Review of Knowledge Management and Assistance Systems

WMA-1	Introduction
WMA-2	Semantic Networks
WMA-3	Description Logics
WMA-4	Logics of Image Interpretation
WMA-5	Rule-based Systems
WMA-6	Configuration Systems
WMA-7	Case-based Problem Solving
WMA-8	Supervised Concept Formation
WMA-9	Data Mining and Clustering
WMA-10	Decision Trees and Neural Networks
WMA-11	Semantic Web, XML and RDF
WMA-12	Ontologies and OWL
WMA-13	Review

1

What is Knowledge?

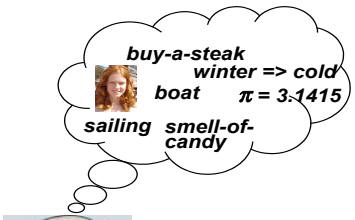

information and skills acquired by education and experience

π = 3,14159 26535 89793 23846

IF winter
THEN cold

marietta.jpg

http://best-steakhouse.com

information and processing methods acquired by programming and machine learning

Characteristics of Knowledge-Based Assistance Systems

- **Relevant knowledge about application domain is represented in a declarative format (as opposed to a procedural format)**
  - enhances readability
  - facilitates change maintenance
- **Domain knowledge and problem-specific knowledge may be separated**
- **Inference services may have general validity and proven correctness**
  - validity of logic-based inferences is well-understood
  - validity of rule-based and handcrafted inferences must be doubted
- **Separation of data and control**
  - enables data-driven processing
  - not cleanly realized in rule-based systems
- **Domain knowledge must be acquired and modelled**
  - "knowledge-acquisition bottleneck"

3

Basic Relations in Semantic Networks

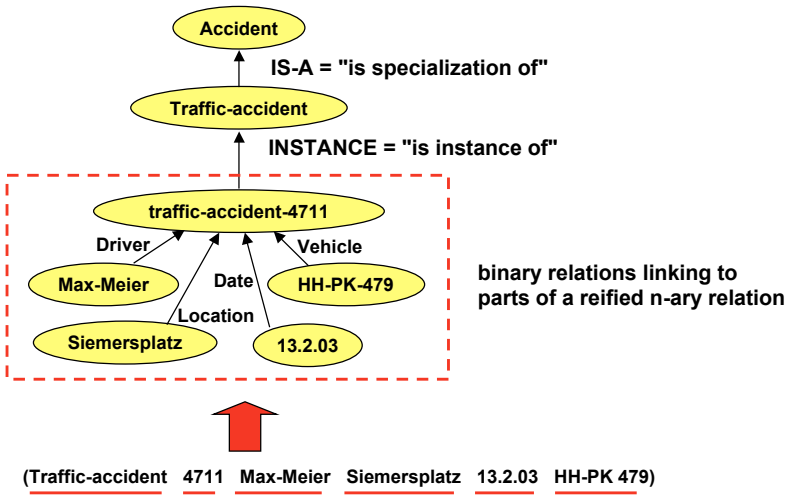

4

Summary of Semantic Networks

- **Intuitive graphical knowledge representation formalism with nodes representing concepts and individuals, and links representing relations**
- **Semantics of relations is well-defined for ISA and INSTANCE, but not clearly defined in general.**
- **Relations between relations cannot be expressed.**
- **The notion of an object and of object properties is not explicitely supported.**
- **Some services (basic information retrieval, basic classification) can be supported by pattern matching.**
- **Generally useful services require additional formalisms such as rules and rule-based inferences, e.g. for axiomatizing domains.**

5

Description Logics for Knowledge Representation

- **DLs are a family of knowledge-representation formalisms**
- **Decidable subset of FOL**
- **Object-centered, roles and features (binary relations)**
- **Necessary vs. sufficient attributes**
- **Inference services**
  - subsumption check
  - consistency check
  - classification
  - abstraction
  - default reasoning
  - spatial and temporal reasoning
- **Guaranteed correctness, completeness, and decidability properties**
- **Highly optimized implementations (e.g. RACER)**
- **Provides inferences for Semantic Web language OWL**

Description Logics vs. Semantic Networks

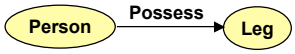

Unclear semantics:

A Specification of person concept restricted by role possess?
B Specification of role possess restricted by domain and range concepts?
C Specification of leg concept restricted by role possess?

