

**AOT
LAB**

Agent and Object Technology Lab
Dipartimento di Ingegneria dell'Informazione
Università degli Studi di Parma



Multi-Agent Systems

Agents

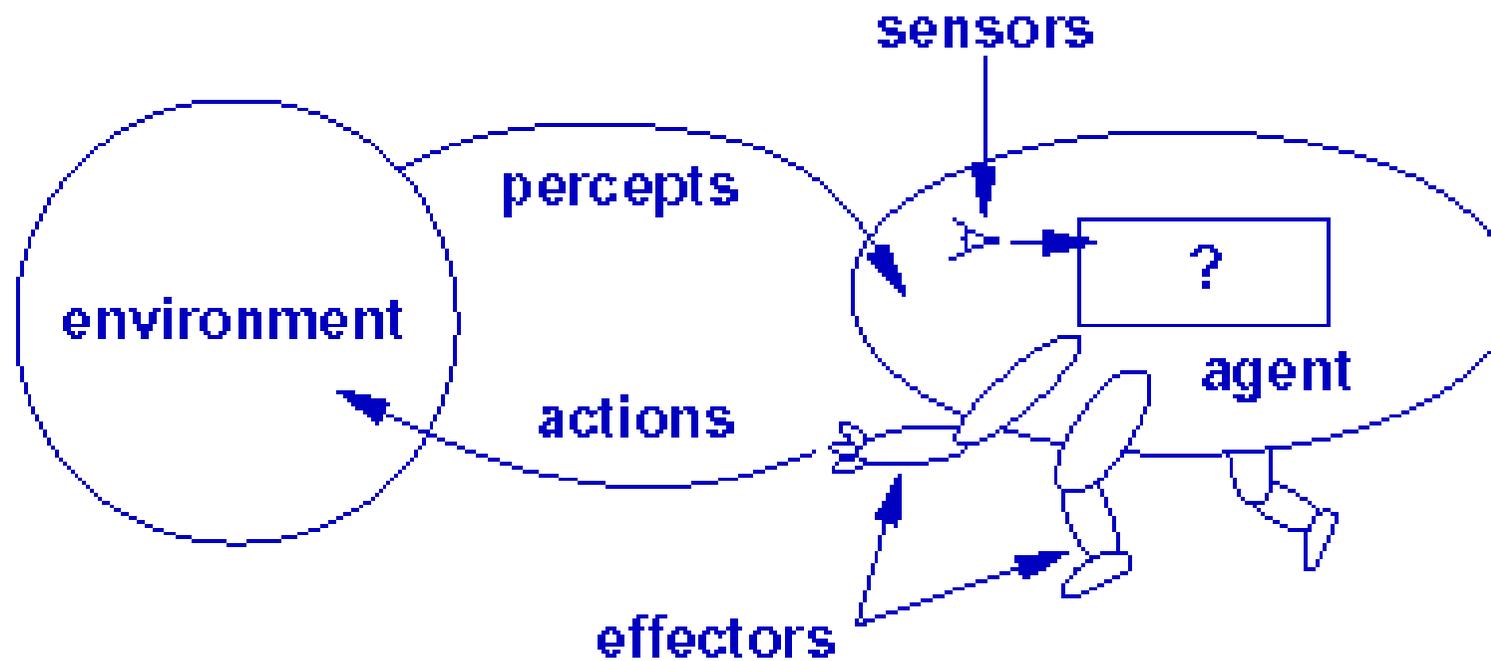
Prof. Agostino Poggi

- ◆ Distributed systems
- ◆ Game theory
- ◆ Software engineering
- ◆ Artificial life
- ◆ Databases
- ◆ Organizational science
- ◆ Artificial intelligence
- ◆ Economics

- ◆ A multi-agent system (MAS) consists of a set of agents that interact (cooperate, coordinate, etc.) with each other

- ◆ Multi-agent systems can be divided in two categories
 - Cooperative MAS
 - Distributed sensing and monitoring, distributed delivery, distributed problem solving, ...

 - Competitive MAS
 - Voting, auctions, ...

