Il ruolo delle ontologie nella rappresentazione della conoscenza

Nicola Guarino
Laboratorio di Ontologia Applicata (LOA)
Istituto di Scienze e Tecnologie della Cognizione (ISTC-CNR)
Trento, Italy

1. The emergence of ontologies in AI

Representation vs. Reasoning

- Representation comes first!
- The very task of representation (i.e. ***modelling***) is left to the user
- AI researchers focus more on the ***nature of reasoning*** than in the ***nature of the real world***

Essential ***ontological promiscuity*** of AI?
(Genesereth and Nilsson 1987)

The need to focus on content

- Philosophers have generally stopped short of trying to actually specify the truth conditions of the basic atomic propositions, dealing mainly with the specification of the meaning of complex expressions in terms of the meanings of elementary ones. ***Researchers in artificial intelligence are faced with the need to specify the semantics of elementary propositions*** as well as complex ones.

  [Woods 1975]

- The majority of work in knowledge representation has been concerned with the technicalities of relating predicate calculus to other formalisms, and with the details of various schemes for default reasoning. There has been almost an aversion to addressing the problems that arise in actually representing large bodies of knowledge with content. The typical AI researcher seems to consider that task to be 'just applications work'. But there are deep, important issues that must be addressed […]: What ontological categories would make up an adequate set for carving up the universe? How are they related? What are the important things most humans today know about solid objects? And so on. In short, we must ***bite the bullet***.

  [Lenat&Guha 90] (our italics).

Kinds, roles, attributions

How many rock kinds are there?

