




Introduction

Statistical Relational Artificial Intelligence
(StaRAI)



$\infty \text{ Presents}(X, P, C) \Rightarrow \text{Attends}(X, C)$

$3.75 \text{ Publishes}(X, C) \wedge \text{FarAway}(C) \Rightarrow \text{Attends}(X, C)$

Contents

1. Introduction

- Notions of AI, agents, StaRAI

2. Modelling Formalisms

- a. *Logic*
- b. *Probabilistic Graphical Models (PGMs)*
- c. *Probabilistic Relational Models (PRMs)*
- Syntax, semantics, inference tasks

3. Inference

- a. *Exact Inference*
 - i. *(Weighted) Model Counting*
 - ii. *Variable Elimination*
 - iii. *Lifted Inference*

- b. *Approximate Inference*

- Inference algorithms

4. Learning

- Relation learning in logic
- Parameter / structure learning for PGMs
- From PGMs to PRMs, PRM learning

5. Sequential Models and Inference

- Syntax, semantics, inference tasks, algorithms

6. Decision-theoretic Models and Inference

- Preferences, utility
- Syntax, semantics, inference task, algorithm

Overview: 1. Introduction

a. Artificial Intelligence (AI)

- Approaches: thinking / acting humanly / rationally

b. Framework: Agent Theory

- Agent
- Task environment
- Agent structure

c. Topic: Statistical Relational AI

- Motivation, context
- Relational examples, outlook on probabilistic relational models (PRMs)

Approaches to Artificial Intelligence (AI)

- All approaches researched
 - Supported and hindered each other
- **Rationality**
 - System is rational if it does the “right thing,” given what it knows

Success measure		
Fidelity of human performance	Ideal performance measure rationality	
Thinking Humanly	Thinking Rationally	Thought processes, reasoning
“The exciting new effort to make computers think . . . machines with minds, in the full and literal sense.” (Haugeland, 1985) “[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning...” (Bellman, 1978)	“The study of mental faculties through the use of computational models.” (Charniak and McDermott, 1985) “The study of the computations that make it possible to perceive, reason, and act.” (Winston, 1992)	
Acting Humanly	Acting Rationally	Behaviour
“The art of creating machines that perform functions that require intelligence when performed by people.” (Kurzweil, 1990) “The study of how to make computers do things at which, at the moment, people are better.” (Rich and Knight, 1991)	“Computational Intelligence is the study of the design of intelligent agents.” (Poole et al., 1998) “AI ... is concerned with intelligent behaviour in artefacts.” (Nilsson, 1998)	


Acting Humanly

- **Turing Test** (Turing, 1950)
 - Computer passes test, if a human, who asks written questions, cannot tell if the the written answers come from a human or not
 - Example: *Eliza*, program for superficially simulating a psychiatrist
 - See also Ch. 26, “Artificial Intelligence – A Modern Approach” by Russel & Norvig, including a discussion whether a computer would really be intelligent if it passed
 - Regarding Eliza: human’s example closure tendencies are more pronounced for emotional/social intelligence aspects
 - Cf. robot *Shakey*: No on who saw Shakey the first time thought it could shoot hoops, yet the first people interacting with Eliza assumed it was a real doctor
 - **Total** Turing Test: includes a video signal to test perceptual abilities, opportunity to pass physical objects

```
Welcome to
EEEEEE LL      IIII 2222222 AAAAA
EE      LL      II   ZZ  AA  AA
EEEEEE LL      II   ZZZ  AAAAAAA
EE      LL      II   ZZ  AA  AA
EEEEEE LLLLLL IIII 2222222 AA  AA

Eliza is a mock Rogerian psychotherapist.
The original program was described by Joseph Weizenbaum in 1966.
This implementation by Norbert Landsteiner 2005.

ELIZA: Is something troubling you ?
YOU:  Men are all alike.
ELIZA: What is the connection, do you suppose ?
YOU:  They're always bugging us about something or other.
ELIZA: Can you think of a specific example ?
YOU:  Well, my boyfriend made me come here.
ELIZA: Is it important to you that your boyfriend made you come here ?
YOU:  He says I'm depressed much of the time.
ELIZA: I am sorry to hear that you are depressed.
YOU:  It's true. I am unhappy.
ELIZA: Can you explain what made you unhappy ?
YOU:  █
```



Acting Humanly

- Subproblems to solve as part of the Turing Test
 - *Natural Language Processing*
 - Communication
 - *Knowledge representation*
 - Store knowledge and observations
 - *Automated reasoning*
 - Answer questions, draw new conclusions
 - *Machine learning*
 - Adapt to new circumstances, detect and extrapolate patterns
 - Total Turing Test
 - *Computer vision*: perceive objects
 - *Robotics*: manipulate objects, move about

The Turing Test covers a majority of disciplines that make up AI nowadays.