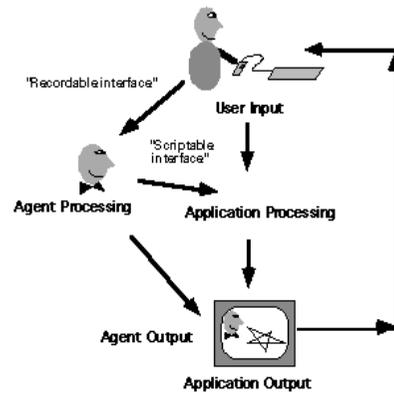
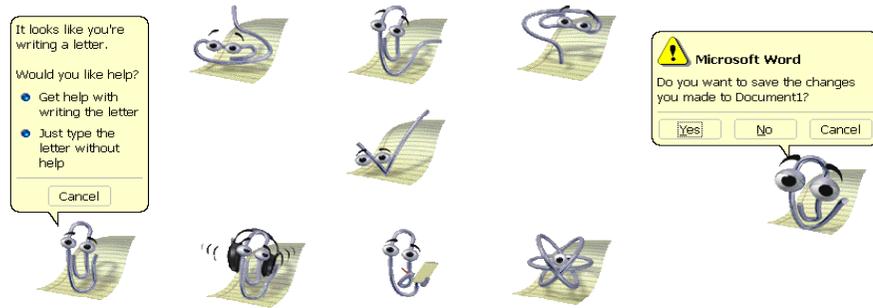


- ◆ Russell and Norvig: an agent is just something that **perceives** and **acts**

- ◆ Weiss: an agent is a computer system that is **situated** in some **environment** and capable of **autonomous** action in order to meet its design **objectives**

- ◆ **Physical or virtual entity**
- ◆ Owning **resources** of its own
- ◆ Capable of **perceiving** and **acting** in its environment
- ◆ Owning only a **partial representation** of its environment
- ◆ Capable of **communicating** directly with other agents
- ◆ Owning **skills** and providing **services**
- ◆ Driven by a set of **tendencies** towards **satisfying its objectives**, taking account of its resources and skills, perception and representations, and communications
- ◆ Capable of **reproducing** itself

- ◆ The main point about an agent is that it is **autonomous**: capable of
 - Acting independently
 - Exhibiting control over its internal state
- ◆ Thus, an agent is a computer system capable of autonomous action in some environment
- ◆ Therefore, trivial (non-interesting) agents are:
 - Any control system can be viewed as an agent (e.g., thermostat)
 - Most software daemons (e.g., background processes in the UNIX operating system)



- ◆ An intelligent agent is a computer system capable of **flexible** autonomous action in some environment

- ◆ By flexible, we mean
 - Reactive

 - Pro-active

 - Social

- ◆ If a program's environment is guaranteed to be fixed
 - Program needs never worry about its own success or failure
 - Program just executes blindly
- ◆ The real world is not like that
 - Things change, information is incomplete
 - Many interesting environments are dynamic
- ◆ Software is hard to build for dynamic domains
 - Program must take into account possibility of failure
 - Ask itself whether it is worth executing!
- ◆ A reactive system is one that
 - Maintains an ongoing interaction with its environment
 - Responds to changes that occur in it (in time for the response to be useful)

- ◆ Reacting to an environment is easy (e.g., stimulus response rules)
- ◆ But we generally want agents to do things for us without continuously asking for them
- ◆ Pro-activeness means
 - Generating and attempting to achieve goals
 - Not driven solely by events
 - Taking the initiative
 - Recognizing opportunities

- ◆ The real world is a multi-agent environment
 - We cannot go around attempting to achieve goals without taking others into account
 - Some goals can only be achieved with the cooperation of others
- ◆ Social ability in agents is
 - Ability to interact with other agents (and possibly humans)
 - Via some kind of agent-communication language
- ◆ Social ability implies coordination
 - Cooperation
 - Negotiate

◆ Rationality

- Agents act in order to achieve their goals, and do not act in such a way as to prevent their goals being achieved at least insofar as their beliefs permit
- Agents act to maximize their achievements

◆ *Benevolence*

- Agents do not have conflicting goals, and that every agent will therefore always try to do what is asked of it

