Expert Systems and Knowledge Representation

Guide to AI and Data Courses · Part II: Branches of AI

Knowledge representation | Rule base | Inference engine | Ontologies | Knowledge acquisition | Neighboring courses

Sercan KÜLCÜ · Giresun University · Computer Engineering

© Sercan Külcü. All rights reserved. · Interactive HTML version of these slides: https://sercankulcu.github.io/files/ai_courses/slides/Chapter_11_Expert_Systems.html

INTRODUCTION

What Question Does It Ask?

> **How is an expert's knowledge written down and used for inference?**

An expert system writes a domain expert's knowledge into a **knowledge base** of rules and facts and draws conclusions from it with an inference engine. Knowledge representation studies how to encode knowledge so that a computer can reason with it: logic, frames, semantic networks and ontologies.

At a glance

APPROACH	INPUT	OUTPUT
symbolic logic and rules	facts and expert rules	conclusion with justification

SUCCESS MEASURE	MATH LOAD	DATA NEEDS
expert agreement, explainability	low: propositional and predicate logic	little; a knowledge base

A Short History

1965
DENDRAL starts at Stanford: the first expert system, inferring molecular structure from mass spectra.

1972
Colmerauer and Roussel develop Prolog: logic programming is born.

1970s
MYCIN recommends antibiotics for blood infections, adding certainty factors to its rules.

1980
R1 (XCON) starts configuring computer orders at DEC: the first big commercial success.

1984
Lenat launches Cyc, aiming to hand-code common-sense knowledge.

2012
Google announces the Knowledge Graph; knowledge graphs spread in search.

THE COURSE

Core Concepts

Knowledge base
Domain facts and IF-THEN rules.

Inference engine
Applies rules to facts to derive new conclusions.

Forward chaining
Start from facts and derive every reachable conclusion.

Backward chaining
Start from a goal and look for facts that support it.

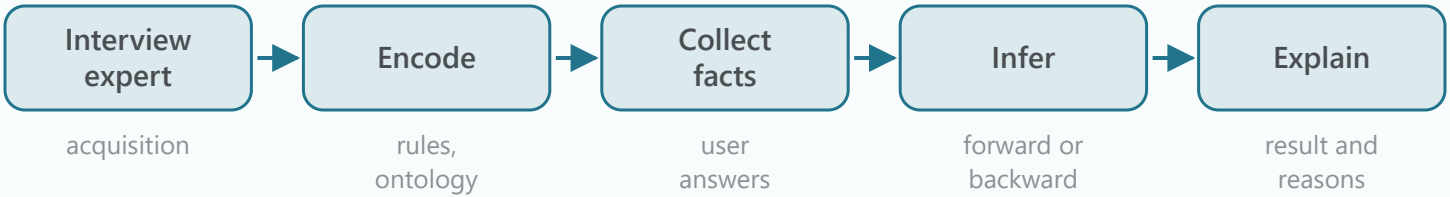
Explanation facility
Shows which rules led to a conclusion.

Frames and networks
Representing objects with slots and relations.

Ontology
A domain's concepts, classes and relations: OWL.

Knowledge acquisition
Eliciting knowledge from experts; the bottleneck.

A Typical Workflow



Example: A fault diagnosis system turns technicians' experience into rules such as "if the engine won't start and battery voltage is low, replace the battery". It asks the user about symptoms, chains the rules to find the fault and shows step by step which rules it used.

THE COURSE

What Does the Course Cover?

1. The knowledge representation and reasoning problem
2. Propositional and first-order logic; resolution
3. Logic programming: Prolog
4. Rule-based systems; forward and backward chaining
5. The Rete algorithm and rule engines: CLIPS, Drools
6. Frames, semantic networks and inheritance
7. Description logics
8. Ontologies: RDF, OWL and Protégé
9. Knowledge graphs and SPARQL queries
10. Uncertainty: certainty factors, default reasoning
11. Knowledge acquisition and expert interviews
12. Explanation, verification and knowledge base upkeep

Skills you will gain

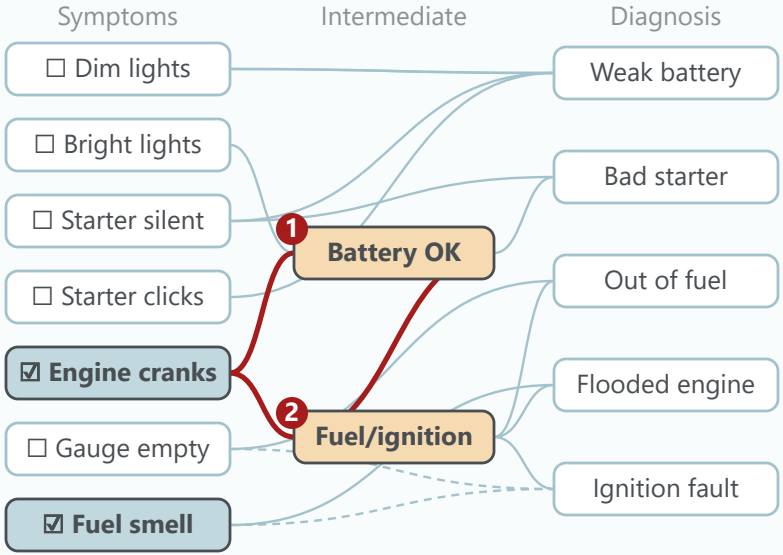
writing rules in CLIPS or Drools · logic programming in Prolog · building ontologies in Protégé · querying knowledge graphs with SPARQL · eliciting expert knowledge as rules

A typical semester; the order and coverage vary between universities.

