Multi-Agent Systems

CPSC 383: Explorations in Artificial Intelligence and Machine Learning
Fall 2025

Jonathan Hudson, Ph.D
Associate Professor (Teaching)
Department of Computer Science
University of Calgary

August 27, 2025

UNIVERSITY OF CALGARY

Outline

- Multi-agent systems
- Properties of Multi-agent systems
- Non-Cooperative Game Theory
- Making Collective Decisions

Multi-Agent Systems

Multi-Agent Systems

- Multi-agent systems (MAS) are used to describe several agents that interact with each other (positively, but also negatively).

- Positive interaction is usually known as **cooperation**,
- **collaboration** is used as a more elaborated word for interaction, and
- **competitive** settings describe usually systems where negative interaction takes place

Organization and Cooperation

- The main reason for human beings to form groups, like teams or larger societies, is that such groups can solve certain tasks better than a single individual.
- Here, better means that either no individual has all the necessary capabilities to solve the task or that the task is solved faster or with a (measurably) better result (synergy).
- Agents working together to achieve results cooperate. In order to achieve cooperation, the agents have to be (or become) organized

Cooperative Problem Solving

- Definition: Cooperative Problem Solving
- The following steps can be identified in each organization for cooperative problem solving:

a) Definition, creation and distribution of (sub)tasks

b) Working on (or solving of) the tasks by the assigned agents

c) Synthesis of the achieved results

- a) to c) are repeated until a satisfying solution of the initial problem is found (produced)

Multi-Agent Systems

- **Definition**: Agent in a MAS
- An agent Ag in a MAS is an agent,
- i.e. $Ag = (Sit, Act, Dat)$,
  - Sit -> situations the agent can be in
  - Act -> things an agent can do
    - Own actions
    - Actions to communicate or coordinate with others
  - Dat -> internal data

Multi-Agent Systems

- **Definition:** Multi-Agent System
- A multi-agent system **Mult** is a tuple
- i.e. $Mult = (Sit, Ag, M_{act}, a)$
  - Sit -> The situations the system can be in
  - Ag -> The agents in the system
  - M_{act} -> the combined actions the agent can all do
  - A function a that defines how agents are assigned actions

Dimensions

Dimensions

- System model
  - individual … team … society
- Granularity
  - fine grained … coarse grained
- Number of agents
  - small … medium … big
- Ability to adapt of agents and the whole MAS
  - fixed … programmable … able to learn …
- Control distribution
  - being controlled … dependent … independent

Dimensions

- Resources
  - limited … rich … unlimited
- Interaction scheme
  - simple … complex
- Solution strategy
  - synthesis … analytical
  - Build-up vs break-down
- Degree of cooperation between agents
  - cooperative and selfless … competitive and hostile

Decision Making

Decision Making

- One Decision Maker
  - Leader
  - Election challenges?
  - Distributed algorithms needed to elect?
- Multiple Decision Makers
  - Consensus
  - Voting
  - Stakeholders
  - Game theory?
  - Distributed algorithms for consensus?

Distributed Algorithms

- The network nodes communicate among themselves in order to decide which of them will get into the "leader" state.
- For that, they need some method in order to break the symmetry among them.
  - Ex. if each node has unique and comparable identities, then the nodes can compare their identities, and decide that the node with the highest identity is the leader.
  - With identical nodes things become much more complicated
- Algorithm goals generally
  - Byes sent
  - Time to complete
- Most common minimal problem is a ring (in a graph), where a token (leader) has been lost -> how to create a new token (and only one)

One decision maker

One decision maker

- multiple actors, it contains only one decision maker.
- **benevolent agent assumption**: agents will simply do what they are told

- **Multi-effector planning:** manages each effector while handling positive and negative interactions among the effectors.

- **Multi-body planning**: effectors are physically decoupled into detached units

Multiple decision makers

Multiple decision makers

- multiple actors, multiple decision makers.
- each have **preferences** and **choose and execute** their own plan

- There are two possibilities:
  1. Common goal for the actors, coordination problem (same direction)
  2. Own personal preference and can be diametrically opposed

Communication

Basic Communication Structures

- Today's computers and computer networks know two general principles to communicate with each other, defined by the physical world of our hardware:

1. communication using shared memory
2. communication using message passing

- In principle, each of these principles can simulate the other principle, perhaps with higher costs.

Communication using Shared Memory

All agents of the MAS have access to a common area of memory (the shared memory). This access allows both for reading and writing. If an agent writes information to this area (often also called blackboard) every other agent can read it. This realizes a 1:m communication and each agent has to decide on its own what to do with such an information:

- n ignore it
- n use it
- n delete it (usually forbidden)

Communication using Shared Memory

Problems/Disadvantages:

- How do we establish a control?
- Agents have to react to all new entries in the blackboard, in order to find the ones relevant for them.
- Only usable for problems that can be divided into subproblems easily and where the solutions of the subproblems can be easily put together to a solution for the initial problem

Communication by Message Passing

- In contrast to blackboards with the fixed communication partner blackboard, communication by sending messages is always an action between two agents (1:1 communication), the sender and the receiver. Both agents have to perform (nearly) simultaneously communication actions:
- The sender sends the message composed by it to the communication channel and the receiver gets the message from the channel.
- Usually the message is composed in a special communication language that has to be understood by both agents (internal data structures usually are not used, in contrast to shared memory).
- The whole communication to a subject typically consists of several messages where the agents take their turns as sender and receiver. This is called a dialog.
- Today's hardware also allows to send a message to all other agents (broadcast).

Communication by Message Passing

Problems/Disadvantages:

- New agents have to register themselves with all other agents
- How to define a communication language?
- What to communicate?
- Agents (processors) have to synchronize
  - potential loss of valuable time

Planning

Multiagent planning

Multiagent planning

- Issue of concurrency: plans of each agent may be executed simultaneously
- Agents must take into account the way in which their own actions interact with the actions of other agents.

Approaches

1. Interleaved execution
2. True concurrency
3. Perfect synchronization

Interleaved execution approach

Interleaved execution approach:

- certain the order of actions in the respective plans will be preserved
- assume that actions are atomic
- must be correct with respect to all possible interleavings of the plans
- does not model the case where two actions actually happen at the same time.
- the number of interleaved sequences is exponential with the number of agents and actions rising

Concurrency/Synchronization Approaches

True concurrency approach:

- do not attempt to create a full serialized ordering of the actions,
- partially ordered

Perfect synchronization approach

- Global clock,
- Same time, same duration, actions are simultaneous.
- simple semantics

Planning

Planning with multiple agents: Cooperation and coordination

- Adopt **convention** before engaging in joint activity
- **Convention**: constraint on the selection of joint plans
- When conventions are widespread, they are called **social laws**
- Agents can use **communication** to achieve common knowledge of a feasible joint plan.
- **Plan recognition:** single action (or short sequence of actions) by one agent is enough for the other to **determine a joint plan** unambiguously

Next…game theory and social welfare

Jonathan Hudson, Ph.D.
jwhudson@ucalgary.ca
https://cspages.ucalgary.ca/~jwhudson/

UNIVERSITY OF CALGARY