Multi-Agent Systems

Albert-Ludwigs-Universität Freiburg

Bernhard Nebel, Felix Lindner, and Thorsten Engesser

Summer Term 2017

Course outline

1. Introduction
2. Agent-Based Simulation
3. Agent Architectures
4. Beliefs, Desires, Intentions
5. Norms and Duties
6. Communication and Argumentation
7. Coordination and Decision Making
   - Distributed Constraint Satisfaction
   - Auctions and Markets
   - Cooperative Game Theory

Motivation

The famous prisoner's dilemma with the following payoff matrix:

	Silent	*Betray*
Silent	$-1, -1$	$-3, 0$
Betray	$0, -3$	$-2, -2$

In games like this one cooperation is prevented, because:

- Binding agreements are not possible
- Pay-off is given directly to individuals as the result of individual action

Cooperative Game Theory

- In many situations…
  - Contracts can form binding agreements
  - Pay-off is given to groups of agents rather than to individuals
- Hence, cooperation is both possible and rational.
- Cooperative game theory asks which contracts are meaningful solutions among self-interested agents.

Characterization (Shoham, Keyton-Brown, 2009, Ch. 12)

[Cooperative game theory is about] how self-interested agents can combine to form effective teams.

Real-world examples

- Political parties form coalitions to ensure majorities. Division of power (ministry posts).
- Companies cooperate to safe ressources.
- People buy expensive things together they could not buy themselves.
- Buildings are built by several people with different capabilities (craftsmen, electricians, architects, ...). Who should earn how much?

Key Questions

- Which coalition should/will form?
- How should the value be divided among the members?

Voting game

- Let A, B, C, D be four political parties in a parliament. They 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 it should be controlled by each of the parties.
- 51 votes are required to pass the bill, if the bill is not passed, every party gets zero to spend.
- Which coalition should form? How much does each of the parties in the coalition get to spend?

Cost sharing game

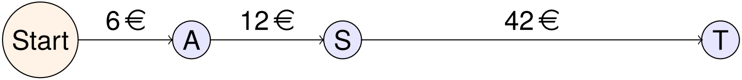

- Alvin drives home by taxi alone: 6 €
- Simon drives home by taxi alone: 12 €
- Theodore drives home by taxi alone: 42 €
- They could form coalitions to share a taxi.
  - How many coalitions will form?
  - If they all decide on sharing a taxi, the total price would be 42 €. Would you be fine if the total price of 42 € is divided by 3? Alternatives?

Cooperative Game Theory: Terminology

Cooperative Game (with transferable utility)

A cooperative game with transferable utility is a pair (N, v):

- N: Set of agents
- Any subset $S \subseteq N$ is called a coalition
- N is the grand coalition
- $v : 2^N \to \mathbb{R}$: characteristic function that assigns a value $v(S)$ to each $S \subseteq N$, $v(\emptyset) = 0$.

- Transferable value assumption:
  - Value of a coalition can be (arbitrarily) redistributed among the coalition's members
  - I.e., value is dispensed in some universal currency
  - Each coalition can be assigned a single value

Voting game

- Agents: $N = \{A, B, C, D\}$
- Coalitions: $\{A\}, \ldots, \{A, B, C, D\}$
- Characteristic function $v : 2^N \to \mathbb{R}$
  - $v(\{A\}) = v(\{B\}) = v(\{C\}) = v(\{D\}) = v(\{B, C\}) = v(\{B, D\}) = v(\{C, D\}) = 0$
  - $v(\{A, B\}) = v(\{A, C\}) = v(\{A, D\}) = v(\{B, C, D\}) = v(\{A, B, C, D\}) = 1$

Cost sharing game

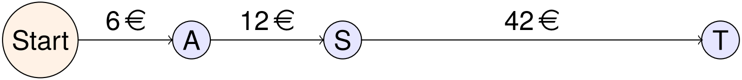

- Agents: $N = \{A, S, T\}$
- Coalitions: $\{A\}, \{S\}, \{T\}, \{A, S\}, \{A, T\}, \{S, T\}, \{A, S, T\}$
- Characteristic function $v : 2^N \to \mathbb{R}$
  - $v(\{A\}) = 6$
  - $v(\{S\}) = 12$
  - $v(\{T\}) = 42$
  - $v(\{A, S\}) = 12$
  - $v(\{A, T\}) = 42$
  - $v(\{S, T\}) = 42$
  - $v(\{A, S, T\}) = 42$

Coalitional game theory: Problems

- Again, which problems to solve?
  1. Which coalition will form?
  2. How should that coalition divide its value among its members?