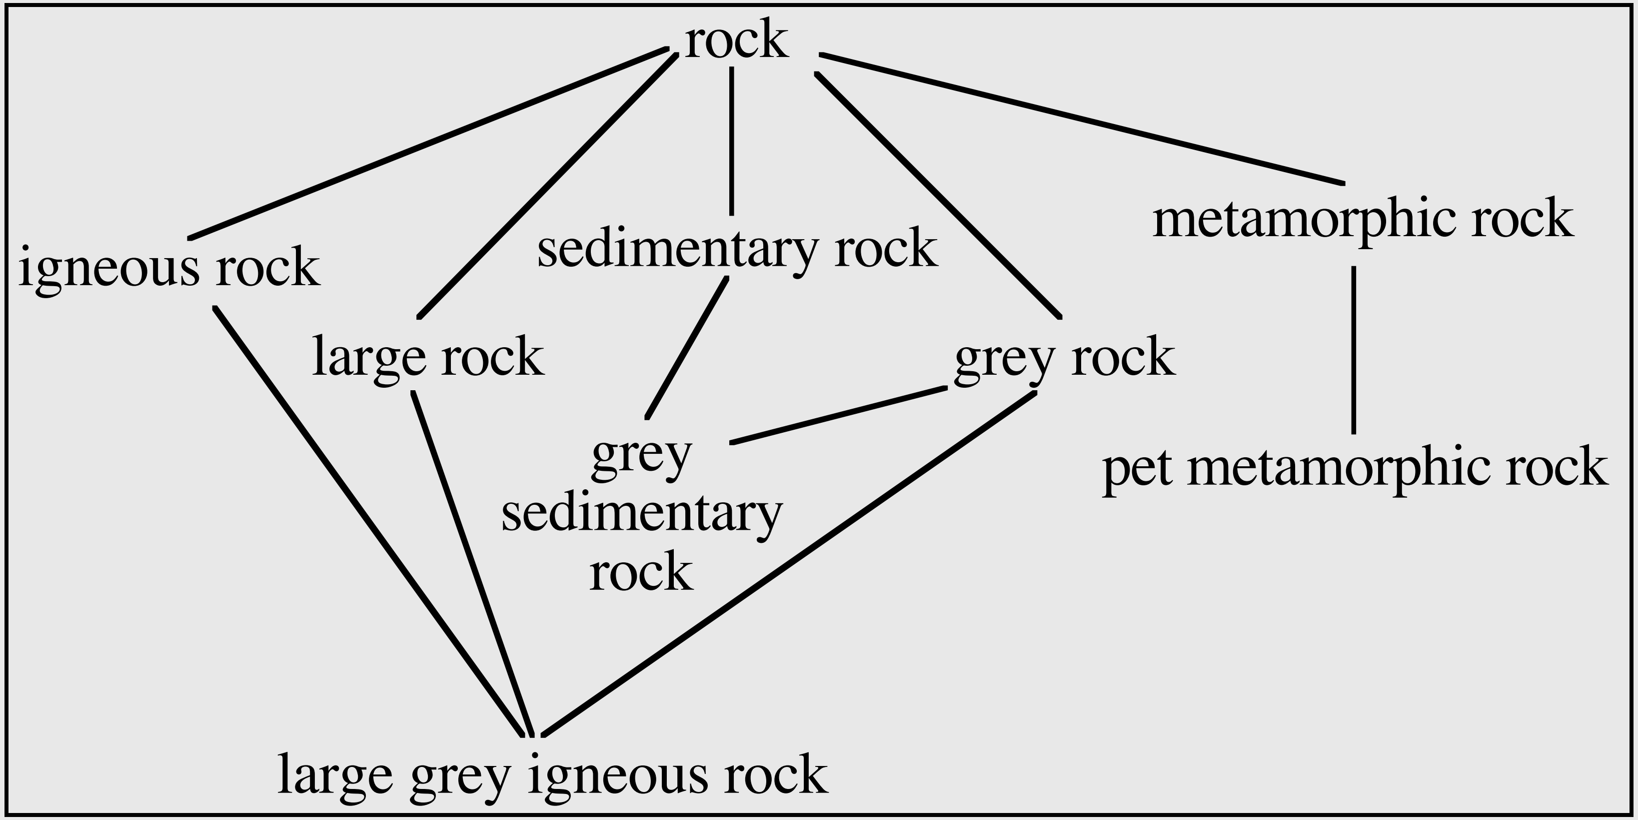

[From Brachman, R ., R. F ikes, et al. 1983. "Krypton: A Functional Approach to Knowledge Representation", *IEEE Computer*]

The answer

- According to Brachman & Fikes 83:
  - It's a dangerous question, only "safe" queries about analytical relationships between terms should be asked
- In a previous paper by Brachman and Levesque on ***terminological competence*** *in knowledge representation* [AAAI 82]:
  - "an *enhancement mode transistor* (which is a *kind* of transistor) should be understood as different from a *pass transistor* (which is a *role* a transistor plays in a larger circuit)"
- These issues have been simply *given up* while striving for logical simplification and computational tractability
- The OntoClean methodology, based on formal ontological analysis, allows us to conclude: ***there are 3 kinds of rocks*** (appearing in the figure)

2. Focusing on content

The focus of ontological analysis: from form to *CONTENT*

- The key problems
  - content-based information access *(**semantic matching**)*
  - content-based information integration *(**semantic integration**)*

- To approach them, content must be studied, understood, analyzed ***as such***, independently of the way it is represented.

- Computer technologies are not really good for that...
- …and users of computer systems are often ***confused by technology***

Ontologies: a ***magic solution?***

No ontology without ***ontological analysis!***

The problem: subtle distinctions in meaning

"Trying to engage with oo many partners too fast is one of the main reasons that ***so many online market makers have foundered***.

The transactions they had viewed as simple and routine actually involved many ***subtle distinctions in terminology and meaning***"

Harvard Business Review, October 2001

A common alphabet is not enough…

- "XML is only the first step to ensuring that computers can communicate freely. ***XML is an alphabet for computers*** and as everyone who travels in Europe knows, knowing the alphabet doesn't mean you can speak Italian or French"

Business Week, March 18, 2002

Standard glossaries can help, but...

- Defining standard vocabularies is ***difficult and time-consuming***
- Once defined, standards ***don't adapt well***
- Heterogeneous domains need a ***broad-coverage vocabulary***
- People don't implement standards correctly anyway
- Vocabulary definitions are often ***ambiguous or circular***

