Knowledge Processing

Franz J. Kurfess

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

Acknowledgements

*Some of the material in these slides was developed for a lecture series sponsored by the **European Community** under the **BPD program** with **Vilnius University** as host institution*

Use and Distribution of these Slides

These slides are primarily intended for the students in classes I teach. In some cases, I only make PDF versions publicly available. If you would like to get a copy of the originals (Apple KeyNote or Microsoft PowerPoint), please contact me via email at fkurfess@calpoly.edu. I hereby grant permission to use them in educational settings. If you do so, it would be nice to send me an email about it. If you're considering using them in a commercial environment, please contact me first.

Overview Knowledge Processing

- Motivation
- Objectives
- Chapter Introduction
  - Knowledge Processing as Core AI Paradigm
  - Relationship to KM
  - Terminology
- Knowledge Acquisition
  - Knowledge Elicitation
  - Machine Learning
- Knowledge Representation
  - Logic
  - Rules
  - Semantic Networks
  - Frames, Scripts
- Knowledge Manipulation
  - Reasoning
  - KQML
- Important Concepts and Terms
- Chapter Summary

Motivation

- the representation and manipulation of knowledge has been essential for the development of humanity as we know it
- the use of formal methods and support from machines can improve our knowledge representation and reasoning abilities
- intelligent reasoning is a very complex phenomenon, and may have to be described in a variety of ways
- a basic understanding of knowledge representation and reasoning is important for the

Objectives

- be familiar with the commonly used knowledge representation and reasoning methods
- understand different roles and perspectives of knowledge representation and reasoning methods
- examine the suitability of knowledge representations for specific tasks
- evaluate the representation methods and reasoning mechanisms employed in computer-based systems

Chapter Introduction

- Knowledge Processing as Core AI Paradigm
- Relationship to KM
- Terminology

Relationship to KM

KP/AI	KM
representation methods suited for KP by computers	representation of knowledge in formats suitable for humans
reasoning performed by computers	essential reasoning performed by humans
mostly limited to symbol manipulation	support from computers
very demanding in terms of computational power	emphasis often on documents
can be used for “grounded” systems	larger granularity
interpretation (“meaning”) typically left to humans	mainly intended for human use

Knowledge Processes

Chaotic knowledge processes

Human knowledge and networking

Information databases and technical networking

Systematic information and knowledge processes

[Skyrme 1998]