The answer to 1) most of the times is the grand coalition N. But this also may depend on how 2) is answered, i.e., the grand coalition will only form if its in the interest of all its members.

Definitions: Superadditive games

Superadditive game

A game (N, v) is superadditive iff for all $S, T \subset N$, if $S \cap T = \emptyset$, then $v(S \cup T) \geq v(S) + v(T)$.

- Assumes that coalitions can work together without interfering with one another (adding someone to a team does not decrease its value).
- Consequently, the grand coalition is among the coalitions with highest value. $\Rightarrow$Unless stated otherwise, it is often assumed that the grand coalition will form.
- For cost-sharing games subadditivity defined in an analog way.

Division of Value

- Notation: $\Psi(S, v) = (\Psi_1(S, v), \ldots, \Psi_k(S, v))$ is a distribution of value to members $1, \ldots, k$ of S.

Feasible distribution

A distribution $\Psi(S, v)$ is feasible for a coalition S iff

$$\sum_{i \in S} \Psi_i(S, v) \leq v(S)$$

Efficient distribution

A distribution $\Psi(S, v)$ is efficient for a coalition S iff

$$\sum_{i \in S} \Psi_i(S, v) \geq v(S)$$

Fair Division

- Goal: Coalition is to divide its value 'fair'.
- Shapley's idea: Members should receive value proportional to their contributions.
- But:
  - Consider $v(N) = 1$ and $v(S) = 0$ for all $S \neq N$.
  - Thus, $v(N) - v(N \setminus \{i\}) = 1$ for every agent i: everybody's contribution is 1 (everybody is indeed likewise essential).
  - Clearly, one cannot pay 1 to everybody
  - Needed: Some way of weighing. How to design it?
  - Next: Axiomatic characterization of properties of a fair value division (due to Shapley).

Symmetry

Definition Interchangeability

Agents i and j are interchangeable relative to v iff they always contribute the same amount to every coalition of the other agents, i.e., for all S that contain neither i nor j, $v(S \cup \{i\}) = v(S \cup \{j\})$.

Axiom Symmetry

For any $S \subseteq N, v$, if i and j are interchangeable then $\Psi_i(S, v) = \Psi_j(S, v)$.

- Agents who contribute the same to every possible coalition should get the same.

Dummy Player and Null Player

Definition Dummy Player

Agent i is a dummy player iff the amount that i contributes to any coalition is $v(\{i\})$, i.e., for all $S \setminus \{i\}$, $v(S \cup \{i\}) = v(S) + v(\{i\})$.
If $v(\{i\}) = 0$, i is called a null player.

Axiom Dummy Player

For any $S \subseteq N, v$ if i is a dummy player then $\Psi_i(S, v) = v(\{i\})$.

Additivity

Axiom Additivity

For any two v_1, v_2, it holds that $\Psi_i(N, v_1 + v_2) = \Psi_i(N, v_1) + \Psi_i(N, v_2)$ for each i, where the game $(N, v_1 + v_2)$ is defined by $(v_1 + v_2)(S) = v_1(S) + v_2(S)$.

Shapley Value Theorem

Theorem

Given a coalitional game (N, v), there is a unique payoff division $\Psi(N, v)$ that divides the full payoff of the grand coalition and that satisfies Symmetry, Dummy Player, and Additivity: The Shapley Value.

Marginal value

Marginal value of agent i

The marginal value of an agent i to any coalition $S \subseteq N$ is defined by $\mu_i : 2^N \to \mathbb{R}$:

$$\mu_i(S) := \begin{cases} v(S \cup \{i\}) - v(S), & i \notin S \\ v(S) - v(S \setminus \{i\}), & i \in S \end{cases} .$$

Shapley Value

Definition Shapley Value

Given a cooperative game (N, v), the Shapley Value divides value according to:

$$\Psi_i(N,v) = \frac{1}{N!} \sum_{o \in \Pi(N)} \mu_i(C_i(o))$$

- $\Pi_n = \{(x_1, \ldots, x_n) | x_i \in N, \forall i, j[i \neq j \Rightarrow x_i \neq x_j]\}$
- $C_i(o)$: set containing only those agents that appear before agent i in o, e.g., $o = (3,1,2)$, then $C_3(o) = \emptyset$, $C_2(o) = \{1,3\}$

