Ontological Resolution in an Open System of Agents

You say “tomato”, I say “*Lycopersicon esculentum.”*

An Open System of Agents

The knowledge sharing effort (**KSE**) of the early to mid nineties initiated research into the mechanisms needed to achieve knowledge sharing between agents.

Most of the work was centered on **syntax**, **semantics**, and the **pragmatics of communication**.

An Open System of Agents

KQML (Knowledge Query and Manipulation Language) and **KIF** (Knowledge Interchange Format) are two examples of the results of the **KSE**.

Much of the work on sharing knowledge (using KQML and KIF) assumed the existence of a **common ontology** and the **relative ease of locating sources of knowledge**, both of which are practical only in **closed systems**.

An Open System of Agents

As the Internet continues to transition from a human-centric system to a machine-centric system, we have by sharing knowledge, a unique opportunity to construct knowledge-based systems with lower cost and greater knowledge than ever.

However, on the scale of the **whole Internet**, we cannot assume the luxuries of a tightly controlled closed-system.

An Open System of Agents

Specifically in an **open system**…

- Locating sources of knowledge becomes a process rather than an event
- Agents are assumed to be social, but not necessarily fully cooperative
- No common ontology is assumed to exist

An Open System of Agents

Although there is no globally mandated ontology in the open system, some group of agents may agree to use a standard ontology.

A group of agents that subscribe to some locally common ontology form a **Community of Practice**.

J. Sowa's "Knowledge Soup"

http://www.jfsowa.com/talks/negotiat.htm

An Open System of Agents

A CoP represents some knowledge specialization. Other agents may not be similarly specialized but may need access to some knowledge contained within a CoP.

To enable knowledge sharing across CoP boundaries, agents need a mechanism for resolving ontological differences.

Example

A personal (consumer) agent that uses the publicly-available **`DMOZ`** (Open Directory Project) (omni) ontology is seeking knowledge about the concept ”**`DMOZ:Lung Cancer`**”

A CoP is formed around the **`MeSH`** (Medical Subject Headings) ontology. This CoP includes four agents representing institutions involved in oncological research.

MeSH is published by the National Library of Medicine (NLM) consists of some 30,000 unique terms. The example following is *extremely* simplified from MeSH. A copy of MeSH in flat ASCII or XML format can be obtained from http://www.nlm.nih.gov/mesh/meshhome.html

DMOZ aka the Open Directory Project is a collaborative WWW categorical index. DMOZ is an *omni* ontology in that it is intended to cover in a shallow manner every field of human knowledge, rather than a specialized ontology for some well defined subdomain. The DMOZ ontology and index are publicly available from http://dmoz.org