◆ *Veracity*

- Agents do not knowingly communicate false information

- ◆ *Learning/adaptation*

- Agents improve performance over time

- ◆ *Mobility*

- Agents are able to move around a computer network

- ◆ Speedup & Efficiency
- ◆ Robustness & Reliability
- ◆ Scalability & Flexibility
- ◆ Cost
- ◆ Distributed Development
- ◆ Reusability

◆ Fundamental features

- An agent is *autonomous*
- An agent is *reactive*
- An agent is *social*

Autonomous Agents

Multi Agent Systems

Intelligent Agents

◆ Useful features

- An agent can be *proactive* (or *goal-directed*)
- An agent can be *adaptive* (or *learning*)
- An agent can be *mobile*

Mobile Agents

Learning Agents

- ◆ Are agents just objects by another name?

- ◆ Objects have the following features:
 - Encapsulate some state

 - Communicate via message passing

 - Have methods, corresponding to operations that may be performed on their state

- ◆ Agents are autonomous
 - Agents embody stronger notion of autonomy than objects
 - In particular, they decide for themselves whether or not to perform an action on request from another agent
- ◆ Agents are smart
 - Capable of flexible (reactive, pro-active, social) behavior
 - Standard object model has nothing to say about such types of behavior
- ◆ Agents are active
 - A multi-agent system is inherently multi-threaded, in that each agent is assumed to have at least one thread of active control

- ◆ Objects **do** actions for **free**
- ◆ Agents **do** actions because they **want** to
- ◆ Agents **do** actions for **money**

- ◆ Aren't agents just expert systems by another name?
- ◆ Expert systems typically disembodied 'expertise' about some (abstract) domain of discourse (e.g., blood diseases)
- ◆ Example: MYCIN knows about blood diseases in humans
 - It has a wealth of knowledge about blood diseases, in the form of rules
 - A doctor can obtain expert advice about blood diseases by:
 - Giving MYCIN facts
 - Answering questions
 - Posing queries

- ◆ Agents are situated in an environment
 - MYCIN is not aware of the world
 - Information is obtained by asking the user questions

