20IS603 Architecture of Intelligent Systems

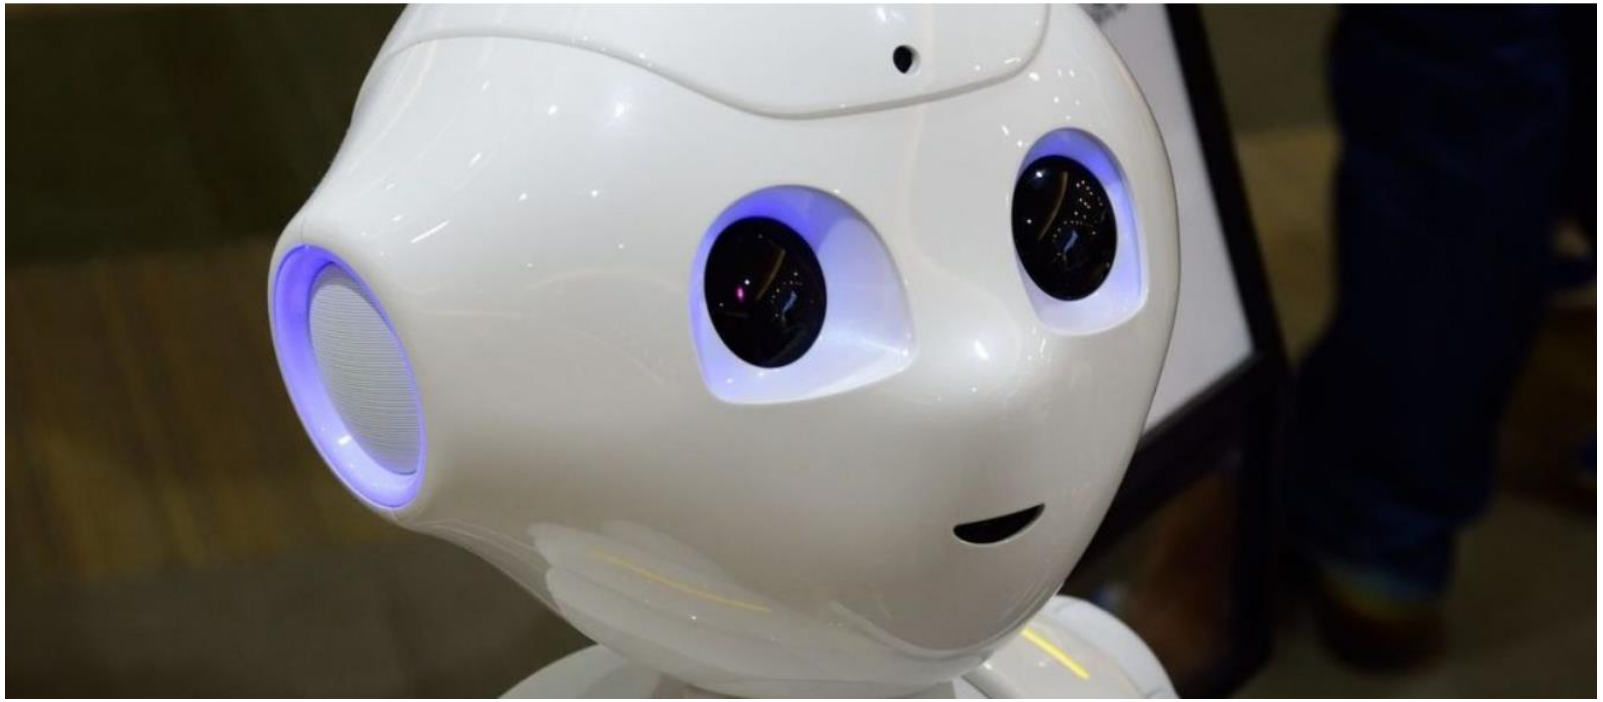

Intelligent Agents

Lecture #10

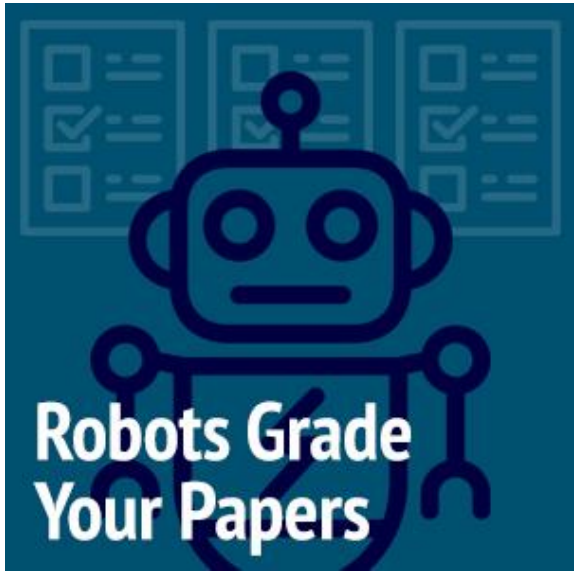
Robots Grade
Your Papers

Frequently Bought Together

Price For All Three: $258.02

Add all three to Cart

- **This item:** The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition (Springer Series in Statistics) by Trevor Hastie
- Pattern Recognition and Machine Learning (Information Science and Statistics) by Christopher M. Bishop
- Pattern Classification (2nd Edition) by Richard O. Duda