+ in all cases unclear cardinalities

With DLs:

(and person (some possess leg))
"objects which are persons and each of which possesses at least one leg"

(and person (all possess leg))
"objects which are persons and each of which only possesses legs"

Each DL concept expression describes a set of objects!

7

Description Logics for Image Interpretation

Using description logics for
- **knowledge representation (visual phenomena, background knowledge)**
- **inferences (meaning assignment, interpretation)**

- **Important application as web service**
  - automatic annotation of images
  - content-based retrieval
  - multimedia content services
- **Current research topics**
  - logics of multimedia interpretation?
  - standard inference services?
  - part of Semantic Web?

Ucraine´s Andrey Sokolovskiy clears 2,38 in Rome

8

Rule-based Expert Systems

Developed 1970 - 1985 to

- **collect and preserve expert knowledge**
- **replace human experts by computer programs**
- **to automatically derive interesting knowledge.**

Basic idea: Represent expert knowledge in terms of IF-THEN rules

Basic structure:

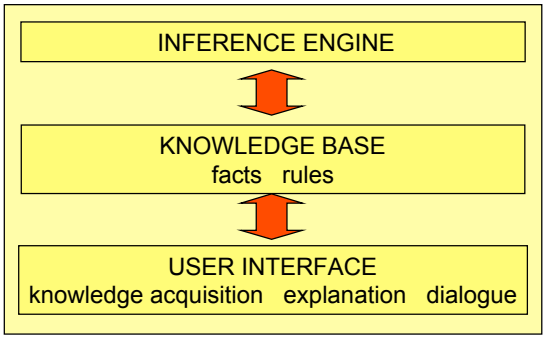

9

Summary of Rule-based Expert Systems

- **Elegant and easy-to-use tool for knowledge-based systems**
  - Intuitive knowledge representation
  - General-purpose inference engine
  - Data-driven control
- **Rule-based expert systems provided first commercial success for AI systems but also first hype and disappointments**
  Prevailing prejudice: AI means "using rules"
- **Knowledge acquisition from human experts proved difficult**
  Creation of the new profession of a "Knowledge Engineer"
- **Systems do not scale up easily, need additional organisation**
  - Programmed control
  - Structure for large numbers of rules
- **Expert-system frameworks lacked integration**

Today´s use of rule-based systems mostly as embedded components

10

Representation of a Configuration Model

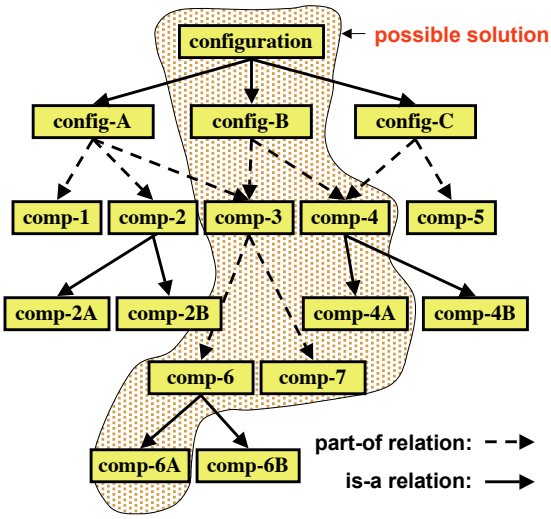

- boxes (frames) specify aggregate and component properties
- has-part relations bind components to aggregates
- is-a relations describe variants of entities
- constraints between entities (not shown) restrict choices and parameter combinations

11

Configuration as Logical Model Construction

Instantiate the configuration model consistent with task description.