Example: DMOZ Ontology

DMOZ

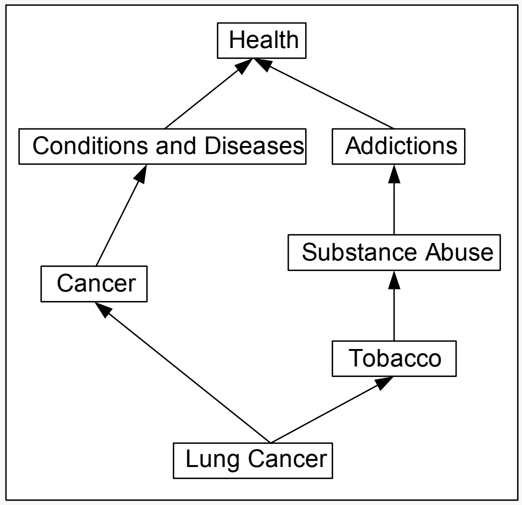

Example: MeSH Ontology

MeSH

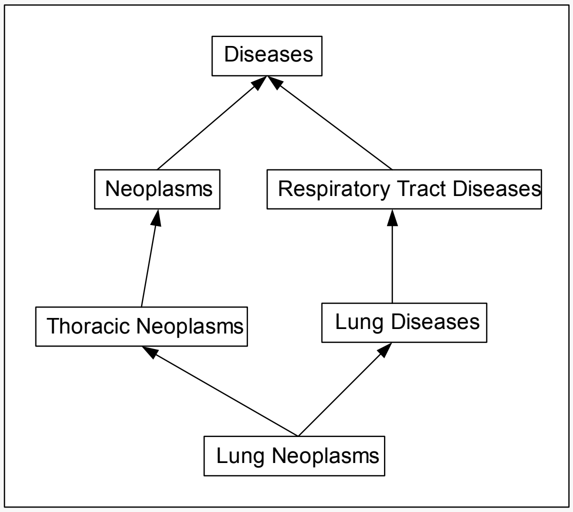

Example: CoP Boundary

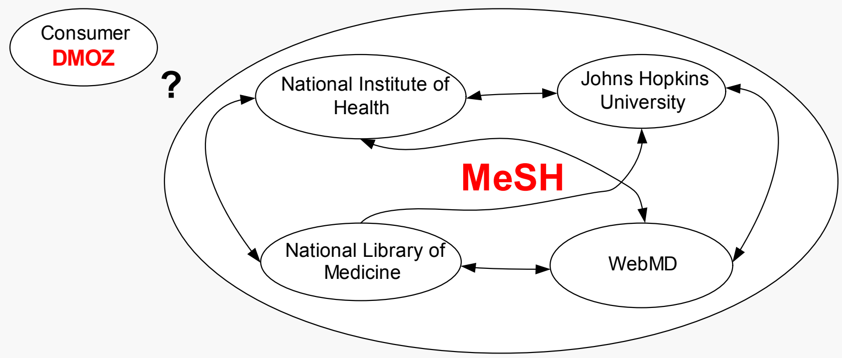

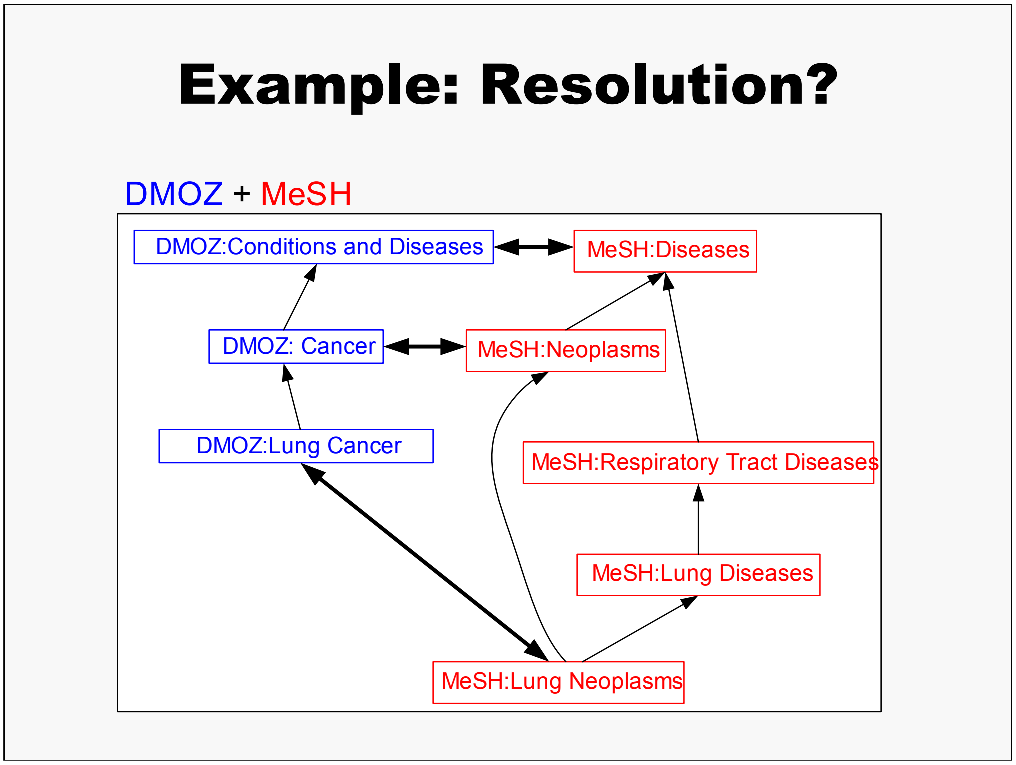

The DMOZ: and MeSH: prefixes are denoting namespaces.