- But:
 - little research effort devoted to pass test
- Instead:
 - Study underlying principles of intelligence

Approaches to Artificial Intelligence (AI)

- All approaches followed
 - Supported and hindered each other
- **Rationality**
 - System is rational if it does the “right thing,” given what it knows

Success measure		Thought processes, reasoning
Fidelity of human performance	Ideal performance measure rationality	
Thinking Humanly	Thinking Rationally	
“The exciting new effort to make computers think . . . machines with minds, in the full and literal sense.” (Haugeland, 1985) “[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning...” (Bellman, 1978)	“The study of mental faculties through the use of computational models.” (Charniak and McDermott, 1985) “The study of the computations that make it possible to perceive, reason, and act.” (Winston, 1992)	
Acting Humanly	Acting Rationally	Behaviour
“The art of creating machines that perform functions that require intelligence when performed by people.” (Kurzweil, 1990) “The study of how to make computers do things at which, at the moment, people are better.” (Rich and Knight, 1991)	“Computational Intelligence is the study of the design of intelligent agents.” (Poole et al., 1998) “AI ... is concerned with intelligent behaviour in artefacts.” (Nilsson, 1998)	

Thinking Humanly

- A “program thinks like a human”
 - Requires a way to determine how humans think → workings of the human mind
 - Given theory of the mind, express theory as computer program
 - If program’s input-output behaviour matches corresponding human behaviour, evidence that some of program’s mechanisms could also be operating in humans
- Approach complementary to AI: *Cognitive Science*
 - Interdisciplinary:
 - Computer models from AI
 - Experimental techniques from psychology
 - Goal:
Construct precise and testable theories of human mind

Approaches to Artificial Intelligence (AI)

- All approaches followed
 - Supported and hindered each other
- **Rationality**
 - System is rational if it does the “right thing,” given what it knows

Success measure		Thought processes, reasoning
Fidelity of human performance	Ideal performance measure rationality	
Thinking Humanly	Thinking Rationally	
“The exciting new effort to make computers think . . . machines with minds, in the full and literal sense.” (Haugeland, 1985) “[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning...” (Bellman, 1978)	“The study of mental faculties through the use of computational models.” (Charniak and McDermott, 1985) “The study of the computations that make it possible to perceive, reason, and act.” (Winston, 1992)	
Acting Humanly	Acting Rationally	Behaviour
“The art of creating machines that perform functions that require intelligence when performed by people.” (Kurzweil, 1990) “The study of how to make computers do things at which, at the moment, people are better.” (Rich and Knight, 1991)	“Computational Intelligence is the study of the design of intelligent agents.” (Poole et al., 1998) “AI ... is concerned with intelligent behaviour in artefacts.” (Nilsson, 1998)	

Thinking Rationally

- Codify thinking → rules
 - Irrefutable reasoning processes
 - Argument structures that always yield correct conclusions when given correct premises
- Field of *Logic*
 - Precise notation for statements about objects in a world and relations among them
 - Programs that could, in principle, solve *any* solvable problem described in logical notation
 - Obstacles:
 - Informal knowledge
 - Unstructured data
 - Uncertainty
 - Solving any solvable problem in practice
 - Limited computational resources

Obstacles apply to *any* attempt to build computational reasoning systems

- Formulated first in logic

Approaches to Artificial Intelligence (AI)

- All approaches followed
 - Supported and hindered each other
- **Rationality**
 - System is rational if it does the “right thing,” given what it knows

Success measure		
Fidelity of human performance	Ideal performance measure rationality	
Thinking Humanly “The exciting new effort to make computers think . . . machines with minds, in the full and literal sense.” (Haugeland, 1985) “[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning...” (Bellman, 1978)	Thinking Rationally “The study of mental faculties through the use of computational models.” (Charniak and McDermott, 1985) “The study of the computations that make it possible to perceive, reason, and act.” (Winston, 1992)	Thought processes, reasoning
Acting Humanly “The art of creating machines that perform functions that require intelligence when performed by people.” (Kurzweil, 1990) “The study of how to make computers do things at which, at the moment, people are better.” (Rich and Knight, 1991)	Acting Rationally “Computational Intelligence is the study of the design of intelligent agents.” (Poole et al., 1998) “AI ... is concerned with intelligent behaviour in artefacts.” (Nilsson, 1998)	Behaviour