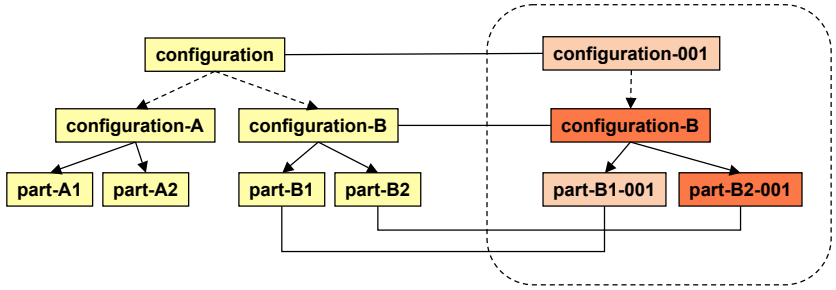

configuration model

constructed logical model involving task description

12

Exploiting Knowledge by Means of of "Cases"

- **Case-based problem solving**
  Last time the light failed I replaced the bulb
- **Supervised learning of concepts**
  This is a positive example for a market where one should expand business
  In this situation one should not expand business
- **Unsupervised discovery of conceptual relations**
  People who buy beer, often also buy chips

"Case-based Reasoning" (CBR) is the discipline in Computer Science dealing with such problems.

13

Case-based Problem Solving

Solving problems with the help of experiences stored as a case database.

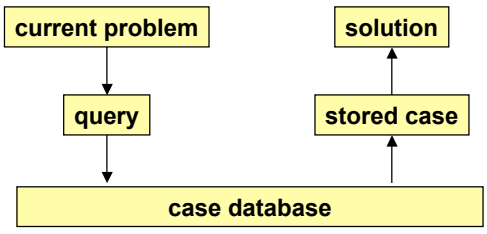

Used by humans for almost all kinds of problem solving:

- **Diagnosis**
- **Design and Construction**
- **Planning**
- **Prediction**
- **Explanation**

How can this work?
Heraklit: Παντα ρει

14

Case Study: Case-based Assistence for Fault Diagnosis in a Flexible Manufactoring System

Study for Deutsche Airbus: Conceptual design for case representation, reasoning and user interface.

Main user modes:

- **Failure handling**
  - Decribing a failure
  - Diagnosing a failure
- **Browsing**
  - Inspection of case database
  - Statistical evaluation
- **Case database administration**
  - Entering new cases and editing cases
  - Ontology extensions

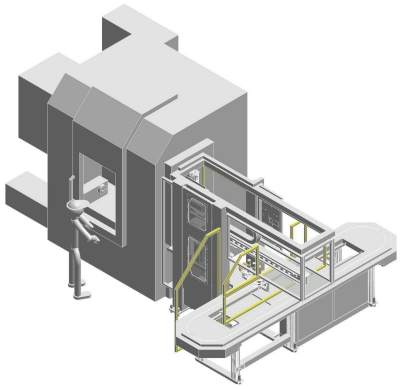

15

Supervised Learning of Concepts from Cases

Given a set of positive and negative examples for an unknown class, what is a conceptual description for that class?

- **Special kind of case-based reasoning (CBR)**

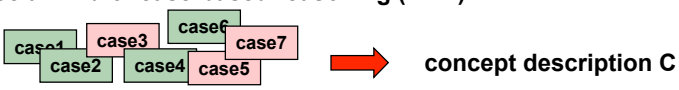

- **Special kind of case-based problem solving where the problem is to determine class membership in a binary classification task.**

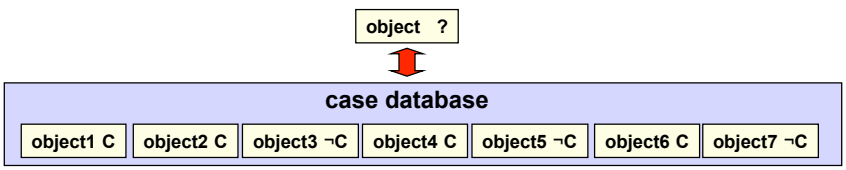

16

Example: Learning to Classify Mushrooms

Learn from positive and negative examples to distinguish poisonous and nonpoisonous mushrooms.

Mushroom description:

Colour	{Red, Grey}
Size	{Small, Large}
Shape	{rOund, Elongated}
Environment	{Humid, Dry}
Height	{loW, hIgh}
Texture	{sMooth, roUgh}
Class	{Poisonous, Nonpoisonous}

Note simple attribute language for the sake of an easy example. VSL can deal with much richer languages.