3. Meanings and signs

The triangle of meaning - 1

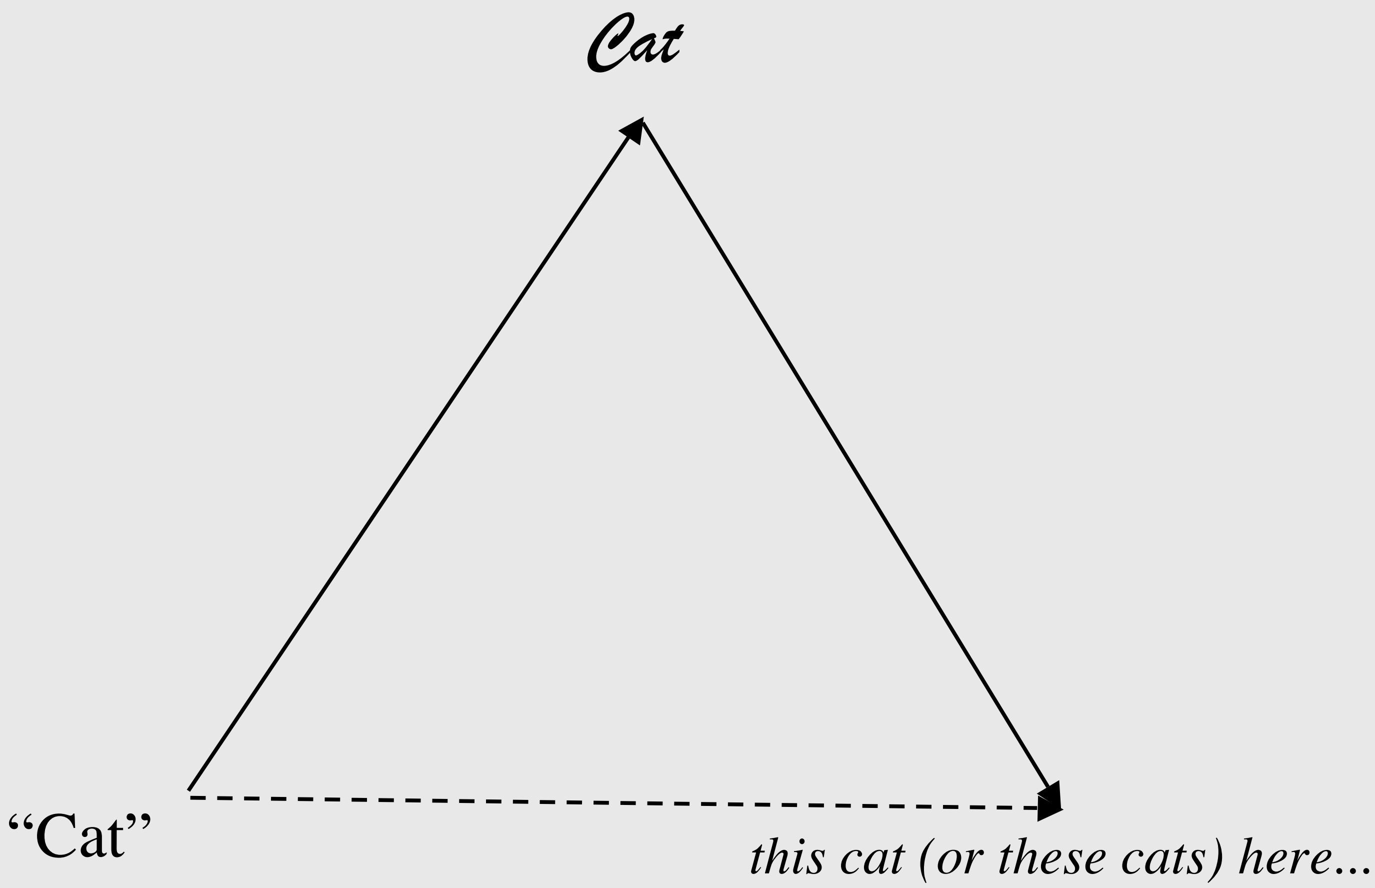

The triangle of meaning - 2

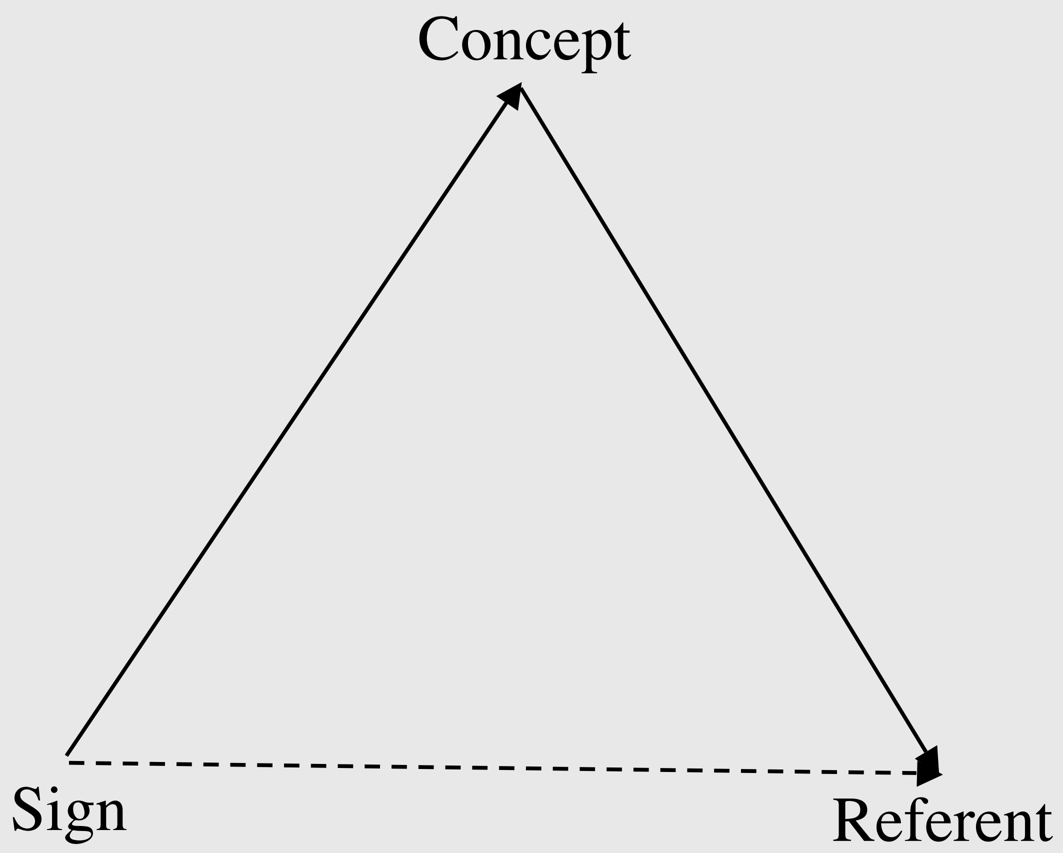

Intension ed extension

- Intension (concept): part of meaning corresponding to general principles, rules to be used to determine reference (typically, abstractions from experience)
- Extension (object): part of meaning corresponding to the effective reference
- Only by means of the ***concept*** associated to the ***sign*** “cat” we can correctly ***interpret*** this sign in various ***situations***
- The sign’s referent is the result of this interpretation
- Such interpretation is a ***situated intentional act***

An example: the concept of *red*

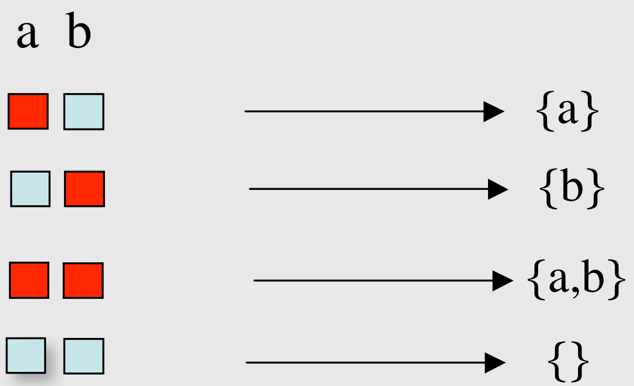

Concepts, properties, and relations

- By means of concepts, we ***ascribe*** properties and relations to things. We can say that concepts *describe* properties or relations.
- Concepts describing relations are also called ***conceptual relations:***
  - *friend-of, father-of...*
- Conceptual relations are NOT sets of tuples! Their extension is a set of tuples.

