

CPE/CSC 580: Intelligent Agents

Franz J. Kurfess

*Computer Science Department
California Polytechnic State University
San Luis Obispo, CA, U.S.A.*



Course Overview

- ❖ **Introduction**
 - ❖ Intelligent Agent, Multi-Agent Systems
 - ❖ Agent Examples
- ❖ **Agent Architectures**
 - ❖ Agent Hierarchy, Agent Design Principles
- ❖ **Reasoning Agents**
 - ❖ Knowledge, Reasoning, Planning
- ❖ **Learning Agents**
 - ❖ Observation, Analysis, Performance Improvement
- ❖ **Multi-Agent Interactions**
 - ❖ Agent Encounters, Resource Sharing, Agreements
- ❖ **Communication**
 - ❖ Speech Acts, Agent Communication Languages
- ❖ **Collaboration**
 - ❖ Distributed Problem Solving, Task and Result Sharing
- ❖ **Agent Applications**
 - ❖ Information Gathering, Workflow, Human Interaction, E-Commerce, Embodied Agents, Virtual Environments
- ❖ **Conclusions and Outlook**

Overview Introduction

- ❖ **Motivation**
- ❖ **Objectives**
- ❖ **What is an Intelligent Agent?**
 - ❖ Intelligence, Autonomy
- ❖ **Towards Agents**
 - ❖ Trends in Computing
- ❖ **Definitions and Concepts**
 - ❖ Agent, Multi-Agent System, Environments
- ❖ **Agent Examples**
 - ❖ Robots and Embodied Agents, E-Commerce, Personal Assistants
- ❖ **Research Issues**
- ❖ **Important Concepts and Terms**
- ❖ **Chapter Summary**

Motivation

- ❖ **introduce the concept of an agent and multi-agent system**
- ❖ **examine intelligent autonomous agents**
 - ❖ main approaches and techniques for the implementation of such agents
- ❖ **investigate multi-agent systems**
 - ❖ main approaches and techniques communication and cooperation in such systems;
- ❖ **explore aspects of a multi-agent society**
 - ❖ automated decision making in multi-agent contexts

Objectives

- ❖ Understand the fundamental concepts in the study of intelligent agents.
- ❖ Be familiar with the basic concepts, methods, techniques, and tools for the use of intelligent agents in computer-based systems.
- ❖ Understand the components and functions of intelligent agents.
- ❖ Apply the principles and methods of intelligent agents to a small-scale practical problem within the framework of a term project.
- ❖ Be prepared for further study in the design, implementation, and application of agent-based systems.
- ❖ Critically evaluate current trends in intelligent agents and their manifestation in business and industry.

Trends in Computing

Ubiquity
Interconnection
Intelligence
Delegation
Human-orientation

Ubiquity

- The continual reduction in cost of computing capability has made it possible to introduce processing power into places and devices that would have once been uneconomic
- As processing capability spreads, sophistication (and intelligence of a sort) becomes ubiquitous
- What could benefit from having a processor embedded in it...?

Interconnection

- Computer systems today no longer stand alone, but are networked into large distributed systems
- The internet is an obvious example, but networking is spreading its ever-growing tentacles...
- Since distributed and concurrent systems have become the norm, some researchers are putting forward theoretical models that portray computing as primarily a process of interaction

Intelligence

- The complexity of tasks that we are capable of automating and delegating to computers has grown steadily
- If you don't feel comfortable with this definition of "intelligence", it's probably because you are a human

Delegation

- Computers are doing more for us – without our intervention
- We are giving control to computers, even in safety critical tasks
- One example: fly-by-wire aircraft, where the machine's judgment may be trusted more than an experienced pilot
- Next on the agenda: fly-by-wire cars, intelligent braking systems, cruise control that maintains distance from car in front...

Human Orientation

- The movement away from machine-oriented views of programming toward concepts and metaphors that more closely reflect the way we ourselves understand the world
- Programmers (and users!) relate to the machine differently
- Programmers conceptualize and implement software in terms of higher-level – more human-oriented – abstractions

Programming progression...

- Programming has progressed through:
 - machine code;
 - assembly language;
 - machine-independent programming languages;
 - sub-routines;
 - procedures & functions;
 - abstract data types;
 - objects;
- to agents.

Global Computing

- What techniques might be needed to deal with systems composed of thousands of processors?
- Don't be deterred by its seeming to be "science fiction"
- Hundreds of millions of people connected by email once seemed to be "science fiction"...
- Let's assume that current software development models can't handle this...

Where does it bring us?

- Delegation and Intelligence imply the need to build computer systems that can act effectively on our behalf
- This implies:
 - The ability of computer systems to act independently
 - The ability of computer systems to act in a way that represents our best interests while interacting with other humans or systems

Interconnection and Distribution

- Interconnection and Distribution have become core motifs in Computer Science
- But Interconnection and Distribution, coupled with the need for systems to represent our best interests, implies systems that can cooperate and reach agreements (or even compete) with other systems that have different interests (much as we do with other people)

So Computer Science expands...

