2. How should coalitions distribute payoff among members?
3. What are the *coalition structures* (partitioning of the agents) that maximizes the system's welfare (e.g., utility/performance)?

Definition 2. A *coalition structure* $S = \{C_1, ..., C_{|S|}\}$ over N is a set of coalitions with $C_i \cap C_j = \emptyset$ for all $i \neq j$; and $\bigcup_{i=1}^{|S|} C_i = N$.

Classes of Coalitional Games

Definition 3. A coalitional game $\langle N, v\rangle$ is *superadditive* if for all $A, B \subset N$ with $A \cap B = \emptyset$ the following holds:

$$v(A \cup B) \geq v(A) + v(B).$$

Thus, in these game types, coalitions never lose value from unifying, and the grand coalition has the highest payoff of all possible coalition structures.

The voting game is superadditive.

Classes of Coalitional Games

... and many more:

- *additive* $v(A \cup B) = v(A) + v(B)$;
- *convex* $v(A \cup B) \geq v(A) + v(B) - v(A \cap B)$;
- *simple* $v(A) \in \{0, 1\}$;
- *constant-sum* $v(A) + v(N \setminus B) = v(N)$;

for all $A, B \subseteq N$ with $A \cap B = \emptyset$.

Payoff Distribution

A central question in cooperative game theory is how to *distribute* the grand coalition's payoff among all agents.

So, why focus on the grand coalition?

First, many studied games are superadditive—in these games, the grand coalition is typically expected to form.

Second, often, agents have no choice but to form the grand coalition (e.g., since all participants in a public project are legally bounded to be included).

Payoff Distribution

Let the function $\psi(N, v) \mapsto \mathbb{R}^{|N|}$ denote a mapping from a coalitional game $\langle N, v \rangle$ to a vector of $|N|$ real values (called a *payoff vector*) where $\psi_i(N, v) \in \mathbb{R}$ denotes the i^{th} such value.

For shorthand, we denote such a payoff vector x, and we let x_i denote the i^{th} element and the *share* of the grand coalition's payoff that agent $i \in N$ receives.

Example 3 (Payoff Distribution)

Suppose we have a coalitional game $\langle N, v\rangle$ with $N = \{1, 2\}$, $v(\{1\}) = 2$, $v(\{2\}) = 3$ and $v(\{1, 2\}) = 4$.

If the grand coalition $N = \{1, 2\}$ was to form, how would you divide the payoff among its members?

- $\langle 4,\ 4\rangle$?
- $\langle 2,\ 2\rangle$?
- $\langle 4 \times \frac{2}{5},\ 4 \times \frac{3}{5}\rangle$?

Payoff Distribution

Relevant questions:

- What is a feasible payoff?
- Is there a fair way for a coalition to divide its payoff?
- Is there a systematic approach that makes it possible to always make a fair and reasonable payoff distribution?

Example 3 (Payoff Distribution)

Definition 4. Given a coalitional game (N, v), the *feasible payoff set* is defined as:

$$\Big\{ \langle x_1, ..., x_{|N|} \rangle \in \mathbb{R}^N : \sum_{i \in N} x_i \leq v(N) \Big\}.$$

So, the feasible payoff set contains all payoff vectors that do not distribute more than the grand coalition's worth.

The *pre-imputation set* Φ is defined as follows:

$$\Phi = \Big\{ \langle x_1, ..., x_{|N|} \rangle \in \mathbb{R}^N : \sum_{i \in N} x_i = v(N) \Big\}.$$

Payoff Distribution

- ~~What is a feasible payoff?~~
- What is a fair way for a coalition to divide its payoff?
- Is there a systematic approach that makes it possible to always make a fair and reasonable payoff distribution?

One approach to defining a *fair* division is to first identify a set of reasonable *axioms* that we agree on.

Payoff Distribution

Shapley's idea: a member should receive payoff proportional to her *marginal contribution*.

This is however not as straight-forward as it may seem...

Definition 5. Player i's *marginal contribution* $\Delta_i(C)$ to the coalition $C \subseteq N \setminus \{i\}$ is defined as:

$$\Delta_i(C) = v(C \cup \{i\}) - v(C).$$

Payoff Distribution

- Suppose $v(N) = 1$ but $v(S) = 0$ if $S \neq N$.
- Then $v(N) - v(N \setminus \{i\}) = 1$ for every $i \in N$: everybody's marginal contribution is 1; everybody is essential to generating any value!
- $\implies$ cannot pay everyone their marginal contribution...

(example by Matt Jackson)

Thus, we will have to use some kind of weighting system—but how should it be designed?

Shapley's axioms give us one answer...