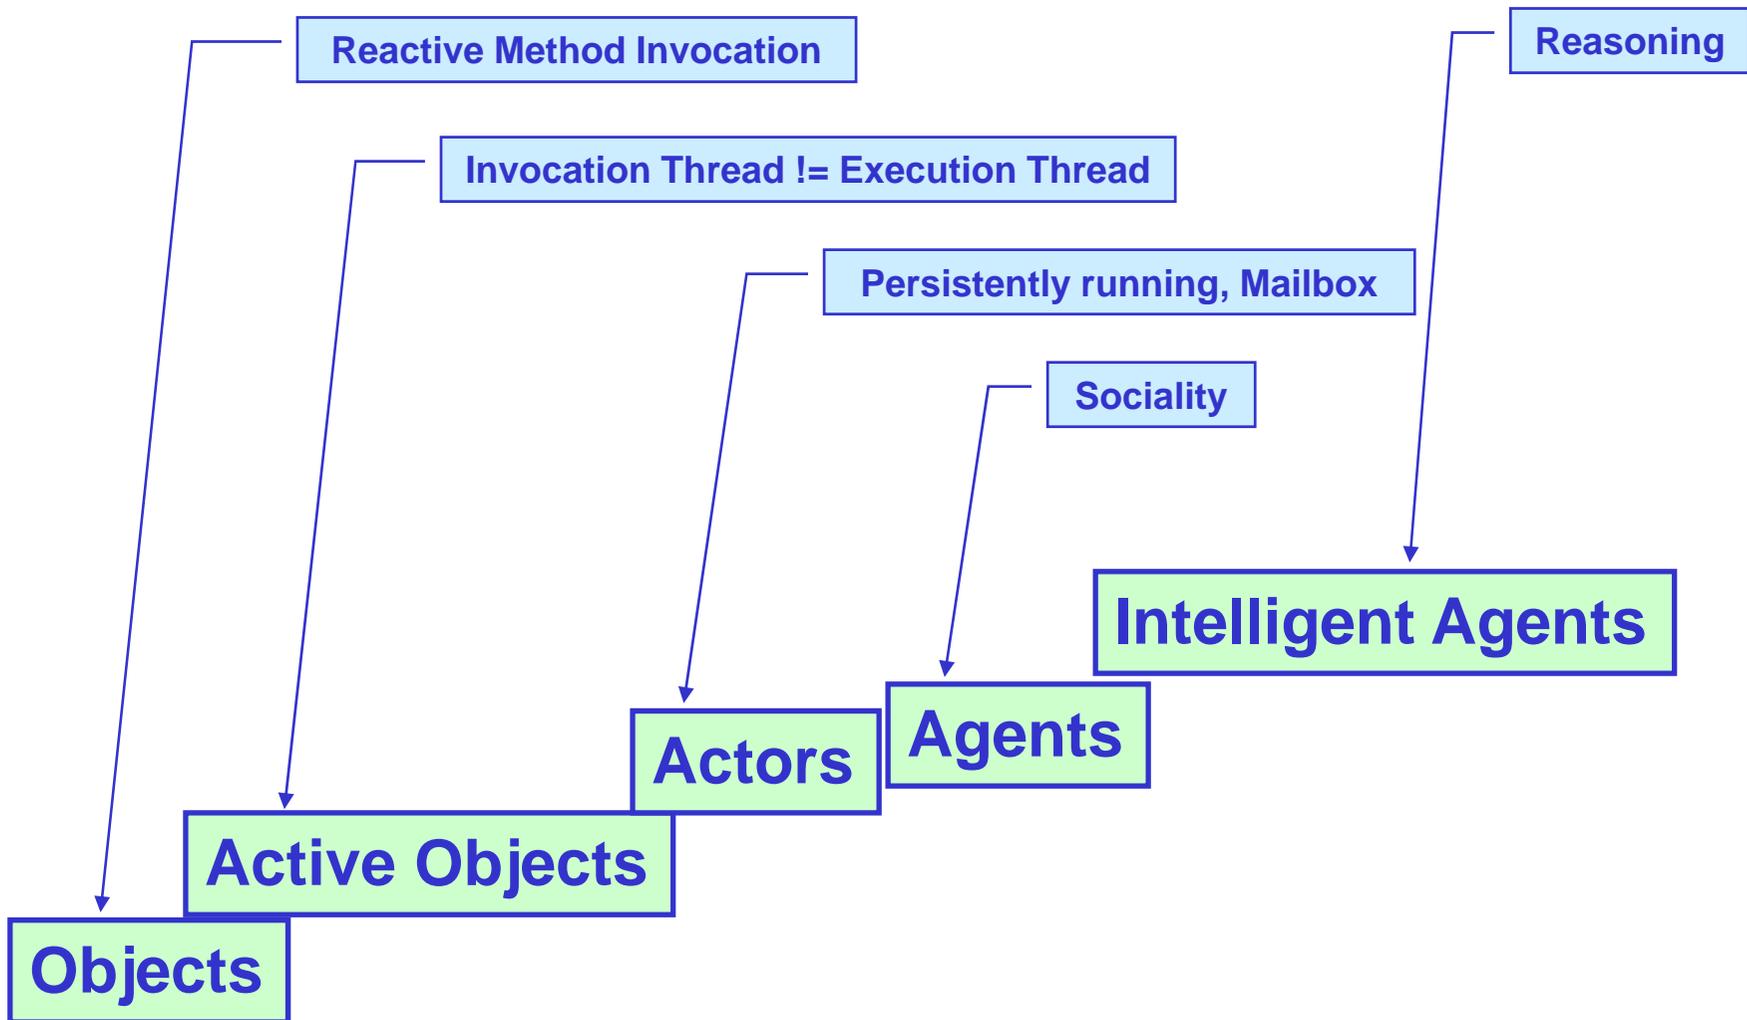
- ◆ Agents act
 - MYCIN does not operate on patients

- ◆ Some real-time (typically process control) expert systems are agents

- ◆ Aren't agents just the goal of the AI project?
- ◆ AI aims to build systems that can (ultimately)
 - Understand natural language
 - Recognize and understand scenes
 - Use common sense
 - Think creatively
 - ...
 - All of which are very hard
- ◆ So, don't we need to solve all of AI to build an agent?