From experience to *conceptualization*

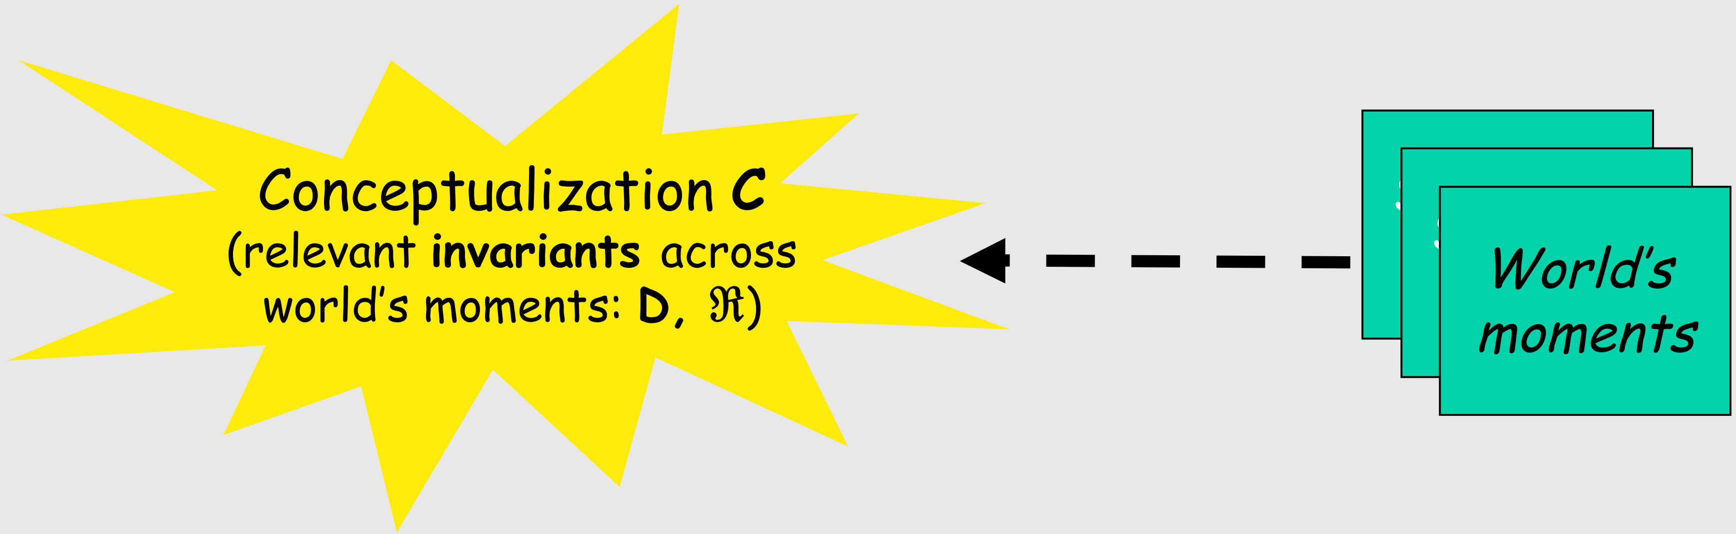

D : cognitive domain

𝔑 : set of *conceptual relations* on elements of D

What is a *conceptualization?* A cognitive approach

- Humans isolate ***relevant invariances*** from physical reality (quality distributions) on the basis of:
  - Perception (as resulting from evolution)
  - Cognition and cultural experience
  - Language

- A set of ***atomic stimuli*** (input pattern) is associated to each world's moment

- Synchronic level: ***spatial invariants***
  - Unity properties are *ascribed* to input patterns: topological and morphological wholes (***percepts***) emerge

- Diachronic level: ***temporal invariants***
  - ***Objects***: equivalence relationships among **percepts** belonging to different moments
  - ***Events:*** unity properties are ascribed to ***percept sequences*** belonging to different moments

4. What is an ontology

Ontology, lexicon, semantics

- Distinctions among contents: ***Ontology***
- Reference to content: ***Lexicon,*** via ***Semantics***

- Every organization, every computer system
  - Makes (implicit) ontologic assumptions
  - Adopt a certain lexicon, to which an ***intended semantics*** is ascribed.

Ontology and Ontologies

- ***Ontology***: the philosophical discipline
  - Study of ***what there is***
    (content *qua* content, even independently of its existence...)
  - Study of the ***nature*** and ***structure*** of "reality"
- ***ontologies***:

Specific (theoretical or computational) artifacts
expressing the ***intended meaning*** of a ***vocabulary***
in terms of ***primitive*** categories and relations describing
the ***nature*** and ***structure*** of a ***domain of discourse***

Gruber: "Explicit and formal specifications of a ***conceptualization***"

What is a conceptualization

- Formal structure of (a piece of) reality *as perceived and organized by an agent*, *independently* of:
  - the ***vocabulary*** used
  - the actual occurence of a specific ***situation***
- Different situations involving same objects, described by different vocabularies, may share the same conceptualization.