Acting Rationally

- Rational agent approach
- **Agent** = something that acts
 - Operate autonomously
 - Perceive environment
 - Persist over a prolonged time period
 - Adapt to change
 - Create and pursue goals
- **Rational** agent
 - One that acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome
 - May include thinking rationally or acting humanly, but *more general*

Advantage: Standard of rationality mathematically well defined

- Better suited to generate agent designs that provably achieve rationality
- Focus of the next slides

Interim Summary

- Four approaches to define AI
 - Acting humanly
 - Turing test
 - Subproblems: Natural language processing, knowledge representation, reasoning, machine learning, computer vision, robotics
 - Thinking humanly
 - Cognitive sciences as an interdisciplinary science between AI and psychology
 - Thinking rationally
 - Logic for knowledge representation and correct reasoning
 - Acting rationally
 - Rational agent as a generalisation of thinking rationally with a mathematical definition of rationality as a formal criterion
 - Also contains subproblems, which the Turing test identifies

Overview: 1. Introduction

A. *Artificial Intelligence*

- Approaches: thinking / acting humanly / rationally

B. **Framework: Agent Theory**


- Agent
- Task environment
- Agent structure

C. *Topic: StaRAI*

- Motivation, context
- Relational examples, outlook on probabilistic relational models (PRMs)


Agent

- Something that **perceives** its environment through **sensors** and **acts** through **actuators**
 - Human agent
 - Sensors: eyes, ears, further organs
 - Actuators: hands, legs, mouth, other body parts
 - Robot agent
 - Sensors: cameras, infra-red sensors, etc.
 - Actuators: motors
 - Software agent
 - Goal: *document retrieval, DR*
 - Sensors: input interface for textual queries
 - Actuators: output interface for documents




Agent Behaviour

- **Agent function**
 - Mapping from percept sequences to actions
$$f : P^* \rightarrow A$$
 - P^* Set of possible percept sequences
 - A Set of available actions
- **Agent program**
 - Implementation of the agent function f
 - Runs on a physical system (**architecture**)
- **Agent = architecture + agent program**



Simple Example

- Vacuum cleaner
 - Two locations: squares *A*, *B*
 - Possible percepts: location; location *clean*, *dirty*
 - Available actions: *right*, *left*, *vacuum*
- Agent function in table
 - Mapping increasingly longer percept sequences to actions
 - Optimisable, as action only depends on last observation (and not the sequence)
 - How table is filled defines the agent
 - Different agents possible



Intro


Percept sequence	Action
[A, Clean]	Right
[A, Dirty]	Vacuum
[B, Clean]	Left
[B, Dirty]	Vacuum
[A, Clean], [A, Clean]	Right
[A, Clean], [A, Dirty]	Vacuum
...	...
[A, Clean], [A, Clean], [A, Clean]	Right
[A, Clean], [A, Clean], [A, Dirty]	Vacuum
...	...

Rationality

Rationality is not omniscience!


- Depends on four things:
 - *Performance measure*, defines criterion of success
 - Agent's prior *knowledge* of environment
 - *Actions* that agent can perform
 - Agent's *percept sequence* to date
- **Rational agent:**
 - For each possible percept sequence, a rational agent should select an *action*
 - expected to maximize its *performance measure*,
 - given evidence provided by *percept sequence* and
 - whatever built-in *knowledge* the agent has.

→ Rational = intelligent



Further Agent Properties

- Properties that agent behaviour can exhibit:
 - Rationality
 - Flexibility
 - Reactive
 - Proactive
 - Social (in multi-agent systems)
 - Autonomy
 - Mobility
 - Veracity
 - Benevolence
 - Ability to learn / adapt



Task Environment

- Task environments

- Essentially the “problems” to which rational agents are the “solutions”
- Specify using **PEAS** description
 - Performance measure
 - Environment
 - Actuators
 - Sensors

Attention: Name collision between
task environment und *environment*.

We will after this part usually only talk about the environment.

→ Properties of task environment

→ Modelling the environment

- Determine agent design

PEAS – Example

- Medical diagnosis system

Agent Type	Performance Measure	Environment	Actuators	Sensors
Medical diagnosis system	Healthy patient, reduced costs	Patient, hospital, staff	Display of questions, tests, diagnoses, treatments, referrals	Keyboard entry of symptoms, findings, patient's answers

- Software agent for *document retrieval, DR*
 - Sensors: input interface for textual queries
 - Actuators: output interface for documents

What makes up the *environment*? What is a fitting *performance measure*?