Customers Who Bought This Item Also Bought

All of Statistics: A Concise Course in Statist... by Larry Wasserman
(8) $60.00

Pattern Classification (2nd Edition) by Richard O. Duda
(27) $117.25

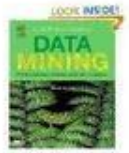

Data Mining: Practical Machine Learning Tools an... by Ian H. Witten
(29) $41.55

Bayesian Data Analysis, Second Edition (Texts in... by Andrew Gelman
(10) $56.20

Data Analysis Using Regression and Multilevel /... by Andrew Gelman
(13) $39.59

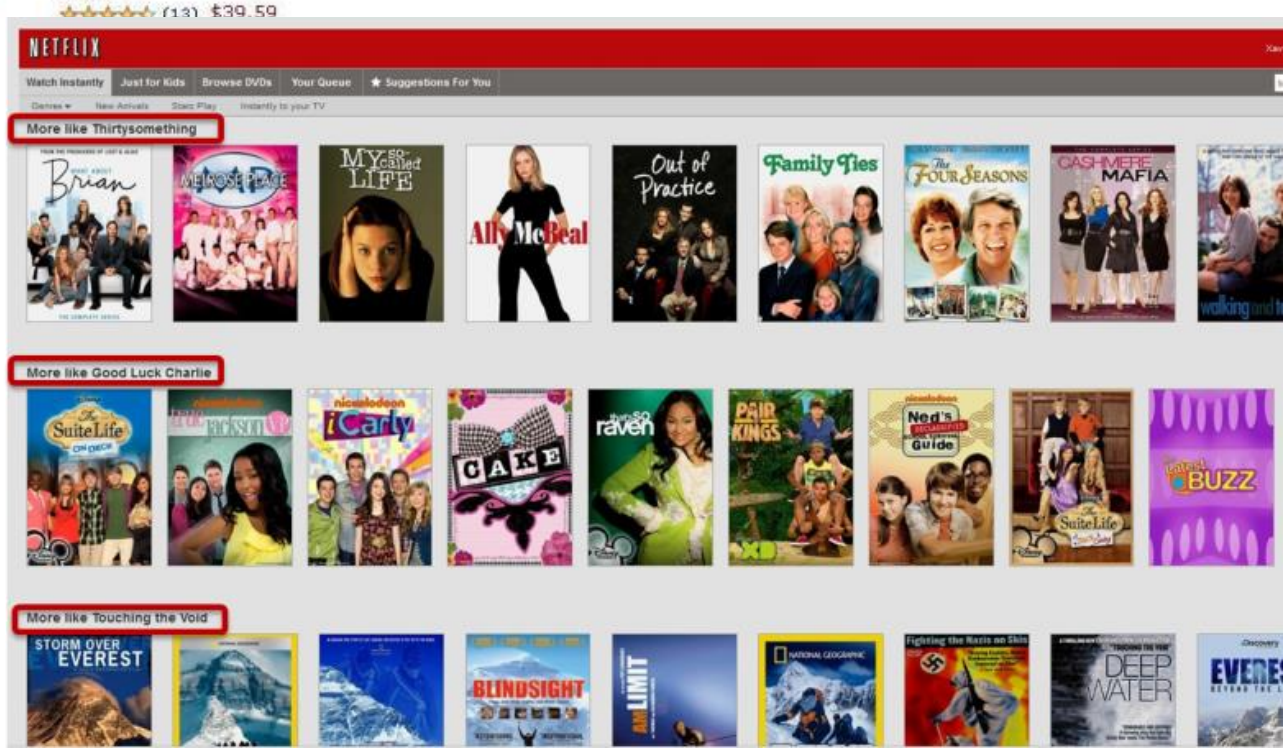

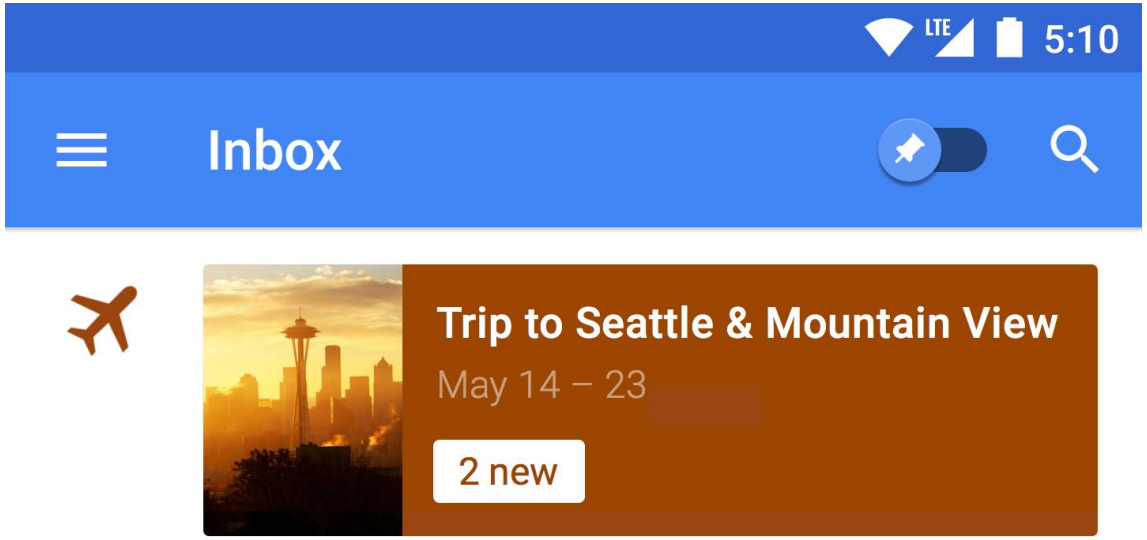
5:10
Inbox
Trip to Seattle & Mountain View
May 14 – 23
2 new

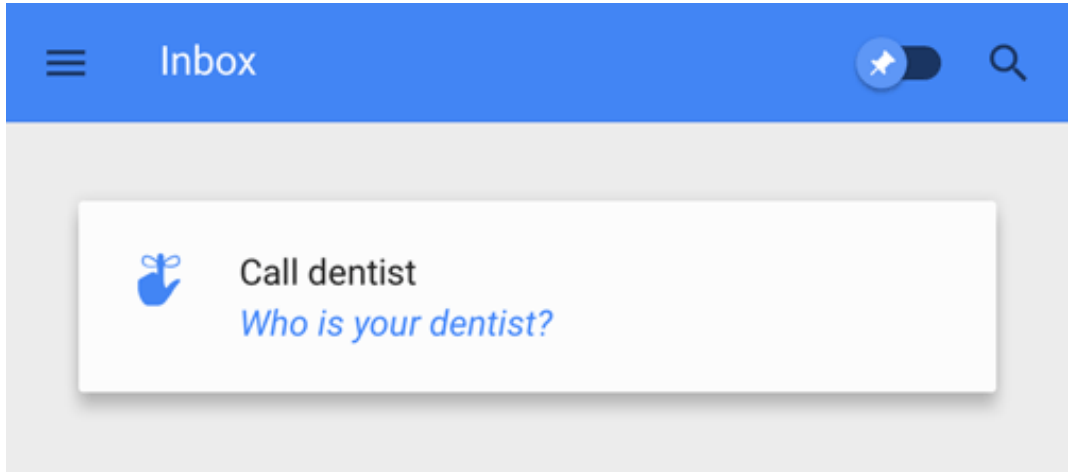
Inbox
Call dentist
Who is your dentist?