- ◆ When building an agent, we simply want a system that can
 - Choose the right action to perform
 - Typically in a limited domain
- ◆ We **do not** have to solve **all** the problems of AI to build a useful agent
 - A **little intelligence** goes a long way!
- ◆ Oren Etzioni, speaking about the commercial experience of NETBOT, Inc
 - We made our agents dumber and dumber and dumber
 - . . . until finally they made money

- ◆ Distributed Artificial Intelligence (**DAI**) is the part of the Artificial Intelligence that evolved and actually is named Multi-Agent Systems
- ◆ DAI defined two classes of systems respectively based on
 - Behavior coordination
 - Several agents coordinate their knowledge and activities by reasoning about the problem solving process
 - Distributed Problem Solving
 - A particular problem is solved by dividing tasks among a number of generally equivalent nodes who divide and share knowledge about the problem
- ◆ MAS actually cover both the two types of systems



- ◆ The need of using and the difficult to build Multi-Agent Systems may depend on the properties of the environment in which the agent should act
- ◆ Russell and Norvig suggest the following classification of environment properties
 - Accessible vs inaccessible
 - Deterministic vs non-deterministic
 - Episodic vs non-episodic
 - Static vs dynamic
 - Discrete vs continuous

- ◆ An accessible environment is one in which the agent can obtain
 - Complete
 - Accurate
 - Up-to-dateinformation about the environment's state
- ◆ Most moderately complex environments (including, for example, the everyday physical world and the Internet) are inaccessible
- ◆ The more accessible an environment is, the simpler it is to build agents to operate in it

- ◆ As we have already mentioned, a deterministic environment is one in which
 - Any action has a single guaranteed effect
 - There is no uncertainty about the state that will result from performing an action
- ◆ The physical world can to all intents and purposes be regarded as non-deterministic
- ◆ Non-deterministic environments present greater problems for the agent designer

- ◆ In an episodic environment, the performance of an agent is
 - Dependent on a number of discrete episodes
 - With no link between the performance of an agent in different scenarios
- ◆ Episodic environments are simpler from the agent developer's perspective because
 - Agent can decide what action to perform based only on the current episode
 - It need not reason about the interactions between this and future episodes

- ◆ A static environment is one that
 - Can be assumed to remain unchanged
 - Except by the performance of actions by the agent
- ◆ A dynamic environment is one that
 - Has other processes operating on it
 - Which hence changes in ways beyond the agent's control
- ◆ The physical world is a highly dynamic environment

- ◆ An environment is discrete if there are a fixed, finite number of actions and percepts in it
- ◆ Russell and Norvig give
 - A chess game as an example of a discrete environment
 - Taxi driving as an example of a continuous one

Agents	Number (min 2) Uniformity (homogenous vs. heterogeneous) Goals (contradicting vs. complementary) Architecture (reactive vs. deliberative) Abilities of sensor & effectors (simple vs. advanced)
Interaction	Frequency (low vs. high) Persistence (short-term vs. long-term) Level (signal-passing vs. knowledge intensive) Flow of control (decentralized vs. centralized) Variability (fixed vs. changeable) Purpose (competitive vs. cooperative)
Environment	Predictability (foreseeable vs. unforeseeable) Accessibility (unlimited vs. limited) Dynamics (fixed vs. dynamic) Diversity (poor vs. rich) Availability of resources (restricted vs. ample)

Examples of Multi-Agent Systems

Agent Type	Percepts	Actions	Goals	Environment
Medical diagnosis System	Symptoms, findings, patient's answers	Questions, tests, treatments	Healthy patient, minimize costs	Patient, hospital
Satellite image analysis system	Pixels of varying intensity, color	Print a categorization of scene	Correct categorization	Images from orbiting satellite
Part-picking robot	Pixels of varying intensity	Pick up parts and sort into bins	Place parts in correct bins	Conveyor belt with parts
Refinery controller	Temperature, pressure readings	Open, close Valves, adjust temperature	Maximize purity, yield, safety	Refinery
Interactive English tutor	Typed words	Print exercises, suggestions, corrections	Maximize student's score on test	Students, course material