Knowledge Cycles

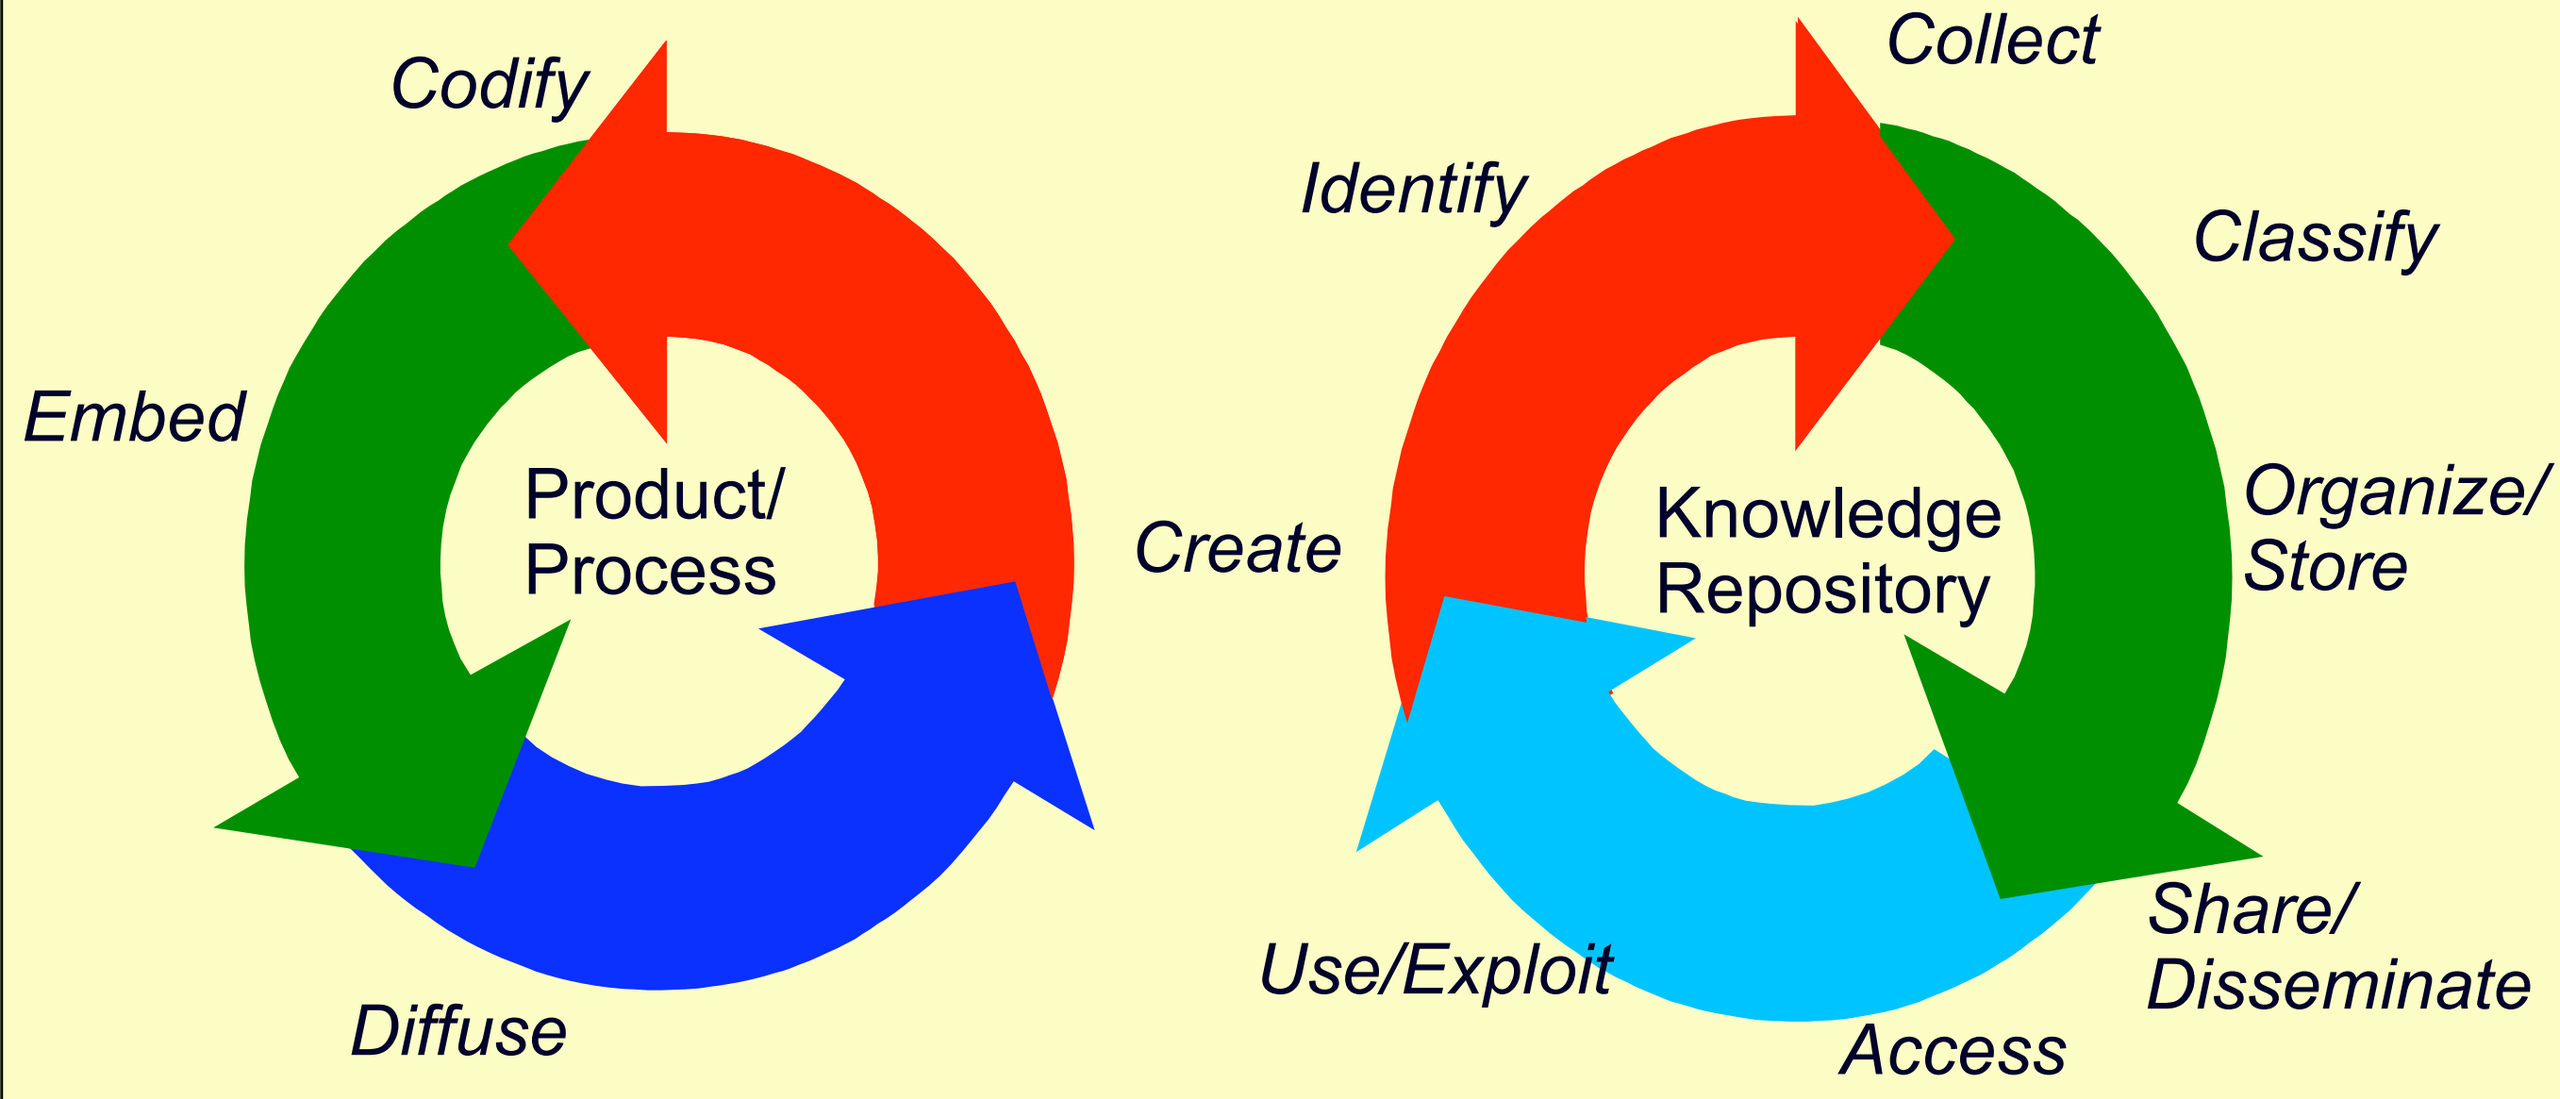

[Skyrme 1998]

Knowledge Representation

- Types of Knowledge
  - Factual Knowledge
  - Subjective Knowledge
  - Heuristic Knowledge
  - Deep and Shallow Knowledge
- Knowledge Representation Methods
  - Rules, Frames, Semantic Networks
  - Blackboard Representations
  - Object-based Representations
  - Case-Based Reasoning
- Knowledge Representation Tools

Roles of Knowledge Representation

- Surrogate
- Ontological Commitments
- Fragmentary Theory of Intelligent Reasoning
- Medium for Computation
- Medium for Human Expression

[Davis, Shrobe, Szolovits, 1993]

KR as Surrogate

- a substitute for the thing itself
- enables an entity to determine consequences by thinking rather than acting
  - reasoning about the world through operations on the representation
- reasoning or thinking are inherently *internal* processes
- the *objects* of reasoning are mostly *external* entities ("things")
  - some objects of reasoning are internal, e.g. concepts, feelings, ...