17

Summary of Version Space Learning

- VSL is a logic-based method for determining all possible concept descriptions based on positive and negative examples.
- Concepts must be described in a description language which allows to establish a generalization hierarchy.
- VSL may be extended for structured objects.
- The version space may collapse if erroneous examples are introduced.
- Probabilistic learning models compete.

18

Data Mining

Data mining means analyzing large data sets to identify and establish new patterns or relationships which in the end prove valid, comprehensible and useful.

Classical Statistics
Machine Learning
Pattern Recognition

} **Data Mining und Knowledge Discovery (DMKD) becomes part of AI (1990)**

Data Mining:
Analysis, generation of hypotheses

Knowledge Discovery:
Evaluation and interpretation of hypotheses

Sources: Görz et al. (Eds.): Handbuch der künstlichen Intelligenz (3. Aufl.), Oldenbourg, 2000
Maimon & Rokach (Eds.): The Data Mining and Knowledge Discovery Handbook, Springer 2005

19

Requirements for Association Rules

Determine rules X -> Y in transactions T where:

$$s(r) := \frac{\left|\left\{t \in T \mid X \cup Y \in t\right\}\right|}{|T|} \geq s_{min}$$ **minimal support**

$$c(r) := \frac{\left|\left\{t \in T \mid X \cup Y \in t\right\}\right|}{\left|\left\{t \in T \mid X \in t\right\}\right|} \geq c_{min}$$ **minimal confidence**

Typical values: $s_{min} = 0{,}01$ $c_{min} = 0{,}5$

{beer, pizza} -> {chips} will be established if
- at least 1% of all customers have bought beer, pizza, and chips, and
- at least 50% of the beer and pizza customers have also bought chips.

20

Discovering Clusters in Data

"Understanding our world requires conceptualizing the similarities and differences between the entities that compose it." Tyron & Bailey 1970

"Clusters" of data objects are hypothetical classes based on similarities and distances. Data objects should be as similar as possible within clusters and as distinct as possible between clusters.

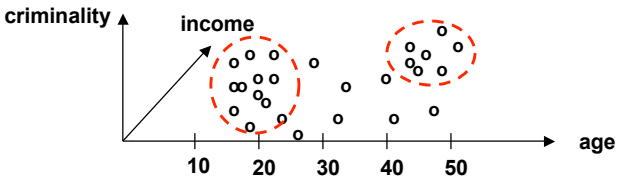

Cluster 1: age 15 - 25, low income, high criminality ("youth criminality")
Cluster 2: age 45 - 55, high income, high criminality ("white-collar criminality")

Data objects are viewed as points in a multi-dimensional feature space. Similarity of data is judged by distance measures.

21

K-means Clustering

- Most popular clustering algorithm
- Searches for local minimum of sum of Euclidean sample distances to cluster centers
- Guaranteed convergence in a finite number of steps
- Requires initialization of fixed number of clusters k
- May converge to local minimum

A	Initialize cluster centers
B	Assign data objects to nearest cluster centers
C	New cluster centers are the mean of assigned data objects
D	Repeat steps B to D until no more changes occur

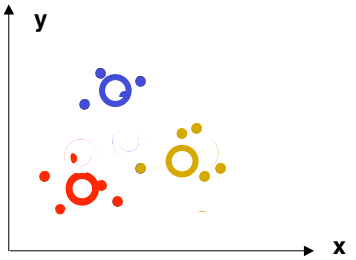

Example:
K-means clustering with k = 3

22

Decision Trees

Decision trees are a popular method for classifying objects by means of a sequence of tests of feature values following a tree structure.