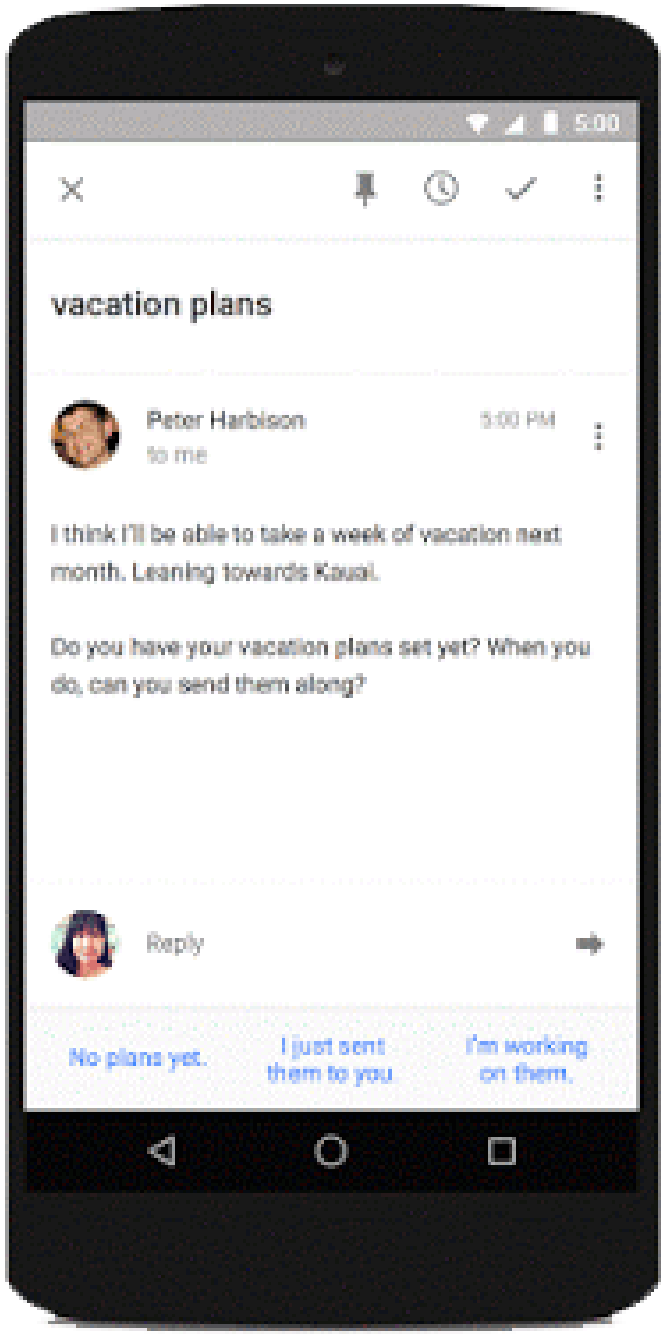
vacation plans
Peter Harbison
5:00 PM
to me
I think I'll be able to take a week of vacation next month. Leaning towards Kauai.
Do you have your vacation plans set yet? When you do, can you send them along?
Reply
No plans yet.
I just sent them to you
I'm working on them.

Intelligent Agents

- People are becoming less and less able to act upon the escalating quantities of information presented to them.
- Build intelligent agents—assistants that take care of specific tasks for us.
- Personalizes itself to your individual requirements by learning your habits and preferences
- Perform an action as instructed - not servants but independent entities in charge of their own actions.
- Agents or Intelligent agents are autonomous entities that manage their own activities

Intelligent Agents (2)

- An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators.
- A human agent has eyes, ears, and other organs for sensors and hands, legs, vocal tract, and so on for actuators.
- A robotic agent might have cameras and infrared range finders for sensors and various motors for actuators.
- The environment could be everything—the entire universe!

Intelligent Agents (3)

- A software agent receives file contents, network packets, and human input (keyboard / mouse / touchscreen /voice) as sensory inputs and acts on the environment by writing files, sending network packets, and displaying information or generating sounds.
- Percept refer to the content an agent's sensors are perceiving.

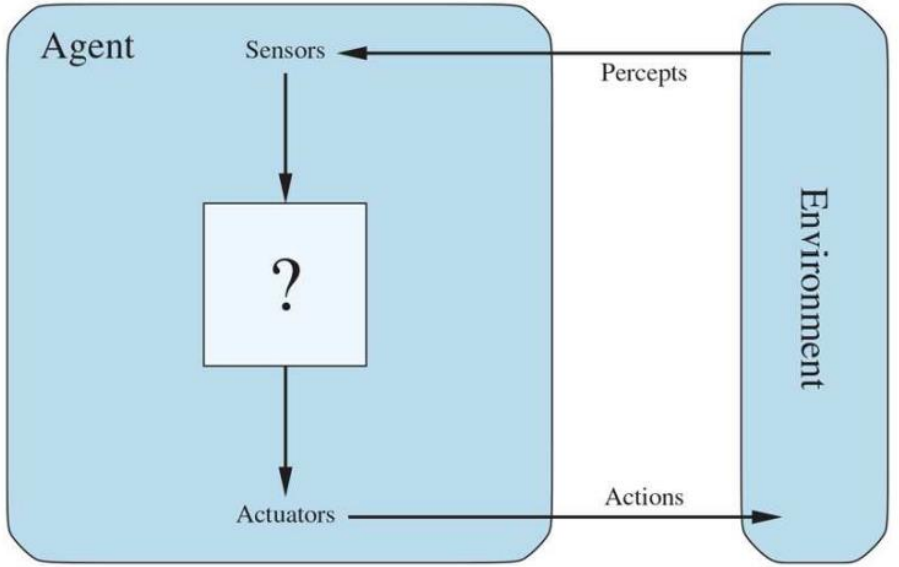

Examples of Intelligent Agents

Agent	Percepts	Actions	Goals	Environment
Medical diagnostic system	Symptoms, test results, patient's answers	Questions, test requests, treatments, referrals	Healthy patients, minimise costs	Patient, hospital, staff
Part-picking robot	Pixels of varying intensity and colour	Pick up parts and sort them into bins	Place parts into correct bins	Conveyor belt with parts, bins
Interactive English tutor	Typed words	Display exercises, suggestions, corrections	Maximise student's exam results	Set of students, exam papers

