

Artificial Intelligence Programming

Ontologies

Chris Brooks

Department of Computer Science
University of San Francisco

Knowledge Engineering

- Logic provides one answer to the question of *how* to say things.
- It does not tell us *what* to say.
- Typically, this is the hard part.
- We want to give our agent enough knowledge to solve all of the problems we are interested in.
- The process of building a knowledge base is referred to as *knowledge engineering*
 - Many of the same principles as software engineering.

Department of Computer Science — University of San Francisco — p.177

Knowledge Engineering

- The knowledge engineering process typically consists of the following steps.
 1. Determine the queries and types of facts that are available/allowed
 - For example, will the agent need to select actions or make conclusions, or just answer questions from a human user?
 - Will we ask existential queries, or just universal queries?
 2. Gather the relevant knowledge
 - Interview experts and find out how the domain works.

Department of Computer Science — University of San Francisco — p.178

Knowledge Engineering

- The knowledge engineering process typically consists of the following steps.
 - 3. Select a vocabulary of classes, functions, predicates and constants
 - This vocabulary is called an *ontology*
 - 4. Encode general knowledge about the domain
 - Formally represent the knowledge gathered in step 2.
 - This may require revisiting step 3.
 - 5. Encode specific queries or problems to be solved.
 - This may require returning to steps 3 and 4.

Department of Computer Science — University of San Francisco — p.377

Ontologies

- An ontology is a vocabulary that describes all of the objects of interest in the domain and the relations between them.
 - All of the relevant knowledge about a problem we are interested in.
- Ontologies allow knowledge about a domain to be shared between agents (including humans).
- This can allow knowledge to be reused more easily.
- Allows an agent to perform inference with its current knowledge.

Department of Computer Science — University of San Francisco — p.477

Vocabulary

- An ontology consists of:
 - A set of *concepts* or *classes*
 - $Guitarist(GeorgeHarrison), \forall x Guitarist(x) \rightarrow Musician(x)$
 - Instances of these classes.
 - Relations between classes of objects, sometimes called *slots* or *properties* or *roles*.
 - $performs(Beatles, WhiteAlbum), containedOn(RockyF, ...)$
 - Restrictions on properties (sometimes called *constraints*)
 - $\forall x, y Artist(x) \wedge Song(y) \wedge containedOn(y, z) \wedge performs(x, y) \rightarrow performsOn(x, z)$

Department of Computer Science — University of San Francisco — p.478

Ontologies vs OO design

- Classes are a primary focus of ontology design.
- In many ways, this looks like object-oriented design.
- We have classes and subclasses, and properties of classes that look like data members.
- However, properties have richer semantics than data members.
- A property may attach to several classes at once, and have subproperties.
- We can specify constraints on the values in a slot's range.
- Properties can exist without being assigned to a class.

Ontology tools

- An popular means for assisting in knowledge engineering is the use of GUI-based tools.
- Details of logic and inference are hidden from the user.
- Protege is one of the more well-known ontology development tools.

OWL

- Web Ontology Language (OWL) is an XML-based language for representing ontologies.
- Built on top of RDF
- Encodes a *description logic*
 - Decidable subset of first-order logic.
- We'll be using a graphical tool (Protege) that uses OWL as its underlying representation.

Types of OWLs

- OWL Lite
 - Does not allow cardinality (except 0 and 1)
 - Primarily classes, simple constraints/rules
 - Efficient inference mechanisms exist
- OWL DL
 - Complete and decidable, but more expressive.
 - Inference mechanisms exist.
 - This is what we'll be using.
- OWL Full
 - Allows metaclasses, richer reasoning about classes
 - No inference mechanisms exist for all of OWL Full.

Uses of OWL

- Organization and indexing of journals and scientific literature
- Human Genome Project (and bioinformatics in general)
- Data integration between DBs in an enterprise
- Automated Web Service description and composition
- and others ...