Voting game: Shapley Value

- Agents: $N = \{A, B, C, D\}$
- Coalitions: $\{A\}, \ldots, \{A, B, C, D\} \in 2^N$
- Characteristic function $v : 2^N \to \mathbb{R}$
  - $v(\{A\}) = v(\{B\}) = v(\{C\}) = v(\{D\}) = v(\{B, C\}) = v(\{B, D\}) = v(\{C, D\}) = 0$
  - $v(\{A, B\}) = v(\{A, C\}) = v(\{A, D\}) = v(\{B, C, D\}) = v(\{A, B, C\}) = v(\{A, B, C, D\}) = 1$
  - $(A, B, C) \to (0, 1, 0)$ because: $v(\{\} \cup A) - v(\{\}) = 0$, $v(\{A\} \cup \{B\}) - v(\{A\}) = 1$, $v(\{A, B\} \cup \{C\}) - v(\{A, B\}) = 0$
  - $(A, C, B) \to (0, 0, 1)$ because: $v(\{\} \cup A) - v(\{\}) = 0$, $v(\{A, C\} \cup \{B\}) - v(\{A, C\}) = 0$, $v(\{A\} \cup \{C\}) - v(\{A\}) = 1$
  - $(B, A, C) \to (1, 0, 0)$, $(B, C, A) \to (1, 0, 0)$
  - $(C, A, B) \to (1, 0, 0)$, $(C, B, A) \to (1, 0, 0)$
  - $\Psi(\{A, B, C\}, v) = (0.66, 0.16, 0.16)$
  - $\Psi(\{A, B, C, D\}, v) = (0.5, 0.16, 0.16, 0.16)$
  - What if $\{A, B\}$ form a coalition?

Taxi scenario: Shapley Value

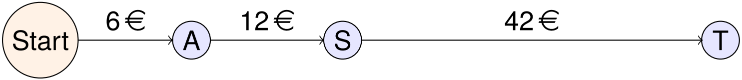

- characteristic function v
  - $v(\{A\}) = 6$
  - $v(\{S\}) = 12$
  - $v(\{T\}) = 42$
  - $v(\{A, S\}) = 12$
  - $v(\{A, T\}) = 42$
  - $v(\{S, T\}) = 42$
  - $v(\{A, S, T\}) = 42$

- Shapley value computation
  - $(A, S, T) \rightarrow (6, 6, 30)$
  - $(A, T, S) \rightarrow (6, 0, 36)$
  - $(S, A, T) \rightarrow (0, 12, 30)$
  - $(S, T, A) \rightarrow (0, 12, 30)$
  - $(T, A, S) \rightarrow (0, 0, 42)$
  - $(T, S, A) \rightarrow (0, 0, 42)$
  - $\Psi(N, v) = (2, 5, 35)$

The Core: Motivation

- Shapley value is concerned with a fair distribution among members of an already formed coalition.
- But it does not consider that some agents may prefer to form a smaller coalition (e.g., the parties A and B can form a smaller coalition and hope for more money each).

The Core: Definition

Core

The core of a cooperative game (N, v) is the set of feasible and efficient distributions of value Ψ, such that for all $S \subseteq N$ it's true that

$$\sum_{i \in S} \Psi_i \geq v(S)$$

- The sum of distributed value to the agents in any subcoalition S is at least as much as they could earn on their own. (Thus, $\leq$ for cost-sharing games).
- E.g., not fulfilled by the Shapley values for the voting game, because $\{A, B\}$ can gain more if they are not in the grand coalition.

Existence of the core

- Is the core always nonempty? No.
  - Voting game example: {A, B}, {A, C}, {A, D}, {B, C, D}, {A, B, C}, {A, B, D}, and {A, B, C, D} have more than 50 votes.
  - If the sum of distributed value given to B, C, D is less than 100 mio. €, then they have interest to form own coalition.
  - Otherwise, A has interest to form a coalition with, e.g., B.
  - Thus, the core is empty.

Uniqueness of the core I

- Is the core always unique? No.
  - Voting game example but 80 votes needed.
  - {A, B, C} and {A, B, D} can reach 80 votes.
  - Any distribution of the 100 mio. € among A and B belongs to the core, because all winnig coalitions need the support of these two parties.