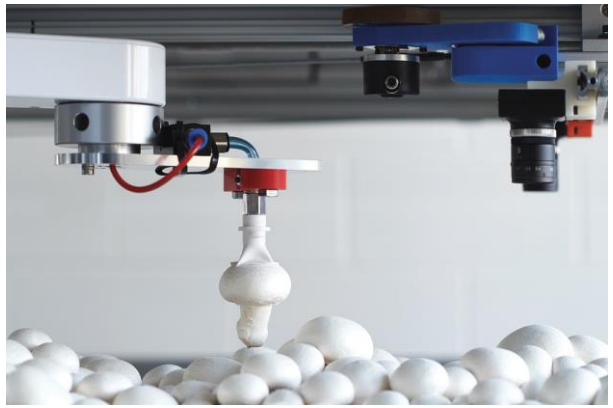

Properties of Agents

- Intelligence
  - Additional domain knowledge that enables them to carry out their tasks even when the parameters of the task change or when unexpected situations arise.
  - Able to learn, from their own performance, from other agents, from the user, or from the environment
- Autonomy
  - Ability to act and make decisions independently of the programmer or user of the agent
- Persistence
  - Functions continuously within its environment
- Social capability
  - Interact with other entities such as agents, humans, and systems

Properties of Agents (2)

- Goal-directed
  - Perform a specific task regardless of their environment.
- Reactivity
  - Perform actions timely corresponding to updated environment.
- Mobility
  - Can travel to remote computers to carry out its designated task
- Versatility
  - Able to carry out many different tasks
- Benevolence
  - Performs action for the welfare of other
- Rationality
  - Decision making is rational

Agent Classification

- Reactive Agents
  - React to events in their environment according to predetermined rules
  - Automatic email filter - reject so-called “junk mail” or “spam”
  - Does not tend to perform well when its environment changes

- Goal-based Agents
  - Acts to try to achieve a goal - done by using search or planning

Agent Classification

- Utility-based Agents
  - Similar to goal-based agent, but in addition to achieve a set of goals, the utility-based agent is also trying to maximize utility value.
  - How well it has satisfied the goal or how efficiently
  - Mars Lander on the surface of mars with an obstacle in its way

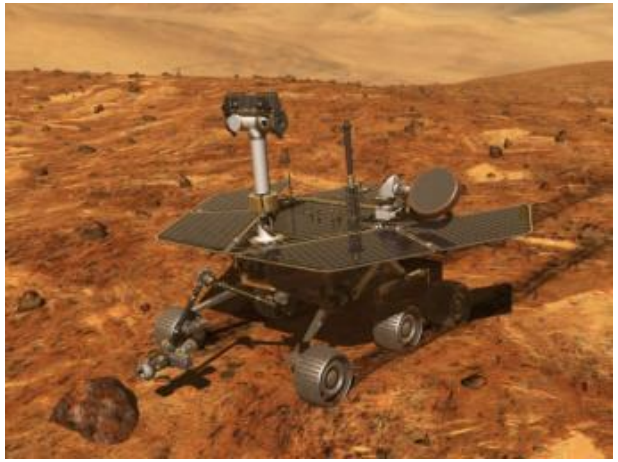

- Interface Agents
  - Autonomous agents, capable of learning in order to carry out tasks
  - Collaborate with the user - take instructions from users and can also learn from feedback from users
  - Behaves as personal assistant

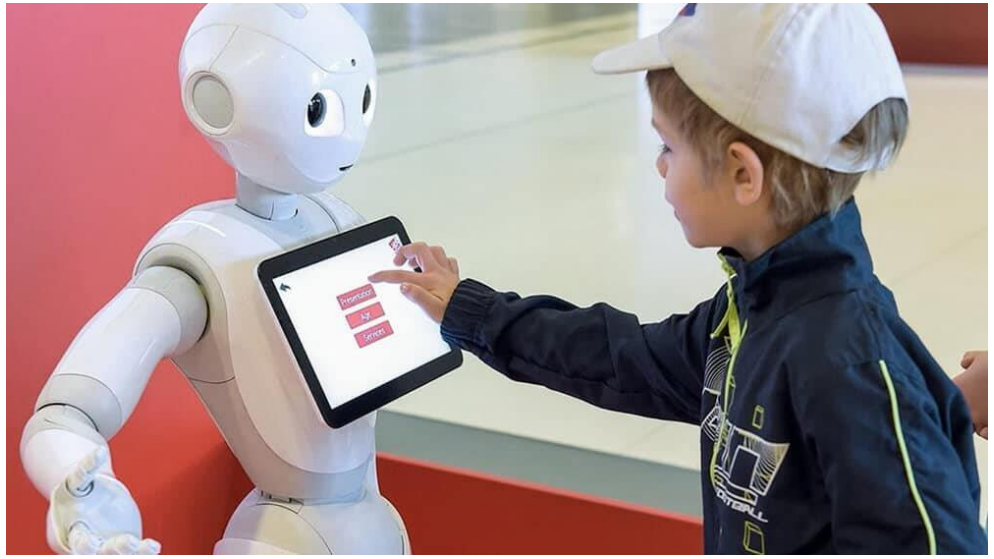

Agent Classification

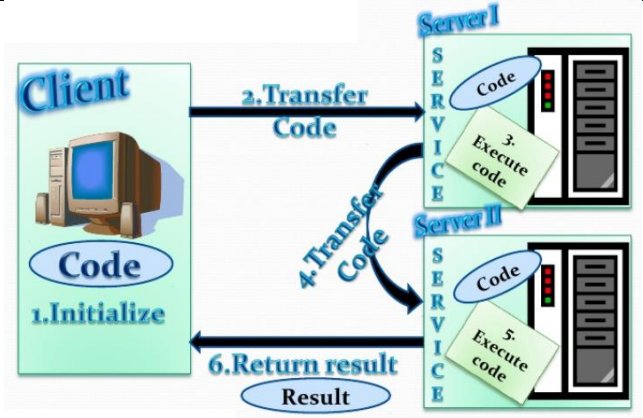

- Mobile Agents
  - Travel from one computer to another, gathering information and performing actions as needed on the basis of that information
  - Telescript developed by General Magic

- Information Agents
  - Used to help a user find, filter, and classify information from the vast array of sources available on the Internet
  - Can be collaborative or can work independently of other agents.

Agent Architectures

Agent = Architecture + Agent program

- Architecture is machinery that an AI agent executes on - with sensors and actuators
- Agent program is an implementation of an agent function
- Agent function is used to map a percept to an action
- An agent program executes on the physical architecture to produce function
- Agent architecture is the brain of the agent as it determines how the knowledge/information is represented in the agent.

Logic-Based Architectures

- Uses symbolic representation for reasoning
- An agent's decision making process is modeled through a set of deduction rules
- Agent's performance is based on its programmed deduction rules and its current database (representing the environment).