The Pizza Tutorial

- The pizza tutorial provides a nice tour of the issues involved in creating an ontology in Protege with OWL.
- We begin by asking *competency questions* - these are questions we'd like our KB to be able to answer.
 - What toppings are on a Margherita pizza?
 - Is the Americana pizza vegetarian?
 - What can I put on my pizza to make it spicy?
 - What pizzas have Tomato on them?

Classes

- As with OO design, we can work top-down, bottom-up, or in a combination of the two.
- Let's start by making a Pizza class.
- We then make PizzaBase and PizzaTopping.
 - Make these disjoint.

Classes

- Use the Wizard to add subclasses of PizzaBase: DeepDishBase, ThinAndCrispyBase
- Use the Wizard to subclass PizzaTopping: MeatTopping, VeggieTopping, CheeseTopping, SeafoodTopping
- Subclass these.

Properties

- Now we can start to create Properties.
- Add a 'hasIngredient' property.
 - In Protege, properties can have subProperties.
 - Add subproperties hasTopping, hasBase.
- We can also add inverse properties.
- Also functional, inverse functional, symmetric, and transitive.
- make hasBase functional.

Domains and ranges

- We can specify that properties can only apply to certain types of objects.
- set the range of hasTopping to be PizzaTopping
- Set the domain of hasTopping to be Pizza.
- Same for hasBase

Property Restrictions

- The primary way that we write rules in Protege is through the use of property restrictions.
- We can use an existential restriction to specify that a Pizza must have a PizzaBase.

Types of pizzas

- Add a subclass of pizza called NamedPizza.
- Add a subclass of this: MargheritaPizza.
- existential restriction: has Tomato and Mozzarella toppings.
- Add Americana Pizza: clone Margherita, add Pepperoni
- Add AmericanaHotPizza: clone Americana, add jalapenos.
- Add SohoPizza: Margherita plus Olives, Parmesan.

Reasoners

- As your ontology gets large, it can be difficult to maintain:
 - Do you have the hierarchy right?
 - Are there contradictions?
 - Are there other relations that are entailed?
- A big advantage of using this sort of tool is the ability to perform this inference automatically.
- Tools for doing this are known as reasoners.
- We'll be using a particular reasoner known as RACER.

Reasoners

- We can use the reasoner to:
 - Check consistency
 - Classify the taxonomy
 - Infer class membership
- Let's add an inconsistent class: InconsistentTopping
 - Subclass of Cheese
 - Restriction: must be a Veggie.
- Run the reasoner.

Necessary and Sufficient Conditions

- So far, we've only specified *necessary* conditions.
- This is not enough to show that something must be a member of a particular class.
- Create a CheesePizza class
- Add a necessary and Sufficient restriction - has a Cheese topping.
 - This is now a *definition* - a CheesyPizza is a Pizza with a Cheese topping.
- Run RACER again.

Universal restrictions

- We specified that a Margherita pizza must have Tomato and Mozzarella, but not that those were the only things it could have.
- To do this, we use universal restrictions.
- Create a VeggiePizza.
- Under necessary and sufficient, indicate that all toppings must be veggie or cheese.
- Run RACER again.

Open World Reasoning

- Why is Margherita (and Soho) not subclassed as Veggie?
- Protege and OWL use open-world reasoning
 - Can't assume something is false just because it's not stated to be true.
- Margherita could have other toppings.
- Add a closure axiom to MargheritaPizza.
- Use the widget to add closure axioms for Soho and Americana.

Value Partitions

- Use the Wizard to create a SpicynessValuePartition
- Subclass this to get Hot, Mild, Medium
- Add hasSpiciness property
- Use the PropertyRestrictions wizard.

Cardinality

- We can also specify the number of toppings a pizza must have.
- Create a CrowdedPizza
- Necessary and Sufficient: minimum cardinality: 3.

Individuals

- We can then add individuals using the individuals tab.
- the Forms tab lets us specify how individuals should be displayed.
- Add a NamedPizza that has tomato and Mozzarella
- Run the reasoner again.

Querying the KB

- The query pane lets us create, save and retrieve queries
- We can specify either AND or OR.
- Try finding all pizzas that have a cheese topping.