Example

After ontological resolution the personal agent can issue a **`DMOZ:Lung Cancer`** query to agents in the **`MeSH`** CoP. The receiving agent in the **`MeSH`** CoP can then resolve **`DMOZ:Lung Cancer`** to **`MeSH:Lung Neoplasms`**, or relax the query to **`MeSH:Lung Diseases`** or **`MeSH:Neoplasms`**

Related Work: DIAMS

- A collaborative information retrieval tool from NASA
- Uses word vectors to represent concept spaces
- Concept spaces can be resolved by word vector comparison
- No formal semantics (no explicit relationships between concept spaces)

Related Work: InfoSleuth

- Dynamic federated information system
- MAS with specialized agents (UI, ontology, etc)
- Uses word vectors for ontological resolution

Related Work: UMDL

- University of Michigan Digital Library
- An open system
- Agents are social, but not necessarily cooperative
- Ontology-based
- Ontologies can be seeded and grown dynamically.
- Many classes of agents

Related Work: DOGGIE

"Teaching Agents to Share Meaning" by A.B.Williams and Z. Ren of the University of Iowa introduced a multi-agent system they called DOGGIE.

- DOGGIE uses a distributed collective memory (DCM) to provide exemplars for semantic concepts.

`DOGGIE` is the Distributed Ontology Gathering Group Integration Environment. A rather contrived but cute acronym.

"Teaching…" was published in the proceedings of the Agents '01 conference sponsored by the ACM.

DOGGIE: Overview

Locating similar semantic concepts…

1) An agent sends a concept query (concept and a list of DCM pointers) to its peers.

2) Some peer receives the query and attempts to **interpret** the objects referred to in the query.

DOGGIE: Overview

3) The peer responds to the agent with **know**, **maybe know**, or **don't know**. If the peer sends **know** or **maybe know**, the response contains a set of sample pointers that describe its semantic concept.

4) The querying agent receives the responses from its peers and attempts to **verify** whether other agents know a similar concept.

DOGGIE: Overview

5) If the querying agent can verify that a peer has a similar semantic concept, it remembers this information by an insertion into its knowledge base.

DOGGIE: Formally

The DOGGIE system is formally described as…

A multi-agent system $\mathbf{A = \{a_1, a_2, \ldots\}}$

Agent $\mathbf{a_1}$ knows a semantic concept $\boldsymbol{\varphi}$, or $\mathbf{K(a1, \varphi)}$

DOGGIE: Formally

$\mathbf{a_1}$ sends a concept-based query (CBQ) to its peers $\mathbf{a_{peer}} \in \mathbf{A}$

The CBQ is a tuple consisting of the semantic concept and a set of DCM pointers to exemplars of the concept.

$\mathbf{CBQ = \langle \varphi, X_{\varphi} \rangle}$

DOGGIE: Formally

For each semantic concept known by $\mathbf{a_i}$, there is an associated set of exemplars that make up the semantic concept $\boldsymbol{\varphi}=\{\mathbf{x_1},\mathbf{x_2},\ldots\}$

For $\boldsymbol{\varphi}$ there exists a function (the classifier) such that $\mathbf{c(x)}=\boldsymbol{\varphi}$

DOGGIE: Formally

c is assumed to be non-computable. A supervised inductive learning algorithm is used to approximate **c**, $\mathbf{h(x)} \approx \mathbf{c(x)}$

The induction algorithm was used for every concept known to agent $\mathbf{a_i}$

$\mathbf{H(x)=\{h_1,\ldots,h_n\}}$ where $\mathbf{h_1(x)=\varphi_1}$ and $\mathbf{h_n(x)=\varphi_n}$

DOGGIE: Formally

For each learned concept $\mathbf{h_i}$ in $\mathbf{H_\varphi(x)}$ there exists a corresponding ratio describing how often this particular concept description correctly determined that an object in the training set belonged to concept φ.