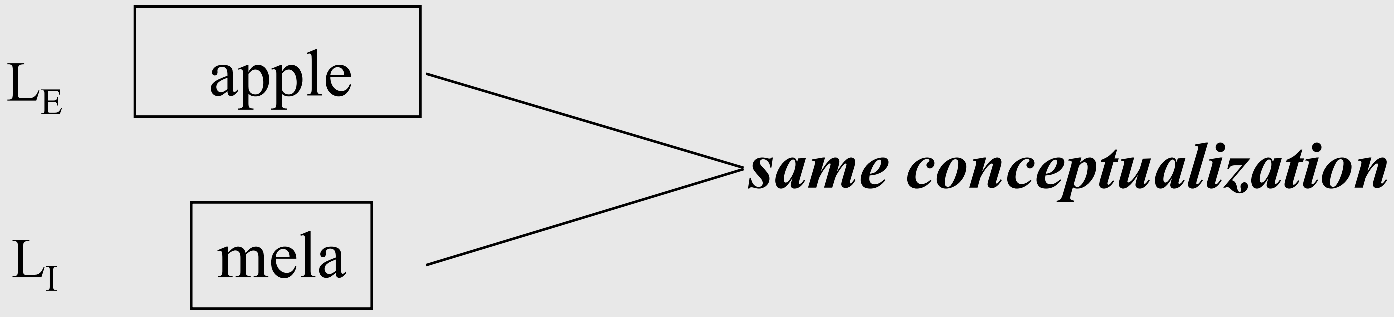

Relations vs. Conceptual Relations

ordinary relations are defined on a ***domain*** D:

$$r_n \in 2^{D^n}$$

conceptual relations are defined on a ***domain space*** <D, W>

$$\rho_n : W \rightarrow 2^{D^n}$$ *(Montague's intensional logic)*

A ***conceptualization*** for D is a tuple **C** = <D, W, ℜ>,
where ℜ is a set of conceptual relations on <D, W>

A ***model*** for a language **L** with vocabulary V is a structure

<**S**, I>, where **S** = <D, **R**> is a ***world structure*** and
I: V→D∪**R** is the usual interpretation function.

A model fixes a particular extensional interpretation of the language. Analogously, we can fix an ***intensional*** interpretation by means of a structure

<**C**, ℑ>, where **C** = <D, W, ℜ> is a conceptualization and
ℑ: V→D∪ℜ is an *intensional interpretation function*.

We call such a structure **K=<C**, ℑ> an ***ontological commitment*** for **L**.
L ***commits*** to **C** by means of **K**.
C is the ***underlying conceptualization*** of **K**.