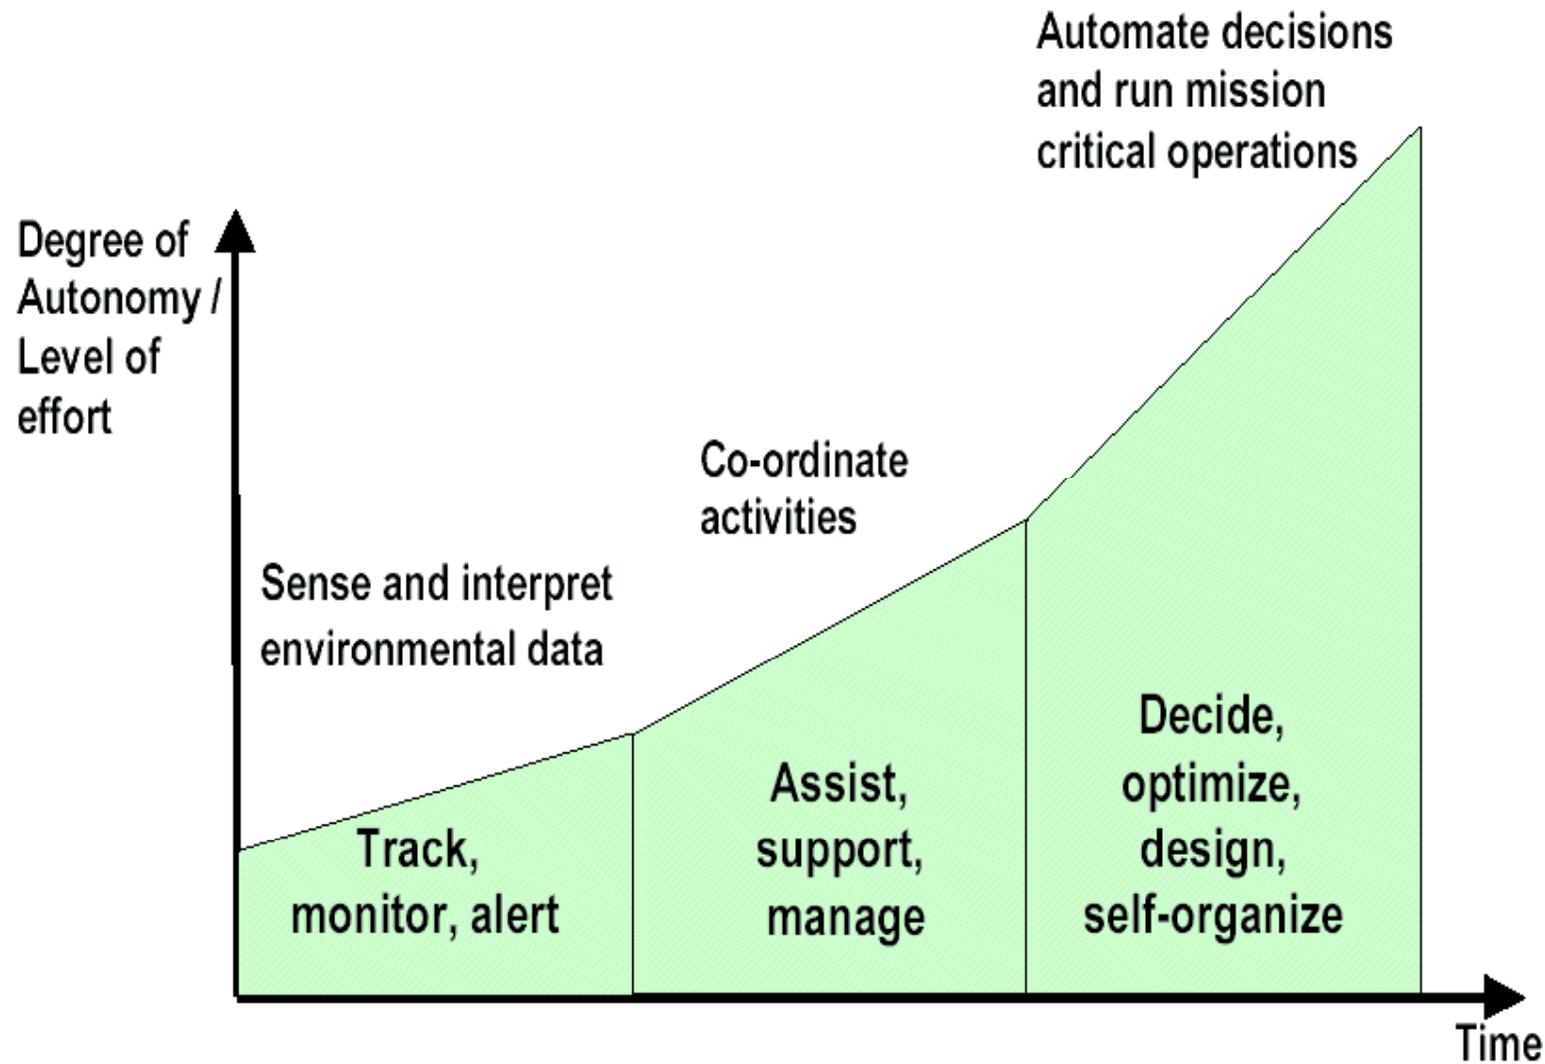
- ◆ Inherent Distribution

- Geographically
- Temporally
- Semantics – requires different ontologies and languages
- Functional – requires different cognitive capabilities

- ◆ Inherent Complexity

- Too large to be solved by single, centralized system

When to Use Multi-Agent Systems



- ◆ To mine big amount of data
- ◆ To assist with decision-making
- ◆ To coordinate self-interested and distributed processes
- ◆ To improve learning
- ◆ To delegate / automate (routine) tasks
- ◆ To personalize functions, anticipate situations and actions