TRY IT

Why Won't the Car Start?

A mechanic's knowledge was written as nine IF–THEN rules. Click the symptoms on the left to mark them; the system fires rules one by one by forward chaining, reaches a diagnosis and explains why.



Example: battery | Example: starter | Clear

1. **K4**: Engine cranks → Battery OK
2. **K6**: Battery OK ∧ Engine cranks → Fuel/ignition

Inference in progress…

Knowledge lives in a **rule base** separate from the code; the inference engine chains the rules. When the mechanic adds a rule, the program does not change, and the system can explain which rules produced each decision.

CONTEXT

Prerequisites

Courses worth knowing before this one, and why they are needed.

Course	Why it is needed
Artificial Intelligence	symbolic AI, logic and rules
Discrete Mathematics	propositional and predicate logic, inference rules
Data Structures	trees, graphs and hash tables: storing rules and facts
Object Oriented Programming with Java	classes and inheritance: a bridge to frames and Drools

CONTEXT

Are You Ready for This Course?

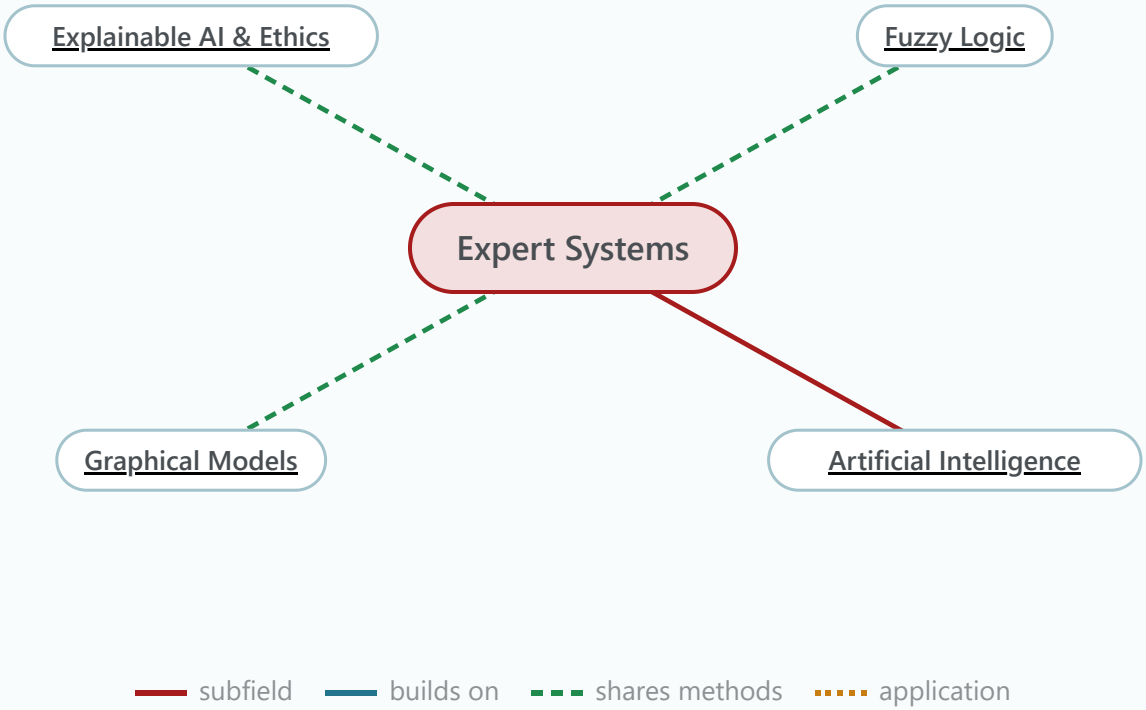
Ask yourself; for every "no", the right-hand column shows where to start.

Ask yourself	If not, start with
Can you write a sentence such as "all humans are mortal" in first-order logic?	Propositional and predicate logic from discrete mathematics
Can you run depth-first and breadth-first search on an example?	Artificial Intelligence
Can you use a hash map and a graph in a program?	Hash tables and graphs from a data structures course
Can you explain classes, objects and inheritance with an example?	An intro to object-oriented programming in Java

If you can say "yes" to most of these, you are ready; one or two gaps can be closed during the semester.