- A vacuum-cleaner world with just two locations. Each location can be clean or dirty, and the agent can move left or right and can clean the square that it occupies

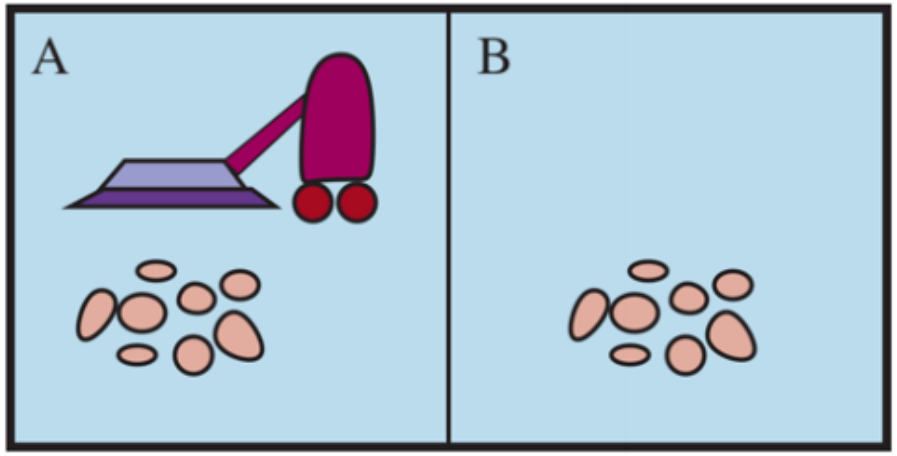

Emergent Behavior Architectures

- Based on *reactive agents*
- Designed for implementing physical robots, which does not involve any centralized intelligence or control mechanism
- Example – Brooks' subsumption architecture
- Behavior modules link actions to observed situations without any reasoning at all.
- The behaviors are arranged into a subsumption hierarchy, where low-level behavior has precedence over higher-level goal-oriented behaviors.
- Rather than having a centralized representation, the subsumption architecture relies on lower-level modules that combine together - these combined modules emerges intelligent behavior

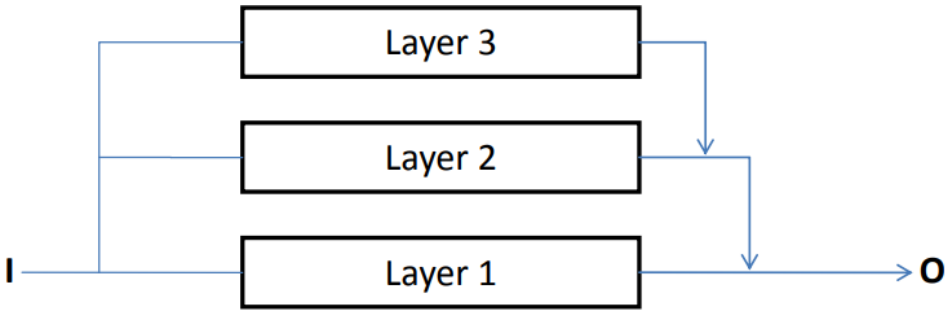

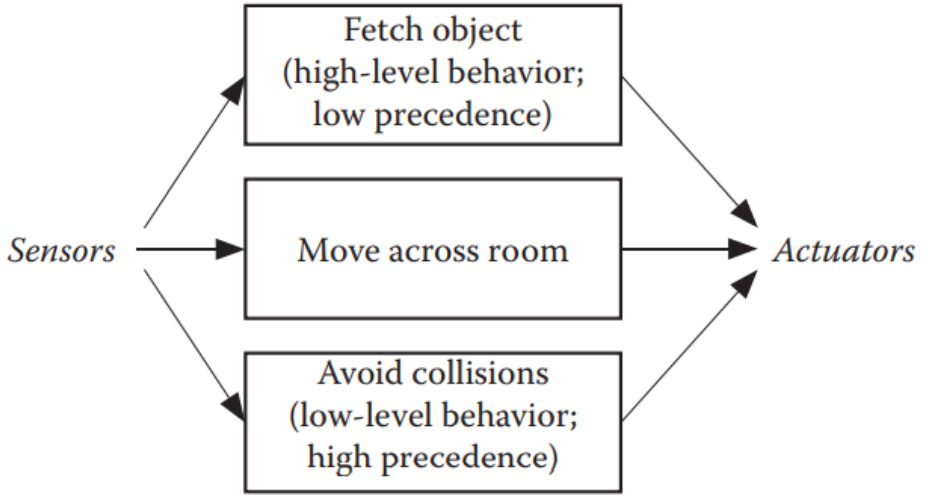

Belief-Desire-Intention (BDI) Architecture

- Modeled with a *knowledge level*
- Represent a symbolic model and make decisions via logical reasoning
- Knowledge of the environment is held as "beliefs" - similar to the set of facts contained in a rule-based production system
- The overall goals are "desires" - the agent would like to reach.
- Selected options that the system commits itself toward achieving its desires are "intentions".
- The process of determining what to do, i.e., the desires or goals, is *deliberation*.
- The process of determining how to do it, i.e., the plan or intentions, is *means–ends analysis*.

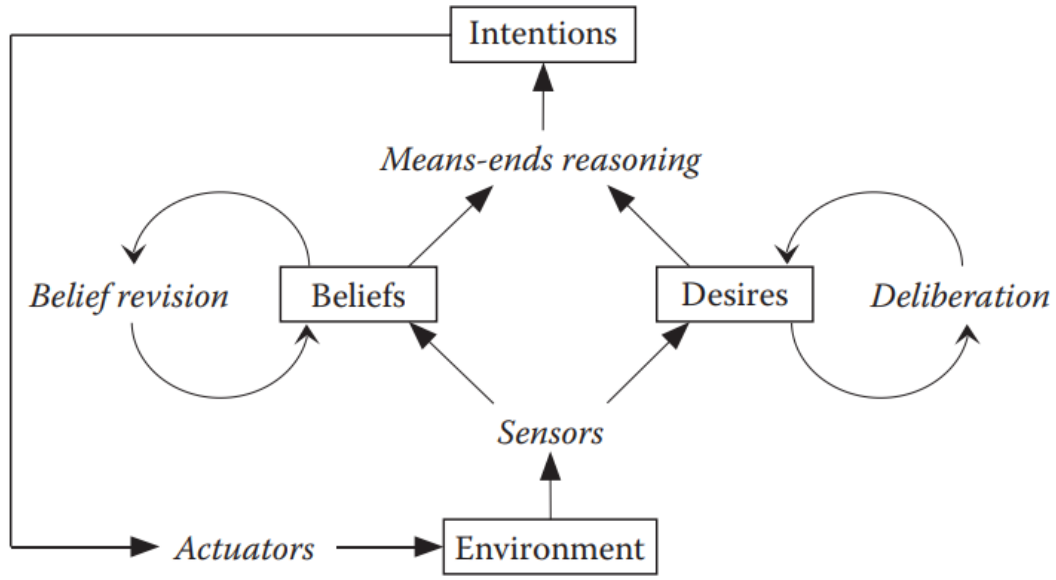

Belief-Desire-Intention (BDI) Architecture (2)