[Davis, Shrobe, Szolovits, 1993]

Surrogate Aspects

- Identity
  - correspondence between the surrogate and the intended referent in the real world
- Fidelity
  - Incompleteness
  - Incorrectness
  - Adequacy
    - Task
    - User

[Davis, Shrobe, Szolovits, 1993]

Surrogate Consequences

- perfect representation is impossible
  - the only completely accurate representation of an object is the object itself
- incorrect reasoning is inevitable
  - if there are some flaws in the world model, even a perfectly sound reasoning mechanism will come to incorrect conclusions

[Davis, Shrobe, Szolovits, 1993]

Ontological Commitments

- terms used to represent the world
- by selecting a representation a decision is made about how and what to see in the world
  - like a set of glasses that offer a sharp focus on part of the world, at the expense of blurring other parts
  - necessary because of the inevitable imperfections of representations
  - useful to concentrate on relevant aspects
  - pragmatic because of feasibility constraints

Ontological Commitments Examples

- logic
  - views the world in terms of individual entities and relationships between the entities
- rules
  - entities and their relationships expressed through rules
- frames
  - prototypical objects
- semantic nets
  - entities and relationships

[Davis, Shrobe, Szolovits, 1993]

KR and Reasoning

- a knowledge representation indicates an initial conception of intelligent inference
  - often reasoning methods are associated with representation technique
    - first order predicate logic and deduction
    - rules and modus ponens
- the association is often implicit
- the underlying inference theory is fragmentary
  - the representation covers only parts of the association
  - intelligent reasoning is a complex and multi-faceted phenomenon

[Davis, Shrobe, Szolovits, 1993]

KR for Reasoning

- a representation suggests answers to fundamental questions concerning reasoning:
  - What does it mean to reason intelligently?
    - implied reasoning method
  - What can possibly be inferred from what we know?
    - possible conclusions
  - What should be inferred from what we know?
    - recommended conclusions

KR and Computation

- from the AI perspective, reasoning is a computational process
  - machines are used as reasoning tools
- without efficient ways of implementing such computational process, it is practically useless
  - e.g. Turing machine
- most representation and reasoning mechanisms are modified for efficient computation
  - e.g. Prolog vs. predicate logic

[Davis, Shrobe, Szolovits, 1993]

Computational Medium

- computational environment for the reasoning process
- reasonably efficient
- organization and representation of knowledge so that reasoning is facilitated
- may come at the expense of understandability by humans
  - unexpected outcomes of the reasoning process
  - lack of transparency of the reasoning process
    - even though the outcome “makes sense”, it is unclear how it was achieved

KR for Human Expression

- a knowledge representation or expression method that can be used by humans to make statements about the world
  - expression of knowledge
    - expressiveness, generality, preciseness
  - communication of knowledge
    - among humans
    - between humans and machines
    - among machines
  - typically based on natural language
  - often at the expense of efficient computability

Knowledge Acquisition

- Knowledge Elicitation
- Machine Learning

Acquisition of Knowledge

- Published Sources
  - Physical Media
  - Digital Media
- People as Sources
  - Interviews
  - Questionnaires
  - Formal Techniques
  - Observation Techniques
- Knowledge Acquisition Tools
  - automatic
  - interactive

Knowledge Elicitation

- knowledge is already present in humans, but needs to be converted into a form suitable for computer use
- requires the collaboration between a domain expert and a knowledge engineer
  - domain expert has the domain knowledge, but not necessarily the skills to convert it into computer-usable form
  - knowledge engineer assists with this conversion
  - this can be a very lengthy, cumbersome and error-prone process

Machine Learning

- extraction of higher-level information from raw data
- based on statistical methods
- results are not necessarily in a format that is easy for humans to use
- the organization of the gained knowledge is often far from intuitive for humans
- examples
  - decision trees
  - rule extraction from neural networks

Knowledge Fusion

- integration of human-generated and machine-generated knowledge
  - sometimes also used to indicate the integration of knowledge from different sources, or in different formats
- can be both conceptually and technically very difficult
  - different “spirit” of the knowledge representation used
  - different terminology
  - different categorization criteria
  - different representation and processing mechanisms
    - e.g. graph-oriented vs. rules vs. data base-oriented

Knowledge Representation Mechanisms

- Logic
- Rules
- Semantic Networks
- Frames, Scripts

Logic

- syntax: well-formed formula
  - a formula or sentence often expresses a fact or a statement
- semantics: interpretation of the formula
  - "meaning" is associated with formulae
  - often compositional semantics
- axioms as basic assumptions
  - generally accepted within the domain
- inference rules for deriving new formulae from existing ones