CONTEXT

Neighbouring Courses



Overlapping Courses

Relation	Course	Shared ground
subfield	This course is a subfield of Artificial Intelligence	symbolic AI, logic and rules
shares methods	Shares methods with Fuzzy Logic and Fuzzy Systems	IF-THEN rule base, inference engine
shares methods	Shares methods with Probabilistic Graphical Models	inference under uncertainty
shares methods	Shares methods with Explainable AI and AI Ethics	explainable decisions through readable rules

DIFFERENCES

Where It Differs (1/2)

Criterion	Expert Systems	Machine Learning	Fuzzy Logic
Core question	How is an expert's knowledge written down and used for inference?	How does a program generalize from examples?	How do we reason with concepts like "hot" or "fast" that have no sharp boundary?
Approach	symbolic logic and rules	statistical learning	partial membership and linguistic rules
Input	facts and expert rules	labeled or unlabeled examples	measurements, expert rules
Output	conclusion with justification	a model that makes predictions	control signal or decision
Success measure	expert agreement, explainability	error on unseen data	control quality, expert agreement
Math load	low: propositional and predicate logic	high: linear algebra, probability, optimization	low-medium: set theory
Data needs	little; a knowledge base	moderate to large	little; expert knowledge

- **Machine Learning:** ML learns knowledge from examples and often gives no reasons; an expert system takes knowledge from experts and explains its result with rules.
- **Fuzzy Logic:** Fuzzy logic lets rules be partly true; this course teaches representing knowledge and chaining rules in general.

DIFFERENCES

Where It Differs (2/2)

Criterion	Expert Systems	Graphical Models	Explainable AI & Ethics
Core question	How is an expert's knowledge written down and used for inference?	How do we model and query many interdependent, uncertain variables?	When can we trust an AI system's decision, and who is responsible for the outcome?
Approach	symbolic logic and rules	probability and graphs combined	human-centered evaluation
Input	facts and expert rules	variables, dependencies, data	model, its data and context of use
Output	conclusion with justification	conditional probability, best explanation	explanation, audit report, principles
Success measure	expert agreement, explainability	likelihood, inference accuracy	interpretability, fairness, accountability
Math load	low: propositional and predicate logic	high: probability	low-medium
Data needs	little; a knowledge base	moderate	the model's own data

- **Graphical Models:** PGMs model uncertainty consistently with probability distributions; expert systems mostly write knowledge as crisp rules.
- **Explainable AI & Ethics:** XAI seeks explanations for learned models after the fact; an expert system explains itself because it is built from readable rules.

PRACTICE

Applications

Clinical decision support
Diagnosis and treatment advice from symptoms and tests.

Fault diagnosis
Finding machine, vehicle and network problems.

Product configuration
Building orders from compatible parts.

Business rules
Loan approval, insurance and compliance checks.

Knowledge graphs
Entity and relation data in search engines.

Semantic web
Linking data through ontologies.

PRACTICE

Risks and Ethics

Bias written into rules

An expert's habits and an institution's old practices become rules; credit or insurance rules can quietly exclude some groups.

Rule conflicts and upkeep

Thousands of rules can shadow and contradict each other; in large systems such as XCON, maintaining the rules became a burden of its own.

Unclear responsibility

After a wrong diagnosis or a rejected application, it is unclear whether the fault lies with the expert, the knowledge engineer or the user.

Errors in knowledge bases

Wrong or outdated facts in open knowledge bases spread to the search and chat systems built on them.

In depth: Explainable AI and AI Ethics.

PRACTICE

Tools and Datasets

Tools and libraries

CLIPS | Drools | SWI-Prolog | Protégé | Apache Jena | Neo4j

Well-known knowledge bases

Wikidata | DBpedia | WordNet | Gene Ontology | SNOMED CT | Protégé Pizza ontology

PRACTICE

Common Misconceptions

Common belief	In fact
Expert systems were a 1980s fad.	Under new names, business rule engines and knowledge graphs are used today in banking, healthcare and search.
An expert system knows as much as the expert.	It knows only what was encoded; it fails quickly outside its domain or when common sense is needed.
Writing the rules is easy.	Much expert knowledge is tacit; turning it into rules (knowledge acquisition) is the hardest and costliest step.
Large language models made knowledge representation obsolete.	Language models can produce false facts; they are paired with knowledge graphs for verifiable results.

NEXT STEPS

After This Course

Courses that build on or extend this one

This course is mostly an end point: it builds on the others.

Related careers

Knowledge engineer | Ontology specialist | Decision support systems developer | Business rules analyst

This course is the gateway to symbolic AI. Next you can take Probabilistic Graphical Models to handle uncertainty consistently, Fuzzy Logic for partial truth, or Explainable AI and AI Ethics for explaining decisions and accountability.

Further Reading

- Jackson, *Introduction to Expert Systems*, 3rd ed., Addison-Wesley, 1999.
- Brachman and Levesque, *Knowledge Representation and Reasoning*, Morgan Kaufmann, 2004.
- Giarratano and Riley, *Expert Systems: Principles and Programming*, 4th ed., Thomson Course Technology, 2005.
- Buchanan and Shortliffe (eds.), *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*, Addison-Wesley, 1984.

The End

Chapter 11: Expert Systems and Knowledge Representation

← Previous chapter · Map of the field · Next chapter →