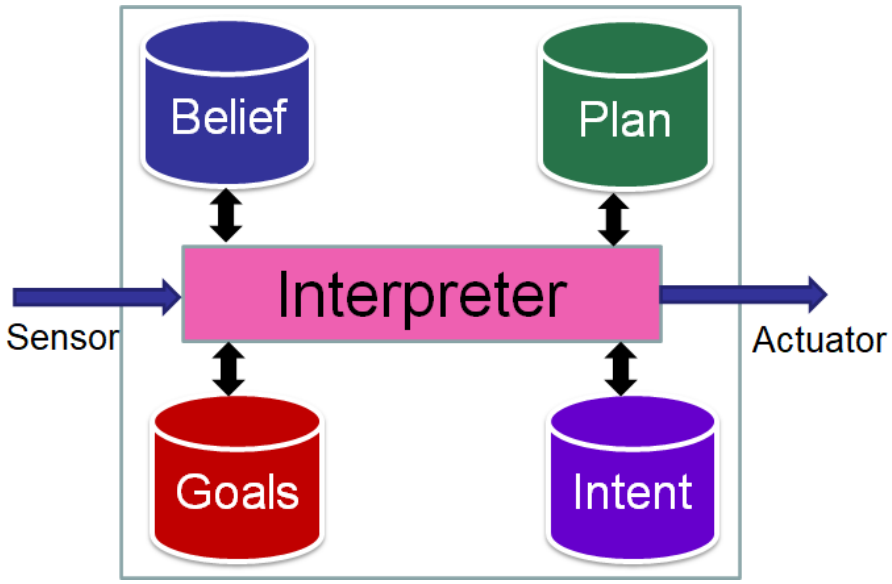

- Practical Reasoning System (PRS) is one of the well known BDI architectures
- PRS is a framework for building real-time reasoning systems that can perform complex tasks in dynamic environments
- Four key data structures: beliefs, desires, intentions and plans, and an interpreter.
- Plans specify some courses of action for the agent in order to achieve its intentions.
- The agent interpreter is responsible for updating beliefs from observations made from the environment, generating new desires (tasks) on the basis of new beliefs, and selecting from the subset of currently active desires to act as intentions.

Layered Architectures

- Combines both the advantages of reactive and logic-based architecture – Hybrid architecture
- Subsystems are decomposed into a layer of hierarchical structure to deal with different behaviors.
- Two types of interaction that flow between the layer - horizontal and vertical
- In the horizontal layer architecture, each layer receives inputs and contributes to the actions and outputs of the agent.
- In a vertical layered architecture, input is passed to one layer, which then passes information on to a further layer

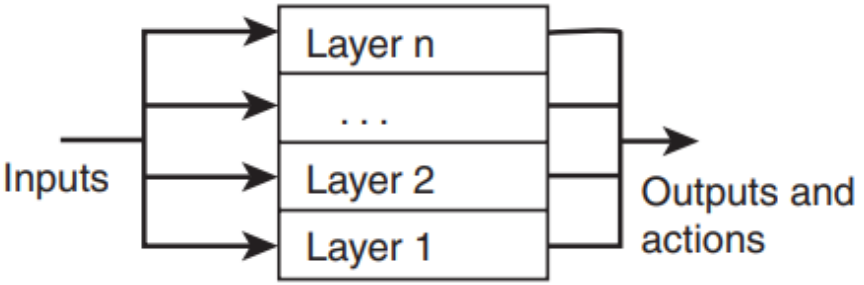

Horizontal layer architecture

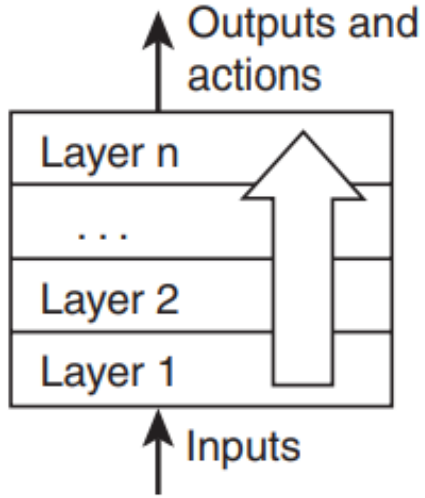

Vertical layer architecture

Layered Architectures

Horizontal layer architecture

- Touring Machines is an example of a horizontal architecture
- Consists of three activity-producing layers:
  - *Reactive layer*: uses situation rules to react to changes in the agent's environment.
  - *Planning layer*: uses a library of plans (called schemas) to determine the behavior of the agent, in order to achieve particular goals.
  - *Modeling layer*: contains a model of the agent and any other agents in the world

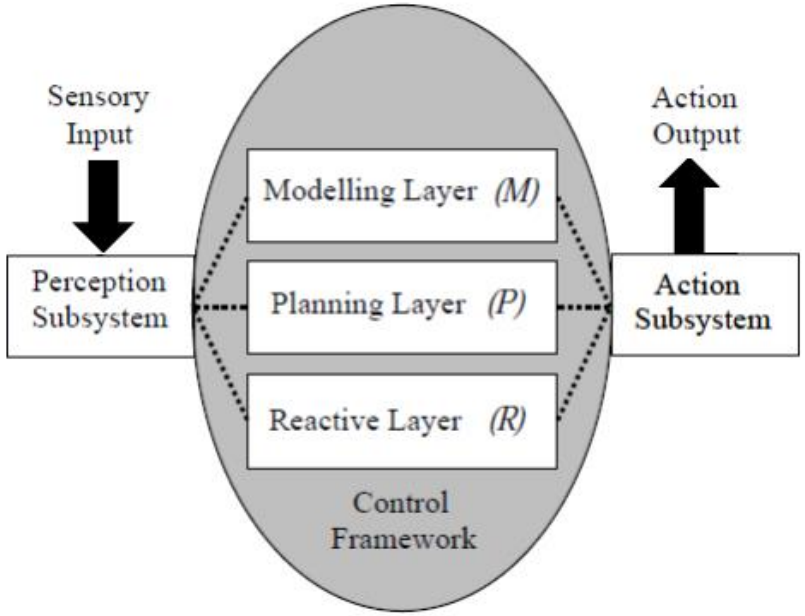

Multiagent Systems

Multiagent Systems

- Multiagent systems (MAS) are often required to solve problems in dynamic and unpredictable environments.
- MAS gives the possibility of mimicking human organizations and societies that contain a collection of individuals with their own personalities.
- Also called as agent-oriented or agent-based systems
- Each agent in a multiagent system can learn independently of the other agents and can also learn from the other agents.
- A multiagent system can be defined as a system in which several interacting intelligent agents pursue a set of individually held goals or perform a set of individual tasks.