Ontologies and intended meaning

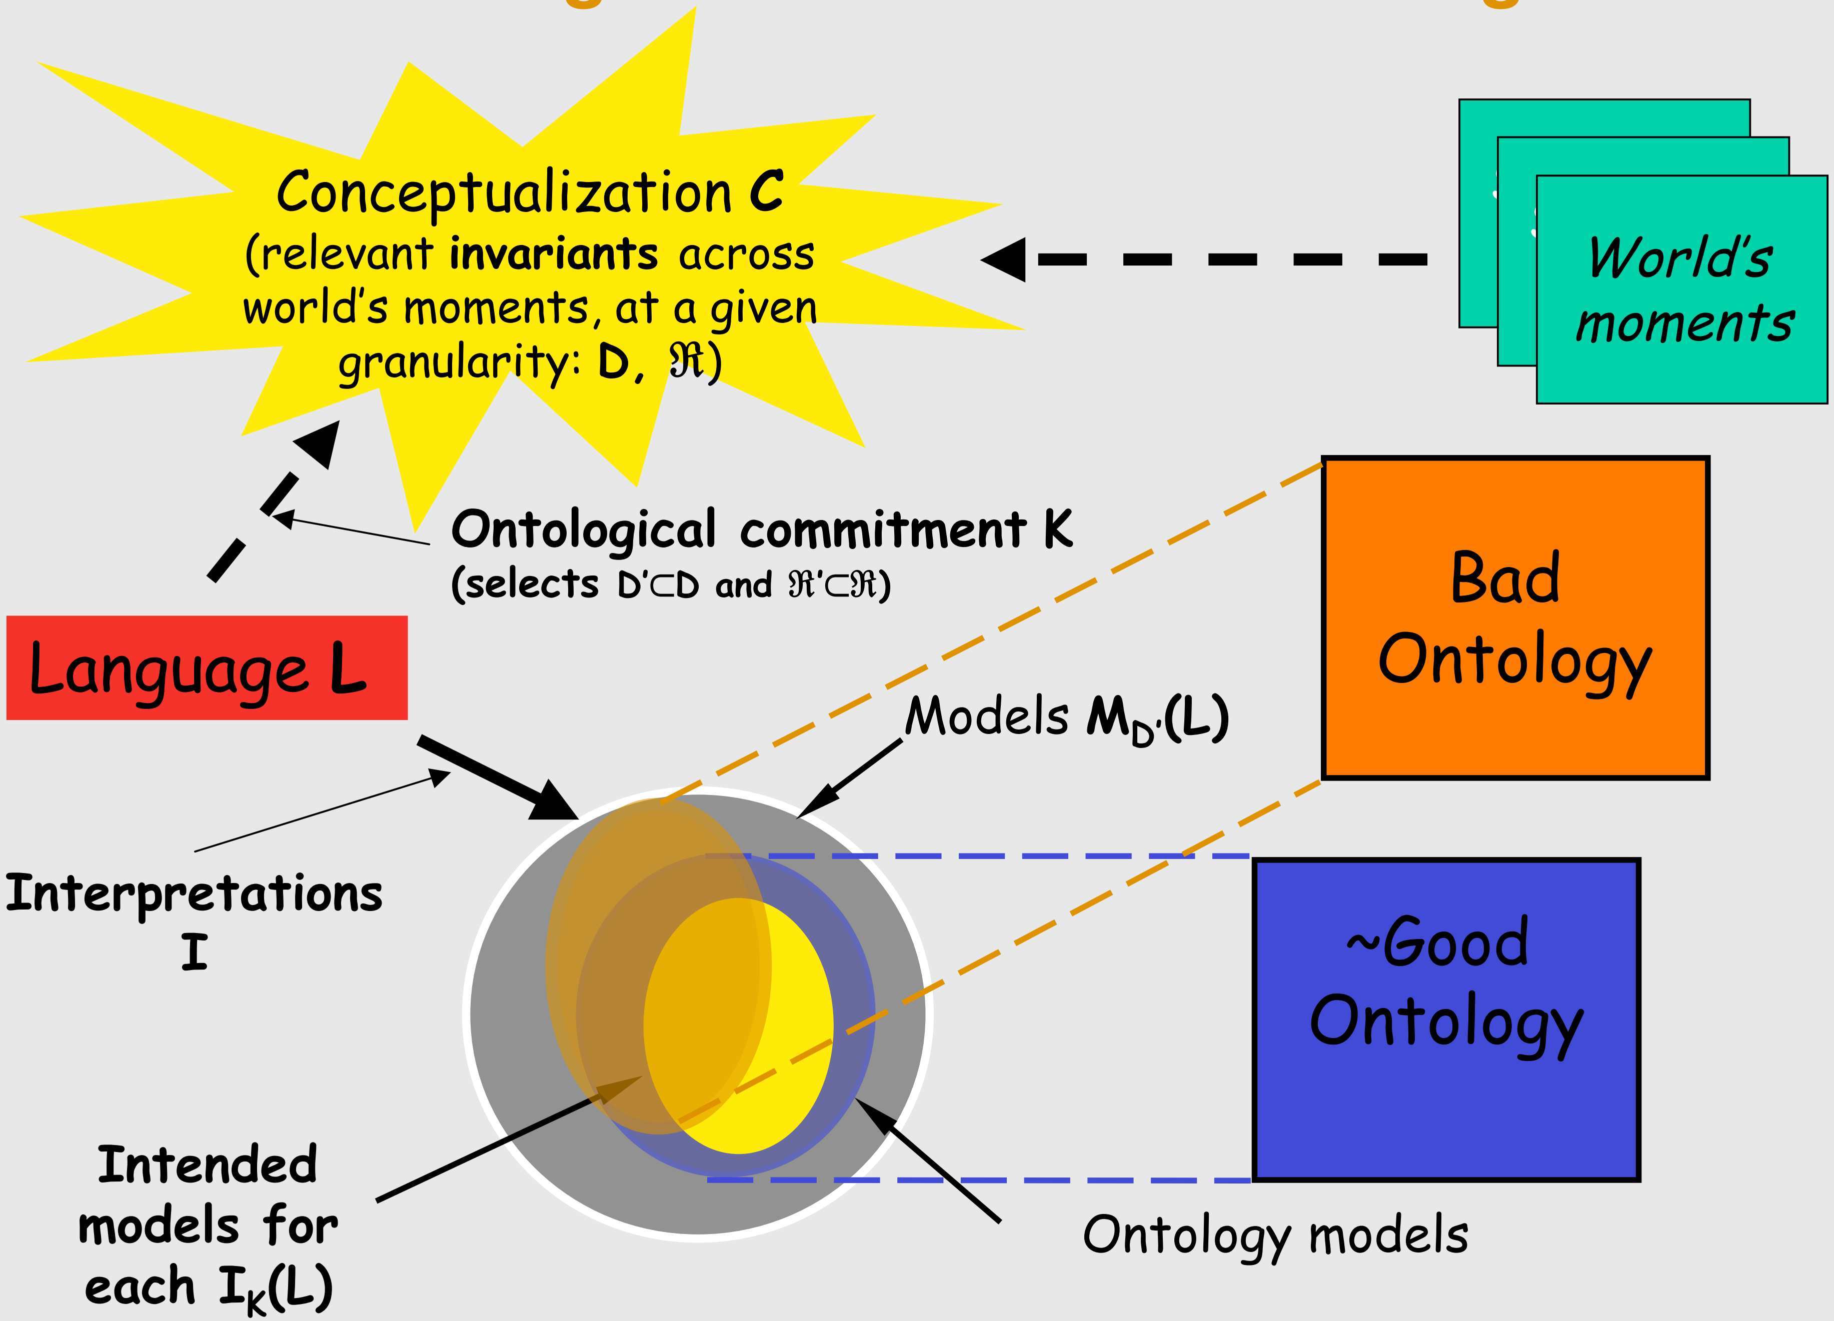

Ontology Quality: Precision and Coverage

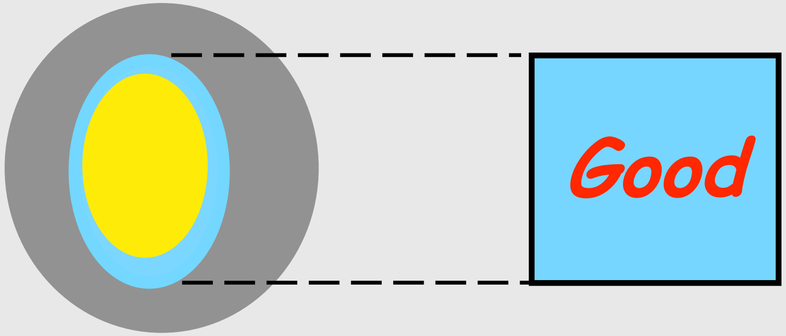