Uniqueness of core II

- How come that in the 80-vote game, the core concept does not grant value to C, D in the grand coaltion $N = \{A,B,C,D\}$, $v(N) = 1$, although one of C, D is needed? Three cases:
  1. Both of C, D get something: $\Psi = (a,b,c,d)$, $c,d > 0$, $a+b+c+d = 1$:
     - Then A, B, C together only get $1-d < v(\{A,B,C\}) = 1$. Likewise, A, B, D together only get $1-c < v(\{A,B,D\}) = 1$. Hence, both these subcoalitions have an incentive to leave the grand coalition.$\Rightarrow$Not in core!
  2. Only C gets something: $\Psi = (a,b,c,0)$, $c > 0$, $a+b+c = 1$
     - Then A, B, D get $1-c < v(\{A,B,D\}) = 1$. Thus, this coaltion has an incentive to leave the grand coalition. $\Rightarrow$Not in core! (Similar argument if D gets s.th. instead of C.)
  3. None of C, D get something: $\Psi = (a,b,0,0), a+b = 1$.
     - Ψ is in the core: $a \geq v(\{A\}) = 0$, $b \geq v(\{B\}) = 0$, $0 \geq v(\{C\}) = 0$, $\ldots$, $1 \geq v(\{A,B\}) = 0$, $1 \geq v(\{A,B,C\}) = 1$, $\ldots\Rightarrow$No subcoalition of N can do better (stability), is not concerned with the individual fate of players.

Simple game, Veto agent

Simple game

A game (N, v) is a simple game iff for all $S \subseteq N, v(S) \in \{0, 1\}$

Veto agent

An agent i is a veto agent iff $v(N \setminus \{i\}) = 0$.

Theorem

In a simple game the core is empty iff there is no veto agent. If there are veto agents, the core consists of all value distributions in which the nonveto agents get 0.

- As we saw in the 80%-voting example: A and B are veto players $\Rightarrow$The core consists of all distributions s.th. A and B get everything and the other get nothing.

Convexity

Convex game

A game (N, v) is convex, iff the value of a coalition increases no slower when these coalitions grow in size, i.e.,
$v(S \cup \{i\}) - v(S) \leq v(T \cup \{i\}) - v(T)$ for all $S \subseteq T \subseteq N, i \in N \setminus T$.

Theorem

In every convex game, the Shapley value is in the core.

Note the relation between convexity and Shapley value via marginal contribution.

Theorem

Every convex game has a nonempty core.

$\Rightarrow$ Fair and stable distributions exist!

Example: Bankruptcy game instance

- Claimants: $N = \{A, B\}$
- Claims: $c_A = 80, c_B = 40$
- Estate: $E = 100$
- $v(C) = \max\{0, \sum_{i \in N \setminus C} c_i\}$
  - $v(\emptyset) = 0, v(\{A\}) = 60, v(\{B\}) = 20, v(\{A, B\}) = 100$

Properties

- This game is convex $\Rightarrow$the Shapley value is in the core.
- Shapley value: $\Psi = (\Psi_A, \Psi_B) = \frac{(60,40)+(80,20)}{2} = (70, 30)$
- In core indeed, because:
  - $\Psi_A = 70 \geq v(\{A\}) = 60$ ☺
  - $\Psi_B = 30 \geq v(\{B\}) = 20$ ☺
  - $\Psi_A + \Psi_B = 70 + 30 \geq v(\{A, B\}) = 100$ ☺

Intermediate Summary

- Cooperative game theory is concerned with what agents can achieve if they form coalitions, viz., binding agreements.
  - Values are given to coalitions first
  - Coalitions redistribute value to their members
- Solution concepts for cooperative games
  - Shapley value: fairness; always exists; unique
  - Core: stability; sometimes exists; not unique
  - For convex games, the Shapley value is in the Core
- Next
  - More compact representations for some types of games
  - Coalition structure formation

Computational aspects

Remember: Given a cooperative game (N, v), the Shapley Value divides value according to:

$$\Psi_i(N, v) = \frac{1}{N!} \sum_{o \in \Pi(N)} \mu_i(C_i(o))$$

- Imagine you wanted to compute the Shapley value of an agent i of a cooperative game (N, v)

  def shapleyValue(N, v, i):

  ...
- How many entries are in v?
- How many steps are necessary to compute Shapley value?