- These issues were not studied in Computer Science until recently
- All of these trends have led to the emergence of a new field in Computer Science:
 - intelligent agents, and
 - multi-agent systems

Definitions and Basic Concepts

**Agent
Multi-Agent Systems**

Definition: Agent

- An agent is
 - a computer system that is
 - capable of independent action on behalf of its user or owner
 - (figuring out what needs to be done to satisfy design objectives, rather than constantly being told)

Definition: Multi-agent Systems

- A multi-agent system is one that consists of a number of agents, which interact with one another
 - In the most general case, agents will be acting on behalf of users with different goals and motivations
 - To successfully interact, they will require the ability to cooperate, coordinate, and negotiate with each other, much as people do

Agent Design

- ❖ construction of individual agents
- ❖ capable of independent, autonomous action
- ❖ can successfully carry out tasks we delegate to them

Agent Society Design

- ❖ **design of collaborating agents**
- ❖ **capable of interacting with other agents**
 - ❖ cooperation
 - ❖ coordination
 - ❖ negotiation
- ❖ **successfully carry out delegated tasks**
 - ❖ delegation may come from humans or other agents
- ❖ **other agents cannot be assumed to share the same interests/goals**

Multi-Agent Systems

- addresses questions such as:
 - emergence of cooperation in societies of self-interested agents
 - languages agents can use to communicate
 - recognition of conflicts among self-interested
 - how can they (nevertheless) reach agreement
 - coordination of activities among autonomous agents
 - to cooperatively achieve goals

Multi-Agent Systems

- these questions are all addressed in part by other disciplines
 - notably economics and social sciences,
- what makes this field unique:
 - emphasizes that the agents in question are computational, information processing entities

Agent Examples

NASA Deep Space 1

Air Traffic Control


Autonomous Vehicles

Internet Agents

Personal Agents

Deep Space 1

- “Deep Space 1 launched from Cape Canaveral on October 24, 1998. During a highly successful primary mission, it tested 12 advanced, high-risk technologies in space. In an extremely successful extended mission, it encountered comet Borrelly and returned the best images and other science data ever from a comet. During its fully successful hyperextended mission, it conducted further technology tests. The spacecraft was retired on December 18, 2001.”
- NASA Web site <http://nmp.jpl.nasa.gov/ds1/>



Autonomous Agents for specialized tasks

- DS1 example is one of a generic class
 - high-risk situations, unsuitable or impossible for humans
- varying degrees of autonomy
 - depending on the situation
 - remote human control may be an alternative, but not always

Air Traffic Control

- “A key air-traffic control system...suddenly fails, leaving flights in the vicinity of the airport with no air-traffic control support. Fortunately, autonomous air-traffic control systems in nearby airports recognize the failure of their peer, and cooperate to track and deal with all affected flights.”
- Systems taking the initiative when necessary
- Agents cooperating to solve problems beyond the capabilities of any individual agent

Autonomous Airplanes

- ❖ **current large airplanes have most of the technologies for autonomous flight**
 - ❖ control of individual airplanes
 - ❖ coordination of activities with other airplanes relies on central authorities
 - ❖ requires dealing with non-autonomous airplanes
 - ❖ older or smaller ones
 - ❖ proposals have been put forward to allow more autonomy
- ❖ **UAVs for military and security applications are in use**
 - ❖ most of them are to a large degree remotely controlled

Internet Agents

- ❖ tons of chat bots
- ❖ sales and marketing agents
- ❖ customer service


IKEA's Virtual Assistant Anna

❖ <http://193.108.42.79/ikea-caen/cgi-bin/ikea-caen.cgi>

CAL POLY

© Franz J. Kurfess

IKEA Help Centre



Are all of your colleagues in different countries called Anna?

You are fully entitled to your opinion. Just ask me what you'd like to know about IKEA and I'll help you.

Go

Apple's Siri

- ❖ **From Apple's Siri FAQ:**
 - ❖ "Siri is the intelligent personal assistant that helps you get things done just by asking. It allows you to use your voice to send messages, schedule meetings, place phone calls, and more. But Siri isn't like traditional voice recognition software that requires you to remember keywords and speak specific commands. Siri understands your natural speech, and it asks you questions if it needs more information to complete a task."
- ❖ <http://www.apple.com/iphone/features/siri.html>



Research Issues

Human-Agent Interaction
Agreements and Negotiation
Scaling Issues
Experiments
Interdisciplinary Aspects

Research Issues

- Human-Agent Interaction
 - How do you state your personal preferences to your agent?
 - How can your agent compare different deals from different vendors?
 - What if there are many different parameters?
- Agreements and Negotiation
 - What algorithms can your agent use to negotiate with other agents (to make sure you get a good deal)?
- Scaling Issues
 - expanding agent use from individuals to large groups
 - automated procurement for corporations or government agencies
- Experiments
 - Robocup
 - Rescue Agent Competitions
 - Trading Agents Competition

Interdisciplinary Aspects

- The field of Multi-Agent Systems is influenced and inspired by many other fields:
 - Economics
 - Philosophy
 - Game Theory
 - Logic
 - Ecology
 - Social Sciences
- Strength
 - infusing well-founded methodologies into the field
- Weakness
 - there are many different views as to what the field is about
- Analogies with artificial intelligence itself

Summary Introduction

- ❖ **Intelligent Agents can be viewed as an evolutionary development in computer science**
 - ❖ ubiquity, distributed systems, communication, collaboration
 - ❖ autonomy, delegation, human-centric orientation
- ❖ **Multi-Agent Systems integrate possibly large numbers of autonomous agents**
 - ❖ resource sharing, negotiation, collaboration, distributed problem solving
- ❖ **Agents have been used in various application areas**
 - ❖ space exploration, autonomous vehicles, personal assistants

Important Concepts and Terms

- ❖ agent
- ❖ agent society
- ❖ agreement
- ❖ delegation
- ❖ environment
- ❖ intelligence
- ❖ multi-agent system
- ❖ negotiation
- ❖ resource sharing
- ❖ ubiquity