

MSc Knowledge Engineering: A List of Topics

Michael Rovatsos

March 17, 2005

Introduction

- Definition and types of knowledge
- What are Knowledge-Based Systems? What is Knowledge Engineering?
- The Knowledge Engineering process
- The human interface
- Critique of KE

Knowledge Acquisition

Inductive Learning

- Definition, what are hypothesis, target concepts, hypothesis spaces
- How are IL methods described?
- Ockham's razor, notions of consistency, realisability, noise overfitting

Decision Tree Learning

- What are DTs? How expressive are they?
- The Decision Tree Learning algorithm
- Attribute selection heuristics
- Validation techniques

Version-Space Learning

- Candidate definitions, describing hypotheses with logical formulae
- Current-best hypothesis search
- Version-Space Learning algorithm
- Updating the version space

Knowledge Representation & Reasoning

Logic Recap

- Propositional logic: truth tables, inference rules, resolution
- First-order logic: predicates, quantifiers, substitution, unification, resolution (idea)

Ontologies

Basics

- Definition, categories, upper ontologies, (multiple) inheritance
- Physical composition (part-of relation, logical minimisation)
- Measurements (quantitative vs. qualitative)
- Substances and objects (individuation problem)
- Expressing change (situation calculus, fluents, the frame problem(s), successor-state axioms)

Category Reasoning Systems

- Semantic Networks (inheritance, relations, problems of binary relations, reification, default reasoning by overriding, shortcomings)
- Description Logics (reasoning about categories with simple logics)

Reasoning with Default Information

- Closed-world assumption/unique names assumption, negation as failure
- Non-monotonic reasoning
- Circumscription (model preference, prioritised circumscription)
- Default logic (default rules, extensions)

Model-Based Reasoning

- A case study in KR&R
- What is MBR?
- The General Diagnostic Engine
- Minimal candidates, candidate discrimination
- Introducing explicit fault models

Reasoning with Uncertainty

- Different kinds of uncertainty
- Probabilistic reasoning (Bayes' rule, belief networks)
- Fuzzy logic (characteristic functions, truth-functional approach, rules for combining fuzzy values, defuzzification)
- Dempster-Shafer Theory (uncertainty vs. ignorance, combining evidence, interval view of degrees of belief)

Knowledge Synthesis

The Amphion System

- Automated software synthesis
- Specification acquisition, program synthesis, domain-specific subsystems
- Deductive synthesis approach

Agents and Multiagent Systems

Basics

- Open vs. closed systems
- Definition of agent, properties of intelligent agents
- The autonomy debate, rationality vs. reactivity (intentional systems, social ability (typology of interaction))
- Definition of multiagent systems
- Research agenda of agent and multiagent systems areas, sub-areas, the programming perspective

Agent Architectures

- Symbolic AI and its problems
- Dispute between behaviour-based and deliberative views of agency
- The BDI architecture (practical reasoning, deliberation and means-ends reasoning, intentions and their properties, issues in BDI)
- Reactive architectures (assumptions, the subsumption architecture, Mars rover example)
- Hybrid architectures (layered approaches, Touring Machines, InteRRaP)

Agent Interaction & Communication

- Categories of agent interaction
- Speech Act Theory (locution, illocution, perlocution, performatives, propositional content)
- Agent Communication Languages (KQML/KIF, FIPA-ACL)
- ACL Semantics (mentalistic semantics, social commitment-based semantics)
- Interaction protocols (protocol design, examples, the contract-net protocol)

Distributed Rational Decision Making

- Decision theory (preferences, expected utility maximisation)
- Game Theory (basics, normal-form games, dominant strategy (equilibrium), best response strategies, Nash Equilibrium, The Prisoner's Dilemma, the evolution of cooperation)
- Mechanism design (criteria: individual rationality, stability, Pareto efficiency, computational efficiency, distribution properties)
- The Revelation Principle (proof!)

- Electronic auctions (English, Dutch/First-Price Sealed Bid, Vickrey Auction, properties of each of them, winner's curse)
- Other methods for distributed rational decision making, critique of game-theoretic approaches

Knowledge Engineering & The Semantic Web

- The Web today, the vision of the Semantic Web
- Semantic Web technologies, the layer cake
- XML, DTDs/XML Schema
- RDF (resources, properties, statements) and RDF Schema (simple lightweight ontologies, semantics)
- OWL (expressiveness, different flavours, shortcomings)
- Critique of the Semantic Web

Knowledge Evolution

Truth Maintenance Systems

- JTMS
- ATMS

Knowledge in Learning

- The knowledge-based inductive learning problem (entailment constraints)
- Explanation-based learning (generalising existing knowledge to cover new situations, procedure, parallel proofs)
- Inductive logic programming (expressiveness, constructive induction, top-down methods, inverse resolution methods, making discoveries with ILP)