Properties of Task Environments

- Fully observable (vs. partially observable)
 - Agent's sensors give access to complete state of environment at each point in time
- Single agent (vs. multiple agents)
 - Single agent acts in environment
 - Can depend on modelling if something in environment is an object or another agent
 - Multi-agent systems:
 - Cooperative vs. competitive
 - Communication possible?
- Deterministic (vs. stochastic)
 - Next state of environment completely determined by current state and action executed by agent
 - Strategic: if environment deterministic except for actions of other agents

Properties of Task Environments

- **Episodic** (vs. **sequential**)
 - Agent's experience divided into atomic "episodes"
 - Episode = agent perceives a single percept and then performs a single action
 - Choice of action in each episode depends only on episode itself
- **Static** (vs. **dynamic**)
 - Environment unchanged while agent is deliberating
 - **Semi-dynamic**: environment itself does not change with passage of time but agent's performance score does
- **Discrete** (vs. **continuous**)
 - Applies to *state* of the environment, to the way *time* is handled, and to the *percepts* and *actions*
 - Finite number of distinct states, percepts, actions, time steps

Special property: **known** (vs. **unknown**)

- Agent's (or designer's) state of knowledge about its environment beforehand
- *No* property of the environment

Properties of Task Environments – Example

Task environment	Observable	Single agent	Deterministic	Episodic	Static	Discrete
Medical diagnosis	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
Self-driving taxi	Partially	Multi	Stochastic	Sequential	Dynamic	Continuous
Chess with a clock	Fully	Multi	Deterministic	Sequential	Semi	Discrete



How about image
classification?

Representation of the Environment


- Depends on properties of the task environment
 - Discrete/continuous, static/dynamic, episodic/sequential, ...
 - Determine expressiveness of representation
 - Increasing complexity and expressiveness
1. Atomic
 - Each state not further dividable (no internal structure)
 2. Factorised
 - Partitioning of state into fixed set of variables / attributes that can take a value
 - Representation: state-variable models, propositional logic, probabilistic graphical models (PGMs)
 3. Structured
 - Objects, relations among them
 - Representation: relational databases + PGMs, first-order logic


Advertisement:
Lecture „Automated Planning and Acting“

This lecture on StaRAI!

Performance Measure: Good Behaviour

- What is the right thing to do?
 - Agent generates sequence of actions given percepts
 - Causes environment to go through sequence of states
 - If sequence desirable / conforms to our expectations, then agent performed well
- **Performance measure** that evaluates any given sequence of **environment states**






Opponent = 0
Adversary (Adv1)

Ties = 0

Victim = 0
Normal (ZooV1)




Performance Measure

- Hard to determine
 - Not one fixed performance measure for all tasks and agents

As a general rule, it is better to design performance measures according to what one actually wants in the environment, rather than according to how one thinks the agent should behave.

Amount
of dirt?

Clean
locations?






Intro

Percept sequence	Action
[A, Clean]	Right
[A, Dirty]	Vacuum
[B, Clean]	Left
[B, Dirty]	Vacuum
[A, Clean], [A, Clean]	Right
[A, Clean], [A, Dirty]	Vacuum
...	...
[A, Clean], [A, Clean], [A, Clean]	Right
[A, Clean], [A, Clean], [A, Dirty]	Vacuum
...	...


Side Note: Current Research

- Idea:
 - **Humans**: intelligent to the extent that **our** actions can be expected to achieve **our** goals
 - ~~• **Machines**: intelligent to the extent that **their** actions can be expected to achieve **their** goals~~
 - **Machines** are **beneficial** to the extent that **their** actions can be expected to achieve **our** goals
- Approach: Performance measure unknown, human as assistant
- Goal: *Provably beneficial AI*
- See for example Stuart Russell




Implementing Agents

- **Agent = architecture + program**
 - Architecture
 - Entity with corresponding physical sensors and actuators
 - Entity = computer, robot, software (interfaces as sensors and actuators)
 - Program
 - Implementation of an agent function
 - Runs on an architecture
 - Both depend on task environment
- Implementation necessary, but how?