Need for Multiagent Systems

- Inherently complex problems
  - Too large to be solved by a single hardware or software system.
  - Well-designed agents will ensure that every circumstance is handled in an appropriate manner, even though it may not have been explicitly anticipated.

- Inherently distributed problems
  - Data and information may exist in different physical locations, or at different times, or may be clustered into groups requiring different processing methods or semantics.
  - Problems require a distributed solution, provided by agents that run in parallel on distributed processors or run concurrently as independent processes on a single processor.

Building a Multiagent Systems

- Multiagent system - dependent on interactions between intelligent agents - when, how, and with whom should agents interact.
- To achieve coherency, multiagent systems can be designed bottom-up or top-down

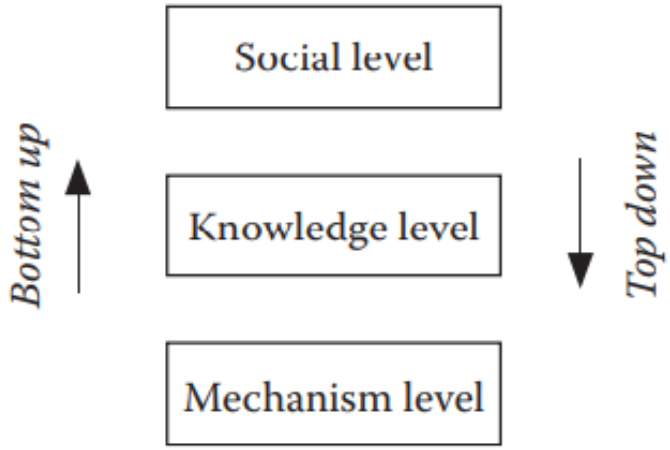

Building a Multiagent Systems (2)

- There may be large numbers of different agents capable of providing different services, possibly located on separate computers.
- Directory Facilitator (DF) agent helps to identify which services an agent can offer and where it is located.
- DF agent is analogous to a Yellow Pages© service.
- DF does not offer services directly related to the problem, but facilitates distributed problem-solving in a multiagent system.
- An agent must actively register with the DF agent in order for its name, network address, and services to be listed.
- The DF agent may act as a broker between the soliciting agent and the recommended service-supply agent, or negotiation may take place directly between them.

Agent Interaction Models

- Contract Net Framework
  - A manager agent generates tasks and is responsible for monitoring their execution.
  - The manager enters into explicit agreements with contractor agents willing to execute the tasks.
  - Individual agents are not designated a priori as manager or contractor - their roles at a given time, and any agent can take on either role dynamically during problem solving.
  - Manager agent advertises the existence of the tasks to other agents - potential contractors evaluate the task announcements and submit bids.
  - The manager evaluates the bids and awards contracts.
  - The manager and contractor are thus linked by a contract and communicate privately while the contract is being executed

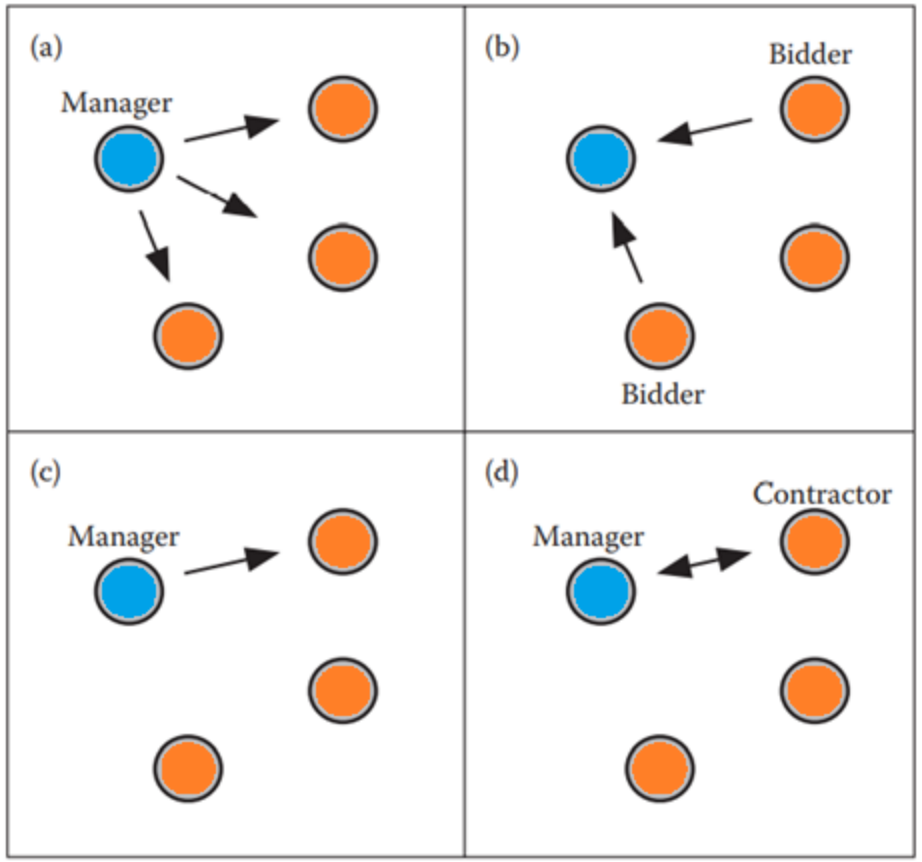

Contract nets: (a) Manager advertises a task; (b) potential contractors bid for the task; (c) manager awards the contract; (d) manager and contractor communicate privately.

Agent Interaction Models (2)

- Cooperative Problem-Solving
  - Top-down model for agent cooperation.

Stage 1: Recognition. Some agents recognize the potential for cooperation with an agent that is seeking assistance, possibly because it has a goal that it cannot achieve in isolation.