KR Roles and Logic

- surrogate
  - very expressive, not very suitable for many types of knowledge
- ontological commitments
  - objects, relationships, terms, logic operators
- fragmentary theory of intelligent reasoning
  - deduction, other logical calculi
- medium for computation
  - yes, but not very efficient
- medium for human expression

Rules

- syntax: **if … then …**
- semantics: interpretation of rules
  - usually reasonably understandable
- initial rules and facts
  - often capture basic assumptions and provide initial conditions
- generation of new facts, application to existing rules
  - forward reasoning: starting from known facts
  - backward reasoning: starting from a hypothesis

KR Roles and Rules

- surrogate
  - reasonably expressive, suitable for some types of knowledge
- ontological commitments
  - objects, rules, facts
- fragmentary theory of intelligent reasoning
  - modus ponens, matching, sometimes augmented by probabilistic mechanisms
- medium for computation
  - reasonably efficient

Semantic Networks

- syntax: graphs, possibly with some restrictions and enhancements
- semantics: interpretation of the graphs
- initial state of the graph
- propagation of activity, inferences based on link types

KR Roles and Semantic Nets

- surrogate
  - limited to reasonably expressiveness, suitable for some types of knowledge
- ontological commitments
  - nodes (objects, concepts), links (relations)
- fragmentary theory of intelligent reasoning
  - conclusions based on properties of objects and their relationships with other objects
- medium for computation
  - reasonably efficient for some types of reasoning

Frames, Scripts

- syntax: templates with slots and fillers
- semantics: interpretation of the slots/filler values
- initial values for slots in frames
- complex matching of related frames

KR Roles and Frames

- surrogate
  - suitable for well-structured knowledge
- ontological commitments
  - templates, situations, properties, methods
- fragmentary theory of intelligent reasoning
  - conclusions are based on relationships between frames
- medium for computation
  - ok for some problem types
- medium for human expression

Knowledge Manipulation

- Reasoning
- KQML

Reasoning

- generation of new knowledge items from existing ones
- frequently identified with *logical* reasoning
  - strong formal foundation
  - very restricted methods for generating conclusions
- sometimes expanded to capture various ways to draw conclusions based on methods employed by humans
- requires a formal specification or implementation to be used with computers

KQML

- stands for Knowledge Query and Manipulation Language
- language and protocol for exchanging information and knowledge

KQML Performatives

- basic query performatives
  - evaluate, ask-if, ask-about, ask-one, ask-all
- multi-response query performatives
  - stream-about, stream-all
- response performatives
  - reply, sorry
- generic informational performatives
  - tell, achieve, deny, untell, unachieve
- generator performatives
  - standby, ready, next, rest, discard, generator
- capability-definition performatives
  - advertise, subscribe, monitor, import, export

KQML Example 1

- query

```
(ask-if
      :sender A
      :receiver B
      :language Prolog
      :ontology foo
      :reply-with id1
      :content
``bar(a,b)'' )
```

- reply

```
(sorry
      :sender B
      :receiver A
      :in-reply-to id1
      :reply-with id2 )
```

agent A (:sender) is querying the agent B (:receiver), in Prolog (:language) about the truth status of ``bar(a,b)'' (:content)

KQML Example 2

- query

```
(stream-about :language
  KIF :ontology motors
  `:reply-with q1
:content motor1)
```

- reply

```
(tell :language
  KIF :ontology motors :in-
  reply-to q1
: content (= (val (torque
  motor1) (sim-time 5)
  (scalar 12 kgf))
(tell :language
  KIF :ontology
  structures :in-reply-to q1
: content (fastens frame12
  motor1))
(eos :in-repl-to q1)
```

agent A asks agent B to tell all it knows about motor1.
B replys with a sequence of tells terminated with a sorry.

Important Concepts and Terms

automated reasoning
belief network
cognitive science
computer science
deduction
frame
human problem solving
inference
intelligence
knowledge acquisition
knowledge representation
linguistics
logic
machine learning
natural language
ontology
ontological commitment
predicate logic
probabilistic reasoning
propositional logic
psychology
rational agent
rationality
reasoning
rule-based system
semantic network
surrogate
taxonomy
Turing machine