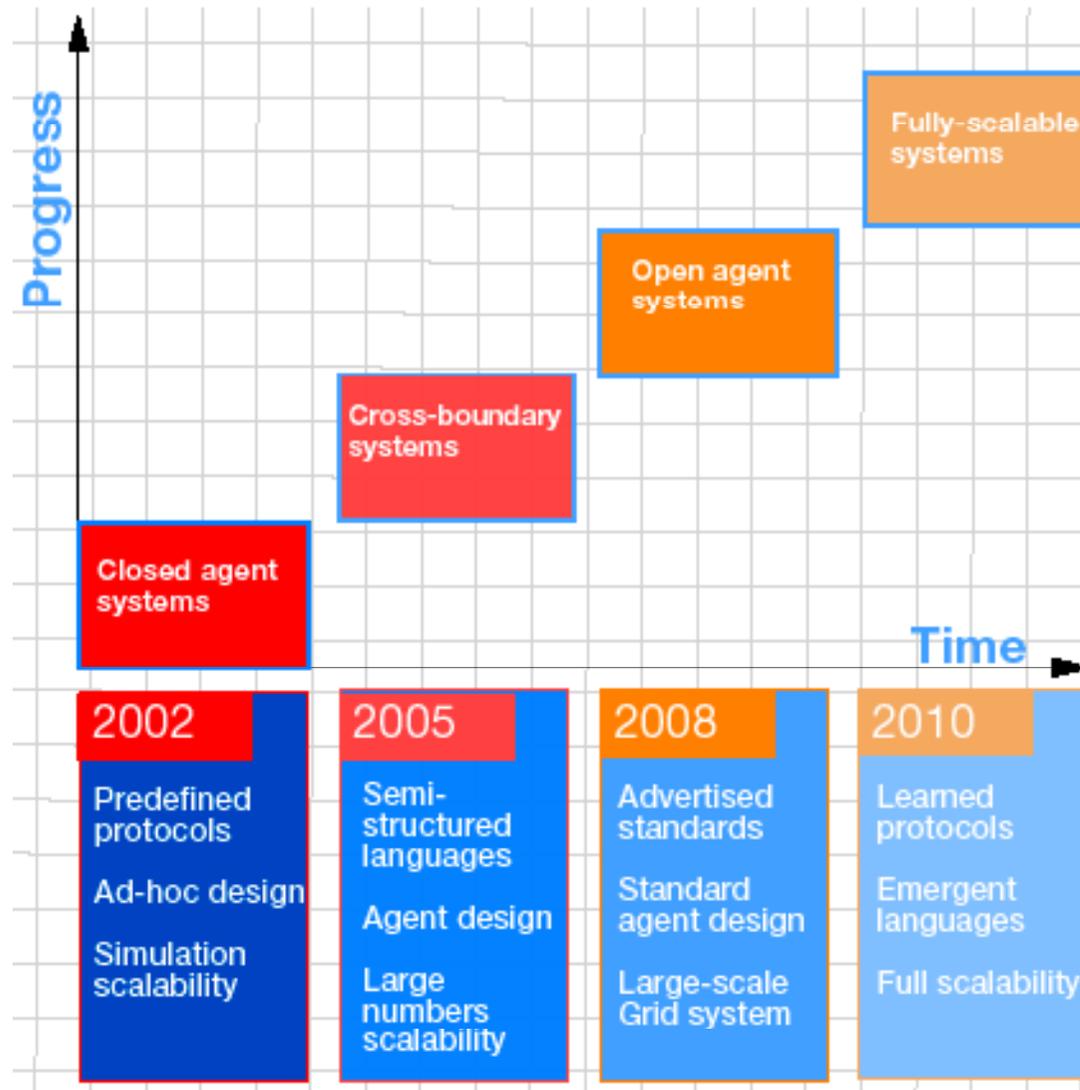


- ◆ Air traffic control
- ◆ Autonomic computing
- ◆ Automatic meeting scheduling
- ◆ Business process management
- ◆ Computer-based collaborative work support
- ◆ Distributed sensing
- ◆ E-commerce
- ◆ E-factory
- ◆ E-learning
- ◆ E-logistics
- ◆ Information gathering and handling
- ◆ Mail filtering
- ◆ Military
- ◆ Modeling and control of transportation systems
- ◆ Network Management
- ◆ Personal Assistant
- ◆ Power systems management
- ◆ Real-time monitoring and control of networks
- ◆ Space shuttle fault diagnosis
- ◆ Space exploration
- ◆ Tutoring systems

- ◆ How to enable agents to decompose their tasks and goals (and allocate sub-goals and sub-tasks to other agents) and synthesize partial results
- ◆ How to enable agents to communicate
- ◆ How to enable agents to represent and reason about the actions, plans, and knowledge of other agents in order to interact with them
- ◆ How to enable agents to represent and reason about the state of their interactions
- ◆ How to enable agents to recognize and handle conflicts between agents
- ◆ How to enable agents to negotiate and contract with each other

- ◆ How to engineer practical multi-agent systems
- ◆ How to effectively balance local computational versus communication
- ◆ How to avoid or mitigate harmful (chaotic or oscillatory) system wide behavior
- ◆ How to form and dissolve organizational structures to meet specific goals and objectives
- ◆ How to formally describe multi-agent systems and the interaction between agents