Feasible Game Types

- Some cooperative games can be treated more efficiently
  - Weighted graph games
  - Weighted voting games
- Centralized algorithm for coalition structure generation

Weighted graph game: Definition

Assumption

The value of a coalition is the sum of the pairwise synergies among agents.

Definition

Let (V, W) denote an undirected weighted graph, where V is the set of vertices and $W \in \mathbb{R}^{V \times V}$ is the set of edge weights; denote the weight of the edge between vertices i and j as $w_{\{i,j\}}$. This graph defines a weighted graph game, where the cooperative game is constructed as follows:

- $N = V$
- $v(S) = \sum_{\{i,j\} \subseteq S} w_{\{i,j\}}$

Weighted graph game: Example

Revenue Sharing game

Consider the problem of dividing the revenues from toll highways between the cities that the highways connect. The pair of cities connected by a highway get to share in the revenues only when they form an agreement on revenue splitting; otherwise, the tolls go to the state.

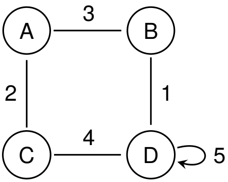

$$
\begin{aligned}
v(\{A,B,C\}) &= 3+2=5 \\
v(\{D\}) &= 5 \\
v(\{B,D\}) &= 1+5=6 \\
v(\{A,C\}) &= 2
\end{aligned}
$$

Weighted graph game: Shapley Value

1. Only N^2 many values to store (adjacency matrix).
2. Shapley-Value sh_i of agent i: $sh_i = w_{\{i,i\}} + \frac{1}{2}\sum_{i \neq j} w_{\{i,j\}}$

Each pair of agents plays a game, in which they are interchangeable. Thus, they get the same value (Symmetry).

Axiom Symmetry

For any $S \subseteq N, v$, if i and j are interchangeable then $\Psi_i(S, v) = \Psi_j(S, v)$.

Value adds up in “bigger” games due to Additivity.

Axiom Additivity

For any two v_1, v_2, it holds that $\Psi_i(N, v_1 + v_2) = \Psi_i(N, v_1) + \Psi_i(N, v_2)$ for each i, where the game $(N, v_1 + v_2)$ is defined by $(v_1 + v_2)(S) = v_1(S) + v_2(S)$.

Weighted graph game: Properties

Theorem

If all the weights are nonnegative then the game is convex.

Remember:

Theorem

Every convex game has a nonempty core.

Theorem

In every convex game, the Shapley value is in the core.

$\Rightarrow$A fair and stable value distribution exists and can be computed in polynomial time w.r.t. to number of agents.

Weighted graph game: Incompleteness

- For a sample game that cannot be represented as a weighted graph game remember the voting game from last lecture:
- Parties A, B, C, D have 45, 25, 15, and 15 representatives, respectively. 51 votes are required to pass the bill.
  - Agents: $N = \{A,B,C,D\}$
  - Coalitions: $\{A\},\ldots,\{A,B,C,D\} \in 2^N$
  - Characteristic function $v : 2^N \to \mathbb{R}$
    - $v(\{A\}) = v(\{B\}) = v(\{C\}) = v(\{D\}) = v(\{B,C\}) = v(\{B,D\}) = v(\{C,D\}) = 0$
    - $v(\{A,B\}) = v(\{A,C\}) = v(\{A,D\}) = v(\{B,C,D\}) = 1$
  - E.g., $v(\{B,C\}) + v(\{B,D\}) + v(\{C,D\}) \neq v(\{B,C,D\})$ ☹

Weighted voting game: Definition

Definition

A weighted voting game $(q; w_1, \ldots, w_n)$ consists of a set of agents $Ag = \{1, \ldots, n\}$ and a quota q. The cooperative game (N, v) is then given by:

- $N = Ag$
- $v(C) = \begin{cases} 1, & \sum_{i \in C} w_i \geq q \\ 0, & \mathit{else} \end{cases}$

Weighted voting game: Example

- Parties A, B, C, D have 45, 25, 15, and 15 representatives, respectively. 51 votes are required to pass the bill.
  - Weighted voting game: $(51; 45, 25, 15, 15)$

Weighted voting game: Properties

- Computing the Shapley value is NP-hard
- But checking if core is non-empty is easy

Remember:

Theorem

In a simple game the core is empty iff there is no veto agent. If there are veto agents, the core consists of all value distributions in which the nonveto agents get 0.