This ratio is the **positive interpretation threshold** for concept φ, or φ_+

The **negative interpretation threshold** is $\varphi_-=1-\varphi_+$

DOGGIE: Formally

If agent $\mathbf{a_2}$ sends $\mathbf{k}$ addresses of its concept $\boldsymbol{\varphi}$ to agent $\mathbf{a_1}$, then agent $\mathbf{a_1}$ uses its set of concept descriptions $\mathbf{H(x)}$ as inference rules and seeks to interpret the example objects sent to it, $\mathbf{X_{a2}=\{x_1,\ldots,x_k\}}$.

The interpretation value $\mathbf{v}$ of concept $\boldsymbol{\varphi_j}$ is the frequency that $\boldsymbol{\varphi_j}$ is inferred $\mathbf{f\varphi_j}$ versus the size of the exemplar set, $\mathbf{k}$. $\mathbf{(f\varphi_j/k)} = \mathbf{v\varphi_j}$

DOGGIE: Formally

Agent a1 then compares the interpretation value $v\varphi_j$ to that concept's positive and negative interpretation value, φ_{j+} and φ_{j-}

$v\varphi_j >= \varphi_{j+}$ **K**

$v\varphi_j >= \varphi_{j-}$ **D**

$\varphi_{j-} < v\varphi j < \varphi j_+$ **M**

DOGGIE: Formally

Where…

K: agent $\mathbf{a_1}$ **knows** the concept φ

D: agent $\mathbf{a_1}$ **does not know** the concept φ

M: agent $\mathbf{a_1}$ **may know** the concept φ

This result as well as an exemplar set taken from agent $\mathbf{a_1}$'s concept φ_j is sent as a response to agent $\mathbf{a_2}$.

New System

By using both intensional and extensional reasoning **DOGGIE** seems to be on the right track to better ontological resolution.

DOGGIE can be generalized in many ways.

I have not yet come up with a catchy acronym for my project.

New System

(1) The classification system that **DOGGIE** uses is boolean. That is for some object $\mathbf{x}$ and a classification predicate $\mathbf{c(x,\varphi)} \in \mathbf{\{0,1\}}$

Relaxing the boolean constraint so that the classification predicate is defined on the continuum, similar to $\mathbf{c(x,\varphi)} \in \mathbf{[0,1]}$ where the value is reflective of the degree of inclusion of object $\mathbf{x}$ in concept $\boldsymbol{\varphi}$

New System

Assume we have some document with the following sentences:

p	r_1	o_1
p	r_2	o_2
p	r_3	o_3
p	r_4	o_4
p	r_1	**q**
p	r_2	**q**

Of the six sentences, all six are about **p** where two are about **q** and **p**.

Should we classify this object as belonging to concept **p**? Concept **q**? Concept **p** and **q**?

What if there were 100 sentences?

New System

(2) **DOGGIE** defines some ontology $\mathbf{O_i}$ as the set of concepts $\{\varphi_1, \varphi_2, \ldots\}$

DOGGIE can include a set of relations between concepts, but there is no explicit way of propagating learned information based on the relations in the ontology.

New System

Consider the previous example with **`MeSH`**. We can look up the grammar for **`MeSH`** (the symbols and the relations between them) but not know the semantics embodied in the ontology.

However, knowing both the grammar and some knowledge of a concept in the ontology (ie, **`MeSH:Lung Neoplasms`**) we then also have **some** knowledge of related concepts.