Agent Structure

- General agent structures for implementing agent functions:
 - Ordered by increasing generalisation
 1. Simple reflex agent
 2. Model-based reflex agent
 3. Goal-based agent
 4. Utility-based agent
 5. Learning agent
 - ❖ Human-aware agent (current research)




Agent Structure: Simple Reflex Agent

- Actions chosen based on current percept
 - Ignores previous percepts
 - No modelling of the environment
- Only correct decision on action if environment fully observable
 - If partially observable, infinite loops possible
 - (Partial) solution:
Choose random action




Agent Structure: Model-based Reflex Agent

- Given partial observability, keep track of not observable part
 - Using **internal state**
 - Depends on percept sequence
 - Encode not observable aspects of the current state
 - Encoding of environment atomic, factorised or structured
- Update internal state with information over
 - How environment evolves independently of agent
 - Effect of actions on environment




Agent Structure: Goal-based Agent

- Goal information useful
 - Description of desirable states
 - Infer from performance measure
 - Conditions for a goal state to fulfil
 - Example: vacuum cleaner
 $\forall x \in Loc : x = clean$
- Combine current state and goal information to choose actions that lead to goal
- Research areas:
 - Search
 - Planning



Agent Structure: Utility-based Agent


- Goal-based: binary distinction between *happy* and *unhappy*
- Utility as a distribution over possible states
 - What we look at later in the lecture
 - Essentially an internalisation of the performance measure
 - If internal utility function *agrees with* external performance measure:
 - Agent that chooses actions to maximize its utility will be *rational* according to the external performance measure



Agent Structure: Learning Agent


- So far: Agents select actions
 - Make decisions about actions
- How do agent programs emerge?
 - Too laborious to do by hand
- Generate learning agent, let it learn
 - Allows agent to operate in initially unknown environments and to become more competent
 - Research area:
Reinforcement Learning

Advertisement: Lecture „Deep Reinforcement Learning“ (PI, 4) by Malte Schilling




Agent Structure: Learning Agent

- 4 conceptual components
 - Learning element
 - Make improvements
 - Performance element (agent so far)
 - Select actions
 - Critic
 - Feedback on results
 - Used by learning element to determine how performance element should be modified to do better
 - Problem generator (for exploration)
 - Suggest actions that will lead to new and informative experiences




Side Note: Current Research

- *Human-aware agent*
 - Interaction with a human during acting (*human-in-the-loop*)
 - Example: *Urban Search and Rescue* with a robot
- Different models considered
 - Original agent model $\mathcal{M}^R$
 - Human model $\mathcal{M}^H$
 - Agent model $\tilde{\mathcal{M}}_r^H$ of human model $\mathcal{M}^H$
 - Agent model $\tilde{\mathcal{M}}_h^R$ of human model $\mathcal{M}_h^R$, which human has of $\mathcal{M}^R$



Side Note: Current Research

- Different models considered:
 - Original agent model $\mathcal{M}^R$
 - Human model $\mathcal{M}^H$
 - Agent model $\tilde{\mathcal{M}}_r^H$ of human model $\mathcal{M}^H$
 - Agent model $\tilde{\mathcal{M}}_h^R$ of human model $\mathcal{M}_h^R$, which human has of $\mathcal{M}^R$
- $\tilde{\mathcal{M}}_r^H$ allows for anticipating human behaviour
- $\tilde{\mathcal{M}}_h^R$ allows for conforming to human expectations
- See, e.g., Subbarao (Rao) Kambhampati
<http://rakaposhi.eas.asu.edu>



Interim Summary

- Agent = architecture + agent program
- Task environment
 - Description: PEAS
 - Performance measure not easy to formalise
 - Properties of task environment
 - Fully observable (vs. partially observable); single agent (vs. multiple agents); deterministic (vs. stochastic), strategic; episodic (vs. sequential); static (vs. dynamic), semi-dynamic; discrete (vs. continuous); special: known (vs. unknown)
 - Environment encoding
 - Atomic, factorised, structured
- Agent structure
 - Simple / model-based reflex agent, goal- / utility-based agent; learning agent

Overview: 1. Introduction

a. Artificial Intelligence

- Approaches: thinking / acting humanly / rationally


b. Framework: Agent Theory

- Agent
- Task environment
- Agent structure


c. Topic: StaRAI

- Motivation, context
- Relational examples, outlook on probabilistic relational models (PRMs)

Connection to AI, Agents, and Environments






Connection to AI, Agents, and Environments



Probabilistic Graphical Models (PGMs)


- *Factorisation* of a full joint according to (conditional) independences in the full joint




<i>Epid</i>	<i>Travel</i>	<i>Sick</i>	<i>P</i>
false	false	false	0.20
false	false	true	0.24
false	true	false	0.28
false	true	true	0.08
true	false	false	0.05
true	false	true	0.06
true	true	false	0.07
true	true	true	0.02

Full joint probability distribution


Problem: Adding Objects and Relations in Propositional Formalisms Clunky...