Stage 2: Team formation. An agent that recognized the potential for cooperative action at Stage1 solicits further assistance. If successful, forms a group having a joint commitment to collective action

Stage 3: Plan formation. The agents attempt to negotiate a joint plan.

Stage 4: Team action. The newly agreed plan of joint action is executed. By adhering to an agreed social convention, the agents maintain a close-knit relationship throughout.

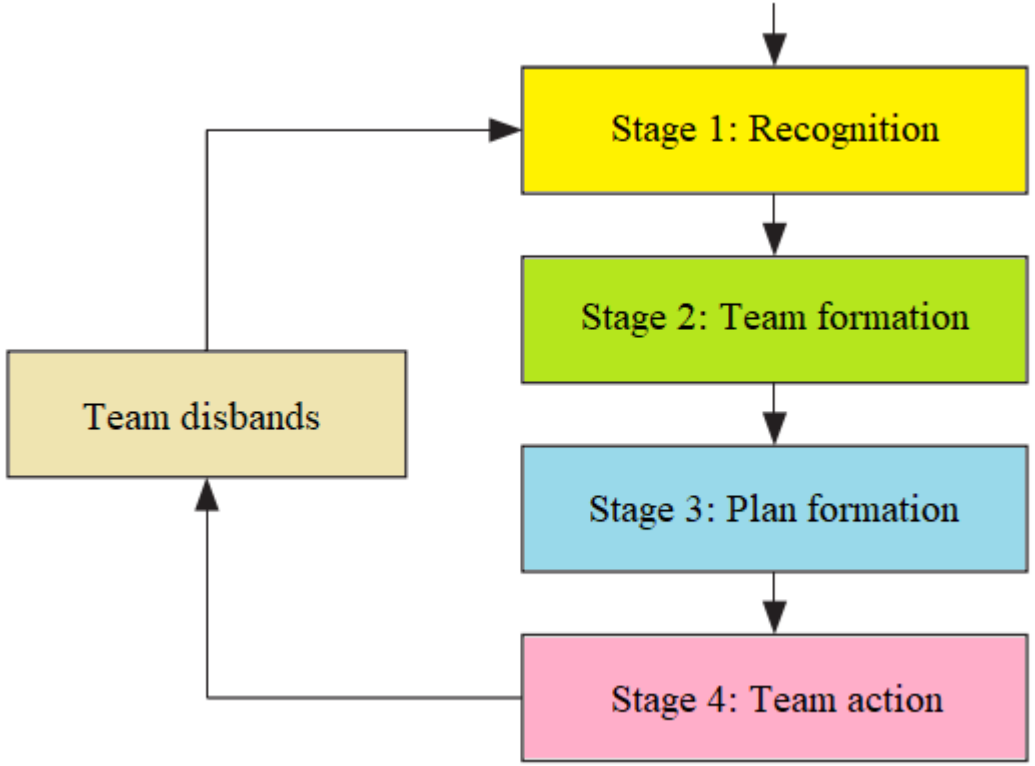

Agent Interaction Models (3)

- Shifting Matrix Management
  - Allows multiple lines of authority, reflecting the multiple functions expected of a flexible workforce
  - The lines of authority is regarded as temporary, typically changing as different projects start and finish.
  - The agents are distinguished by their different motives, functionality, and knowledge.

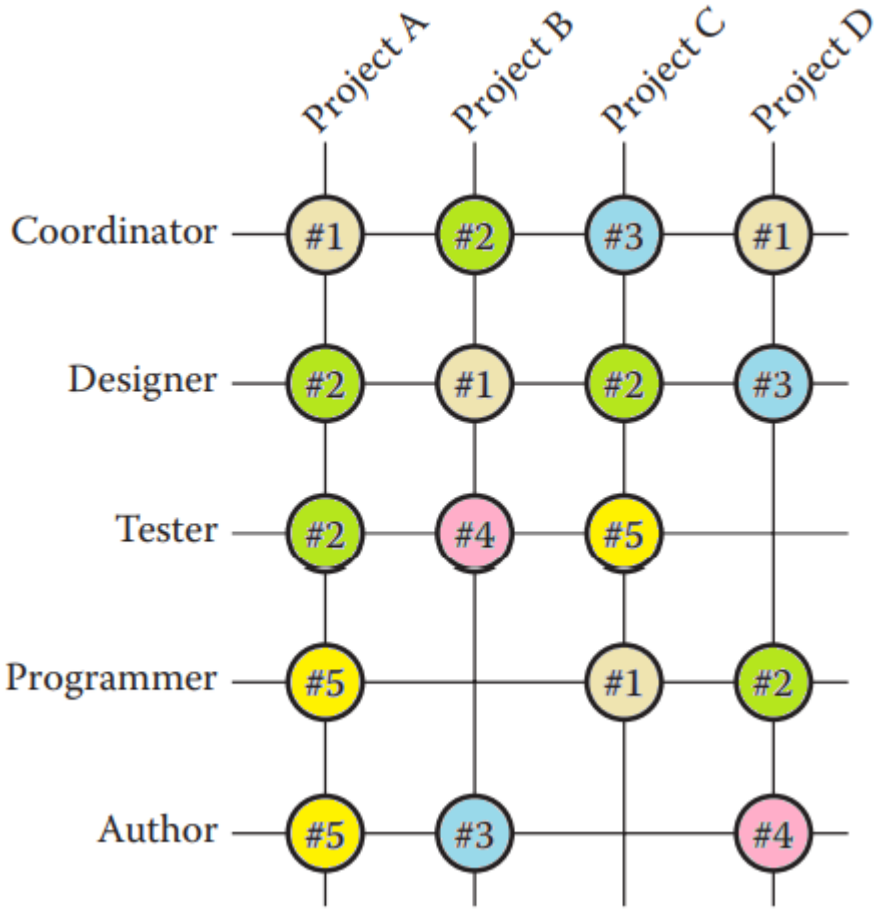

Agent Interaction Models (4)

- Shifting Matrix Management

Stage 1: Goal selection. Agents select the tasks they want to perform, based on their initial mental states.

Stage 2: Action selection. Agents select a way to achieve their goals.

Stage 3: Team formation. Agents that are seeking cooperation attempt to organize themselves into a team - requires an agreed code of conduct, a basis for sharing resources, and a common measure of performance.

Stage 4: Team planning. Workload is distributed among team members

Stage 5: Team action. Team plan is executed by the members under the team's code of conduct.

Stage 6: Shifting. Marks the disbanding of the team, involves shifting agents' goals, positions, and roles. Each agent updates its probability of team-working with other agents - iteration through the six stages takes place until all the tasks are accomplished.