Payoff Distribution (Axiom 1: Symmetry)

Definition 6. $i \in N$ and $j \in N$ are *interchangeable* if they always contribute the same amount to every coalition of the other agents. That is, for all $C \subset N$ with $i \notin C$ and $j \notin C$, $v(C \cup \{i\}) = v(C \cup \{j\})$.

Axiom 1 (Symmetry)**.** For any v, if $i \in N$ and $j \in N$ are interchangeable, then $\psi_i(N, v) = \psi_j(N, v)$.

Payoff Distribution (Axiom 2: Dummy player)

Definition 7. Player i is a *dummy player* if the amount that i contributes to any coalition is exactly the amount that i is able to achieve alone. That is, for all $C \subset N$ with $i \notin C$, $v(C \cup \{i\}) - v(C) = v(\{i\})$.

Axiom 2 (Dummy player)**.** For any v, if $i \in N$ is a dummy player, then $\psi_i(N, v) = v(\{i\})$.

Payoff Distribution (Axiom 3: Additivity)

Axiom 3 (Additivity). For any two v_1 and v_2, we have for any player $i \in N$, that:

$$\psi_i(N, v_1) + \psi_i(N, v_2) = \psi_i(N, v_1 + v_2),$$

where the game $(N, v_1 + v_2)$ is defined by:

$$(v_1 + v_2)(C) = v_1(C) + v_2(C)$$

for $C \subseteq N$.

The Shapley Value

Theorem 1. *Given a coalitional game (N, v), there is a* ***unique*** *(i.e., exactly one) pre-imputation $\phi \in \Phi$ that satisfies the axioms of* ***Symmetry****,* ***Dummy player*** *and* ***Additivity****.*

(recall that Φ is the set of pre-imputations)

This unique payoff division is called the *Shapley value.*

(proof provided in e.g., *A Course in Game Theory* by Osborne and Rubinstein)

The Shapley Value

Definition 8. Given a coalitional game (N, v), the *Shapley value* ϕ_i of player i is given by:

$$\phi_i = \frac{1}{|N|!} \sum_{C \subseteq N \setminus \{i\}} |S|!(|N| - |S| - 1)! \Big(v(C \cup \{i\}) - v(C) \Big).$$

The Shapley Value

$$\phi_i = \frac{1}{\boxed{|N|!}} \sum_{C \subseteq N \setminus \{i\}} |S|!(|N| - |S| - 1)!\Big(v(C \cup \{i\}) - v(C)\Big).$$

The number of ways to "build" the grand coalition
= number of permutations of *N*.

The Shapley Value

$$\phi_i = \frac{1}{\boxed{|N|!}} \sum_{C \subseteq N \setminus \{i\}} \boxed{|S|!}(|N| - |S| - 1)!\Big(v(C \cup \{i\}) - v(C)\Big).$$

The number of ways to "build society"
before i was added.

The number of ways to "build" the grand coalition
= number of permutations of N.

The Shapley Value

$$\phi_i = \frac{1}{\boxed{|N|!}} \sum_{C \subseteq N \setminus \{i\}} \boxed{|S|!}\boxed{(|N| - |S| - 1)!}\Big(v(C \cup \{i\}) - v(C)\Big).$$

The number of ways to "build society"
before i was added.

The number of ways to "build society"
after i was added.

The number of ways to "build" the grand coalition
= number of permutations of N.

The Shapley Value

The Shapley value captures the "average marginal contribution" of agent i (over the distinct sequences with which the grand coalition can be constructed).

Thus, if $\mathcal{P}$ denotes the set of all possible permutations (orderings) of N, and $C_i(P) \subset N$ is the set of players preceding i in $P \in \mathcal{P}$, we have an equivalent and more succinct way of defining the Shapley value:

$$\phi_i = \frac{1}{|\mathcal{P}|} \sum_{P \in \mathcal{P}} \Delta_i(C_i(P)).$$

Example 4 (Shapley Value)

Suppose we have a game $\langle N, v \rangle$, with players $N = \{1, 2, 3\}$ and value function $v(C) = |C|^2$, for which we want to calculate player 1's Shapley value.

