Artificial Intelligence: Past and Present

Alex Kutsenok
April 15, 2005

About Me

- B.S. in Computer Science at the Rose-Hulman Institute Of Technology
- Currently a Ph.D. student at Michigan State University
- My Research Area: Artificial Intelligence
  - Specifically: Swarm Intelligence

What I Will Talk About Today

- Artificial Intelligence!
  - The key concepts of the past
  - Where research is headed today

What is AI?

- Artificial Intelligence is about programming computers to solve problems
- The science of designing intelligent decision-makers
- Want the computer to be as autonomous as possible, work quickly, and be reliable
- Want to minimize the time and effort it takes to program the computer

Some applications of Artificial Intelligence

- Playing chess
- Unmanned planes used by the US army
- Robot assistants for microsurgery
- A computer on board NASA's spacecraft that schedules maintenance tasks
- Diagnosing diseases from symptoms

- Driving a Car across California without human supervision ???
  - When do you think this will happen?

So how do you make a car to drive itself across California?

- Start with a normal car
- Add some cameras that become its “eyes”
- Add a laptop and connect it to the cameras, the steering wheel, and the pedals
- Write a program that
  - takes as input information from the cameras and
  - makes decisions about steering the wheel and pressing the pedals

Designing a Decision Maker

- Given some information
- Have multiple actions to choose from
- Have to decide which actions to carry out based on the information you have

- Agent= Intelligent Decision Maker

Discussion of Traditional AI Approaches

- I will give a description of the basic categories
- Each approach is very effective for some problems
- Each approach has drawbacks that make it unsuitable for other problems

Traditional Ways of Designing Agents

- Rule-Based
- Searching
- Logic manipulation
- Probabilistic Reasoning
- Learning

Rule-Based Agents

- Have pre-programmed rules for what to do
- Ex:

1) If turning, turn on the turn-signal
2) If changing lanes, first look behind you
3) If see deer ahead, hit the breaks.

…

- Used in ATM Machines, TVs, …
- Problems?

Problems with Rule-Based Agents

- Need many rules for every possible situation
  - What if you see a cow?
- All the Rules have to be written by the designer, that's a lot of work
- Rules are very specific, lack generality
  - A rule-based car-driving agent would have no idea how to ride a bike or walk

Searching Agents

- They think ahead to find a solution
  - Like using yarn to find the way out of a maze
- Useful for problems like chess where it helps to think 5 moves ahead
- It takes a lot of time to find solutions
  - Effective for "simple" problems like chess
  - Can take a near-infinite amount of time for more complex problems
- Can this kind of agent handle the car driving problem?

Logical Agents

- Use rules of logic to think about facts
- Ex: If know that A-> B and that B->C then can infer that A->C
- This is really effective for proving math theorems but not practical for driving a car
  - You don't want to be proving to yourself that you should break when you see a deer
  - Deer-> Large Animal -> Impact will damage car -> Glass will break -> …
  - Takes a lot of time, like searching

Probabilistic Reasoning Agents

- Useful for when not all the information is available
- A human doctor has a limited amount of information about a patient's health
  - Can't see inside the patient and know 100%
  - Uses symptoms to find the most probable cause of illness
- Probabilistic Agents act similarly
  - Handle uncertainty by using expected values, Bayes' Theorem, and other probability formulas
  - Find the P(Y given X) = P(X given Y) * P(Y) / P(X)

Drawbacks of Probabilistic Agents

- These math formulas don't describe how humans actually make decisions
- Probabilistic information is available in only a few problem domains
  - Definitely not in car driving
  - If you see deer and want to calculate whether it is a threat, need to know whether the deer will freeze on the road or run across before your car comes too close
  - No way to know that probability!

Learning Agents

- Seems like a great idea
- Agent starts small and gains knowledge as it experiences the environment
- The pressure to design an agent that can immediately start solving a problem is lifted from the shoulders of the designer
- The designer only needs to create a learning mechanism for the agent
- The rest of the work will be done by agent

Different Kinds of Learning

- Supervised Learning
  - Learns from previous cases it is given
- Unsupervised Learning
  - Tries things on its own and figures out what it is effective
  - Learning Mechanism must reward positive behavior
  - Vacuum Cleaner Example

Drawbacks of Learning Agents

- Speed of learning is an issue
  - The more complex a problem, the longer it takes for an agent to be trained
- Need for representation
  - The way information is represented affects the performance potential of the agent

Representation Example

- Let's say we know vacuum cleaner will work in a 10 by 10 ft room
- We can break up room into 1ft by 1ft squares and say that each square has some dirt or no dirt
- Or we can break up room into 1 in by 1 in squares
  - This makes the vacuum cleaner much more precise
  - But now it has much more information to process
- The way information is represented plays a crucial role in the way all AI agents work

Representation for Learning Agents

- It is hard for a designer to make a representation that
  - Gives the agent access to the most relevant information
  - Without making the agent's understanding too narrow or limited
- Representation decisions have to be made before the agent starts learning
  - So we see that the designer still bears a large weight on his/her shoulders
  - Learning agent's success depends on representation decision
  - the designer is still the "brains" of the operation

Current Work in AI

- New Applications of Traditional Approaches
- Extending Traditional Approaches
- Combining Traditional Approaches
- Creating New Approaches and Evaluating Them

My specific research area: Swarm Intelligence

- Swarm Intelligence (SI) is a sub-field of AI
- Instead of having 1 agent be responsible for the decision-making, have a team of agents
- SI explores the usefulness of teams of simple agents in solving various problems
- These agents are often inspired by insects, such as ants and bees

Why Do We Think Ants Are Smart

- One ant is pretty dumb
- But a colony of ants can solve hard problems
- Example: Foraging For Food
  - Ants can find the shortest path to a food source
  - Inform others about it and show them how to get there
  - Do this with a very simple method of communicating
    - Pheromone trails

Why Are the Ants Important

- The Food Foraging problem is a lot like networking problems that have to be solved for the Internet to work
- In both cases, it is important to find shortest routes
- In the early 90s, SI work began when a very successful algorithm was designed based on how ants forage for food
  - Ant System by Marco Dorigo
  - The Ant System is very popular now and has been used to solve many engineering problems

My work

- I have developed a SI design approach called Swarm AI
- I have applied it to different problems
- Swarm Soccer (real-time domain)
  - How to control a team of soccer players
- Swarm Queens to solve the N-Queens problem (classic CS/Math problem)
  - How to arranges queens on a N-by-N board so they don't threaten each other

What Swarm AI is About

- Divide the problem into parts
- Give each part to an agent
- Agents are simple and have a local perspective
- Agents communicate with each other indirectly
- The solution emerges from the interaction of these agents

Pros and Cons of Swarm Intelligence

- Pro: Swarm AI agents work very quickly
- Pro: The designer has to create small agents, not one big one that must solve the whole problem
- Con: Behavior can be unpredictable, further work must be done to better understand how groups of agents interact

Questions?

Thank you for your time

If I have time…talk about

- Evolutionary Computation
  - Many different agents compete
  - Most successful ones survive
  - Pass genes to off-spring
  - Have cross-over and mutation
- Developmental Learning

The Future

- How far are we from human-level intelligence?