- ◆ How to ensure multi-agent systems are correctly specified
- ◆ How to realize *intelligent processes* such as problem solving, planning, decision making, and learning in a multi-agent system context
- ◆ How to enable agents to collectively carry out such processes in a coherent way

Future Trends (Agentlink II)



Future Trends (Forrester Research)

	Phase 1 (2003-2005)	Phase 2 (2006-2008)	Phase 3 (2009+)
	Agent-monitored	Agent-managed	Agent-optimized
Main purpose of using agents	- Sense/Interpret environmental data	- Coordinate activities with partners	- Decentralize and automate decisions
Killer app	- Multi-agent simulation systems	- Multi-agent execution systems	- Multi-agent optimization systems
Agent attributes exploited	- Goal-driven - Cooperation ability - Reactivity	- Mobility - Proactivity	- Autonomy - Adaptability - Learning
Application domains	- Financial risk management - Bioinformatics - National security	- Store-level replenishment - Predictive maintenance	- Adaptive supply networks - Domotics - Nanotechnology
Enabling sciences/standards	- Complexity science - Grid computing - Organic IT	- Web services (BPEL4WS) - AJML	- X Internet (e.g., RFID) - Semantic Web - Swarm Intelligence
Key vendors to watch	- IBM, Lost Wax, NuTech Solutions, Searchspace, Tryllian	- Agentis Software, BTextact Technologies, CGE&Y, HP, IBM, living systems, Microsoft, SAP	- Fujitsu, IBM, Microsoft, Oracle, SAP, Sony, Wipro

Source: Forrester Research, Inc.