New System

(3) Instead of assuming a global ontology, **DOGGIE** takes the other extreme and assumes that no two agents share the same ontology. (That is there are no CoP of size greater than 1.)

This is in essence a constraint on **DOGGIE**'s reasoning and can be removed to generalize the process of ontological resolution.

New System

If we relax this constraint and allow CoPs of arbitrary size then…

(a) interaction with a single agent in the CoP gains knowledge of the ontology that is central to the CoP and hence some knowledge about all agents in the CoP.

If our personal agent talks only to the NLM and learns that DMOZ:Lung Cancer is related to MeSH:Lung Neoplasms, this information is directly portable when communicating to Johns Hopkins, etc. No further negotiation is necessary *if* we are confident of the relation.

New System

(b) knowledge of the CoP (both its ontology and the capabilities of the member agents) is limited by **sampling error**.

DOGGIE does take into account error in classification with its positive/negative interpretation value. A more rigorous statistical approach may yield some benefit.

At the extreme how much can one infer about an entire domain from one or two exemplars?

System Formulation

A set of pointers to objects in a distributed collective memory

$DCM = \{o_1, o_2, \ldots\}$

A set of agents…

$A = \{a_1, a_2, \ldots\}$

System Formulation

A set of ontologies…

$\Phi = \{\varphi_1, \varphi_2, \dots\}$

An ontology $\varphi \in \Phi$ is the DAG $\{T_\varphi, R_\varphi\}$ where
$T_\varphi = \{t_1, t_2, \dots\}$ is the set of terms and
$R_\varphi = \{<t_i, t_j>, \dots\}$ is the set of relations between terms.

System Formulation

A mapping **`exemplars(a,φ,t)`** that given agent **`a`**, ontology φ and the concept $\mathbf{t} \in \mathbf{T}\varphi$ produces the set of tuples $\{(w_1, o_1), \ldots\}$ where $o_1 \in DCM$ and $w_1 \in [0,1]$ is the degree of inclusion of o_1 in the concept represented by **`t`**.

System Formulation

Agent **a** is said to know ontology φ

$\mathbf{K(a,\varphi)} \Leftrightarrow \forall t \in T_{\varphi}[|\mathrm{exemplars}(a,\varphi,t)|>0]$

A Community of Practice around ontology φ,
$A_{\varphi}=\mathrm{CoP}(A,\varphi)$ is a subset of **A** induced by
$a \in A_{\varphi} \Leftrightarrow K(a,\varphi)$

Problem Formulation

We seek some function that given a term $\mathbf{t} \in T_{\varphi 1}$ and a foreign ontology φ_2 the function produces a set of tuples of the form $\{(c_1, t_1), \ldots\}$ where c_1 is the degree of inclusion of $\mathbf{t}$ in the concept $t_1 \in T_{\varphi 2}$

or…

$$\mathtt{map}_{t,\varphi 1,\varphi 1} = \mathtt{commensurate}(t, \varphi_1, \varphi_2)$$

More Problems

If we have some metric of introspective knowledge, ie. A measure of how much is known about components of some foreign ontology, then we can use this knowledge in guiding our selection of exemplars in the learning dialog with another agent.

Ie. We wish to maximize the knowledge gain, while minimizing sampling error.

In other words, the very choice of which (sub)set of exemplars we use to communicate a concept influences an agent's perception of the intended meaning of the concept. If we know something about the other agent, and we are attempting to explain a concept in our ontology in terms of the other agent's ontology, then we can use this knowledge of the other agent to choose a good exemplar subset to refer to in the learning dialog.

Contributions

The expected contributions of this project are:

- An automated method of resolving differences in ontologies
- A reasoning method that includes CoPs.
- A measurement of introspective knowledge about an ontology
- Given a subset of the DCM, we will formulate an optimal lesson plan to teach an agent some semantic concept in some ontology.