- Check if agent i is veto agent:
  1. Draw up $C = N \setminus \{i\}$
  2. Check that both hold:
     - $\sum_{j \in C} w_j < q$, i.e., no winner without i
     - $\sum_{j \in C \cup \{i\}} w_j \geq q$, i.e., winner with i

Next

- Some cooperative games can be treated more efficiently
  - Weighted graph games
  - Weighted voting games
- Centralized algorithm for coalition structure generation

Coalition Structure Formation

- Agents can use their capacity to compute Shapley values to try to optimize their local payoff.
- If, however, there is a central component that knows of all the agents, this component can attempt to maximize social welfare of the whole system.

Coalition Structure

A coalition structure is a partition of the overall set of agents N into mutually disjoint coalitions.

Example, with $N = \{1,2,3\}$:

- Seven possible coalitions:

$$\{1\},\{2\},\{3\},\{1,2\},\{2,3\},\{3,1\},\{1,2,3\}$$

- Five possible coalition structures:

$$\{\{1\},\{2\},\{3\}\},\{\{1\},\{2,3\}\},\{\{2\},\{1,3\}\},$$

$$\{\{3\},\{1,2\}\},\{\{1,2,3\}\}$$

Coalition Structure Formation

Given game $G = (N, v)$, the socially optimal coalition structure CS^* is defined as:

$$CS^* = \text{argmax}_{CS \in \text{ partitions of } N} V(CS)$$

where

$$V(CS) = \sum_{C \in CS} v(C)$$

Unfortunately, there are exponentially more coalition structures over the sets of agents N then there will be coalitions over N
$\Rightarrow$ Exhaustive search is infeasible!

Coalition Structure Graph

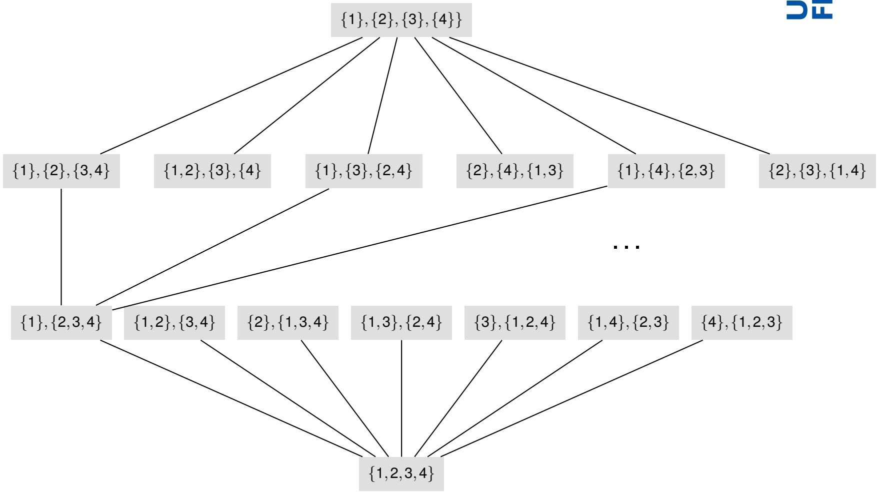

Coalition Structure Graph: Search

- Observation: At the first two levels every coalition is present.
- Let CS' be the best structure we find in these levels.
- Let CS^* be the best structure overall (as defined earlier).
- Let $C^* = \mathrm{argmax}_{C \subseteq N} v(C)$ the coalition with highest possible value.

Then:

- $V(CS^*) \leq |N| v(C^*) \leq |N| V(CS')$
- $\Rightarrow$ in worst case, $V(CS') = \frac{V(CS^*)}{|N|}$

Algorithm:

1. Search first two bottom levels, keep track of best one.
2. Continue with breadth-first search beginning with top level.

Summary

- Weighted graph games
  - Compact representation of games with additive values
  - Efficient compuation of Shapley values
- Weighted voting games
  - Compact representatuon of certain simple games
  - Efficient computation of a core value distribution
- Coalition Structure Formation
  - Centralized search-based algorithm to find a partition of agents into coalitions maximizing overall value.
  - Provable bounds of solution quality.

Course outline

Literature

- M. Wooldridge, An Introduction to MultiAgent Systems, 2nd Edition, John Wiley & Sons, 2009.
- Y. Shoham, K. Layton-Brown, Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations, Cambridge University Press, 2009.
- Game Theory Online, Youtube Channel, `https://www.youtube.com/user/gametheoryonline`