Player 1's marginal contributions are:

- $\Delta_1(\emptyset) = v(\{1\}) - v(\emptyset) = 1 - 0 = 1$
- $\Delta_1(\{2\}) = \Delta_1(\{3\}) = v(\{1, 2\}) - v(\{2\}) = 4 - 1 = 3$
- $\Delta_1(\{2, 3\}) = v(\{1, 2, 3\}) - v(\{2, 3\}) = 9 - 4 = 5$

Example 4 (Shapley Value, cont'd)

$$
\begin{aligned}
\phi_1 &= \frac{1}{3!} \sum_{C \subseteq \{2,3\}} |C|!(3 - |C| - 1)!\Delta_1(C) \\
&= \frac{1}{3!} \sum_{C \subseteq \{2,3\}} |C|!(2 - |C|)!\Delta_1(C) \\
&= \frac{1}{3!}\Big(2!\Delta_1(\emptyset) + 1!\Delta_1(\{2\}) + 1!\Delta_1(\{3\}) + 2!\Delta_1(\{2,3\})\Big) \\
&= \frac{1}{3!}\Big(2 \times 1 + 1 \times 3 + 1 \times 3 + 2 \times 5\Big) = \frac{1}{6} \times 18 = 3
\end{aligned}
$$

Revisiting Payoff Distribution

- What is a feasible payoff?
  $\implies$ a pre-imputation.
- What is a fair way for a coalition to divide its payoff?
  $\implies$ the Shapley value.
- Is there a systematic approach that makes it possible to always make a fair and reasonable payoff distribution?
  $\implies$ the Shapley value.

The Core

Now that we have a fair way of dividing the grand coalition's payment among its members, we may ask: would the agents be *willing* to form the grand coalition given the way it will divide payments (e.g., according to the Shapley value)?

Example 5 (Voting Game Core)

(recall our voting game example) A parliament is made up of four political parties: A, B, C, and D, which have 45, 25, 15, and 15 representatives, respectively.

They are to vote on whether to pass a \$100 million spending bill, and how much of this amount should be controlled by each of the parties.

In this example, A and B have an incentive to defect from the grand coalition and form their own coalition. This way, they can both earn more (e.g., by splitting the \$100 million between themselves 75/25).

The Core

- What payment divisions would make the agents want to form the grand coalition?
- When do such payment divisions exist?

The Core

Definition 9. A payoff vector $\{x_1, ..., x_{|N|}\}$ is in the core of a coalitional game (N, v) if and only if:

$$\sum_{i \in C} x_i \geq v(C)$$

for all $C \subseteq N$—in other words, a payoff vector is in the core if it implies that no agents have an incentive to defect from the grand coalition.

Note that the core (in cooperative games) is analogous to the Nash equilibrium (in non-cooperative games).

The Core

This definition raises a few new questions:

- Is the core always non-empty?
- Is the core always unique?

The Core

The answer to both questions is no.

For example, in our voting-game example, the core is empty, since if the sum of payoffs to parties B, C and D is less than \$100 million, they have an incentive to deviate from the grand coalition. On the other hand, if B, C and D get the entire payoff, A gets \$0, and will have an incentive to deviate.

(recall A, B, C, and D have 45, 25, 15, and 15 representatives respectively)

The Core

If 80% majority vote is instead required, the core is non-unique, since the winning coalitions are $\{A, B, C\}$ and $\{A, B, D\}$. Any complete distribution of the \$100 million among the parties A and B is in the core.

(recall A, B, C, and D have 45, 25, 15, and 15 representatives respectively)

Example 6 (UN Security Council)

The UN security council has 15 members:

- 5 permanent: China, France, Russia, UK and US.
- 10 temporary.

The permanent members can veto solutions.

(example by Matt Jackson)

Example 6 (UN Security Council)

We can represent this as a cooperative game:

- China, France, Russia, UK, US are represented as players $\{1, 2, 3, 4, 5\} \subseteq N$.
- $v(C) = 1$ if $\{1, 2, 3, 4, 5\} \subset C$ and $|C| \geq 8$;
- $v(C) = 0$ otherwise.

Example 6 (UN Security Council)

Three-player game with similar structure:

- 1 permanent member with a veto and 2 temporary members.
- $v(C) = 1$ if $1 \in C$ and $|C| \geq 2$.
- $v(C) = 0$ otherwise.

Example 6 (UN Security Council)

$v(C) = 1$ if $1 \in C$ and $|C| \geq 2$, $v(C) = 0$ otherwise.

Core:

1. $x_1 + x_2 \geq 1$ (otherwise $\{x_1, x_2\}$ would deviate);
2. $x_1 + x_3 \geq 1$ (otherwise $\{x_1, x_3\}$ would deviate);
3. $x_1 + x_2 + x_3 = 1$ (all payoff should be divided); and
4. $x_i \geq 0$ (we cannot force anyone to join a coalition);

$\implies x_1 = 1$, i.e., $x = \langle 1, 0, 0 \rangle$.

Summary

- In *cooperative games*, the focus is on the coalition.
- A common assumption in these games is that of *transferable utility*.
- Payoff distribution can be done in many ways, but *the Shapley value* is the only value that satisfies the axioms *Symmetry*, *Dummy player* and *Additivity*.
- The Shapley value is based on marginal contributions.
- *The core* is a notion of stability which describes whether coalitions have an incentive to deviate.

www.liu.se