Problem: ... and Models Explode with Them




Propositional → First-order View



Symmetries in graph / relations in scenario

Connection to AI, Agents, and Environments



From Logics to Probabilistic Relational Models

- Propositional: Descriptions about world
 - Form of constraints on an environment

Presents $\Rightarrow$ Attends

Presents(eve, paper1, IJCAI) $\Rightarrow$ Attends(eve, IJCAI)

- First-order: Objects, relations among them
 - Groundings to get to propositional logic

Presents(X, P, C) $\Rightarrow$ Attends(X, C)

- Denote either true or false statements

- Approach to soften constraints: Introduce weights to denote that statements hold to a certain degree over all possible worlds
 - Example: Markov Logic Network
 - We will see more of them later


Hard constraint

∞ *Presents(X, P, C) $\Rightarrow$ Attends(X, C)*

3.75 *Publishes(X, C) $\wedge$ FarAway(C) $\Rightarrow$ Attends(X, C)*


Soft constraint,
weight = **exp(3.75)**

Connection to AI, Agents, and Environments



Probabilistic Relational Models (PRMs)

- Random variables for combinations of individuals in populations
 - Build a probabilistic model before knowing (all of) the individuals
 - Learn the model for one set of individuals
 - Apply the model to existing and new individuals
 - Allow complex relationships between individuals
- Exchangeability:
 - Before we know anything about individuals, they are *indistinguishable*, and so should be treated identically.



$\infty \text{ Presents}(X, P, C) \Rightarrow \text{Attends}(X, C)$

3.75 $\text{Publishes}(X, C) \wedge \text{FarAway}(C) \Rightarrow \text{Attends}(X, C)$

- Uncertainty about:
 - Properties of individuals
 - Relationships among individuals
 - Identity (equality) of individuals
 - Existence (and number) of individuals
- Depicted formalisms: Parfactor graphs, Markov logic networks

PRMs for Modelling the Environment in this Lecture

- Different types of PGMs for a (relational) factorised representation of the environment and decision making
- Possible to model environment with following properties
 - Fully or partially observable
 - Single agent
 - Stochastic
 - Episodic or sequential
 - Static
 - Discrete or continuous
- Not considered in this lecture
 - Multiple agents
 - Deterministic, strategic
 - Dynamic
- Approaches exist to deal with such environments to a certain degree

Attention: In the PGM literature, the notions of *static* and *dynamic* are used instead of the notions of *episodic* and *sequential*, while the notions of *static* and *dynamic* are covered by the so-called *Markov assumption* or *(Non-) Markovian abstraction*

Interim Summary

- Environment characterised by
 - Objects and relations between them
 - Uncertainty
- PRMs combine both logic and probability theory
 - Models covered here can model the following environment properties
 - Fully or partially observable
 - Single agent
 - Stochastic
 - Episodic or sequential
 - Static
 - Discrete or continuous

Contents

1. Introduction

- Notions of AI, agents, StaRAI

2. Modelling formalisms

- Logic*
 - Probabilistic graphical models (PGMs)*
 - Probabilistic relational models (PRMs)*
- Syntax, semantics

3. Inference

- Logical inference*
 - Probabilistic inference*
 - Exact inference
 - Approximate inference
- Inference tasks, algorithms

4. Learning

- Relation learning in logic
- Parameter / structure learning for PGMs
- From PGMs to PRMs, PRM learning

5. Sequential Models and Inference


- Syntax, semantics, inference tasks, algorithms

6. Decision-theoretic Models and Inference

- Preferences, utility
- Syntax, semantics, inference task, algorithm

Contents in this Lecture Related to *Utility-based Agents*

- Further topics
 2. Modelling formalisms
 3. Inference (in episodic models)
 4. Learning (of episodic models)
 5. Sequential models and inference
 6. Decision-theoretic models and inference



Overview: 1. Introduction

a. Artificial Intelligence (AI)

- Approaches: thinking / acting humanly / rationally

b. Framework: Agent Theory

- Agent
- Task environment
- Agent structure

c. Topic: Statistical Relational AI

- Motivation, context
- Relational examples, outlook on probabilistic relational models (PRMs)

→ Modelling Formalisms