A decision tree is a tree where:
- the non-leaf nodes are labeled with attributes,
- all arcs out of a node labeled with attribute A are labelled with each of the possible values of the attribute A,
- the leafs of the tree are labeled with classifications

Example :

Decision tree for classifying a person as reading or skipping a book based on several attributes:
- author known or unknown
- thread new or old
- length long or short
- place at home or at work

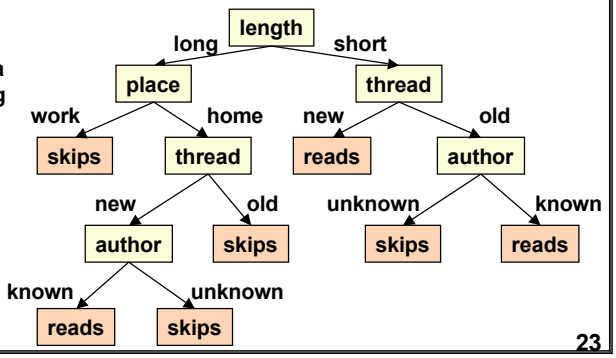

23

Maximizing Information Gain

Select attributes in the order of maximal information gain.

H(G)	entropy of source regarding goal attribute G with distribution according to example set
$H(G\|A=a_i)$	entropy of same source based on subset of examples where attribute A has value a_i
q_i	fraction of example set where attribute A has value a_i
IG	information gain by asking for the attribute value of A

$$IG = H(G) - \sum_i q_i H(G \mid A = a_i)$$

24

Typical Applications for Neural Networks

NNs are cash cows (Goldesel) for engineers:

Feed examples and obtain classifier!

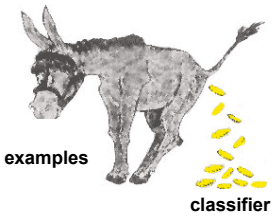

Useful primarily for applications which are difficult to analyze for humans:

- Speech recognition, e.g. determining the identity of a speaker
- Lipreading
- Image understanding, e.g. classifying x-rayed luggage as suspicious
- Event recognition, e.g. dangerous patterns in air traffic
- Predict which job an applicant is best suited for
- Diagnose diseases

25

Semantic Web Layer Cake

(Term proposed by Berners-Lee)

Layers are developed bottom-up and standardized by "recommendations" of the World Wide Web Consortium (W3C)

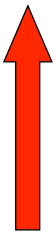

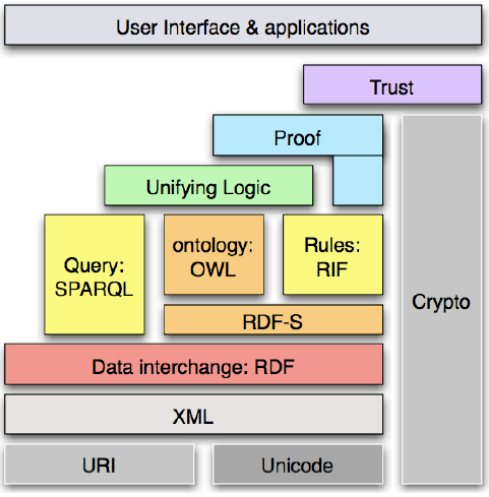

26

What is an Ontology?

In Greek philosophy:

Ontology is the study of being and existence, of the nature of reality

In Computer Science:

An ontology is a formal explicit specification of a shared conceptualization of a domain.

A specification consisting of
- **classes**
- **relations between classes**
- **individuals**
- **axioms.**

"Ontology" is often misused for a taxonomy of concept names without specification of formal relations between classes.

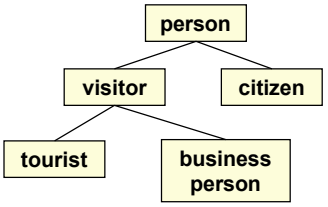

27

OWL Inference Services

- **No specific recommendations of W3C regarding inference services**
- **General agreement about important services:**

Determine
- **class equivalence**
- **subclass relationship**
- **disjunctiveness**
- **global consistency (satisfiability)**
- **class consistency**

A class is inconsistent, if it is equivalent to owl:Nothing

Example for obviously inconsistent class:

```
<owl:Class rdf:about="#book">
    <owl:subClassOf rdf:resource="#publication"/>
    <owl:disjointWith rdf:resource="#publication"/>
</owl:Class>
```

28

Thank you for your attention!

And good luck with your examination!

29