High precision, max coverage

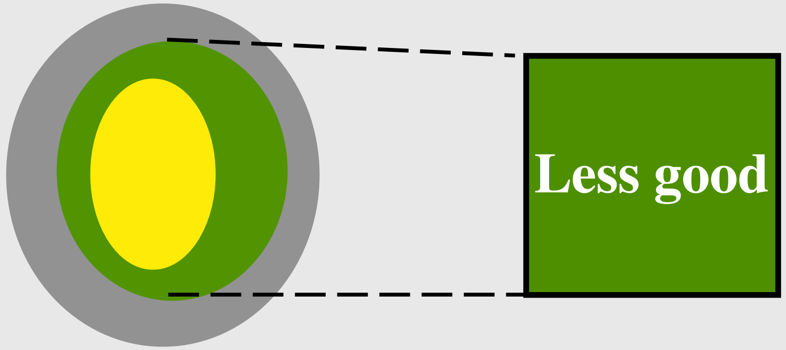

Low precision, max coverage

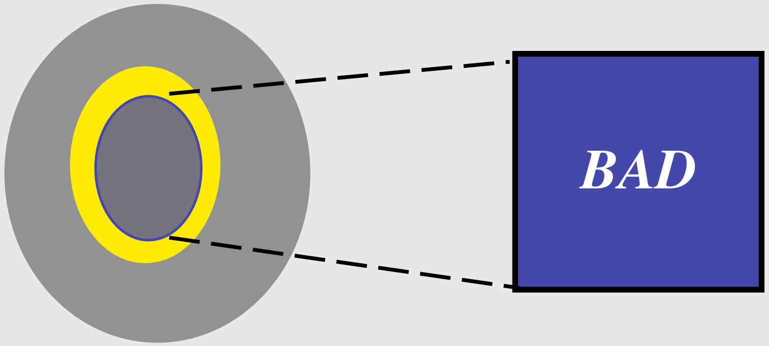

Max precision, limited coverage

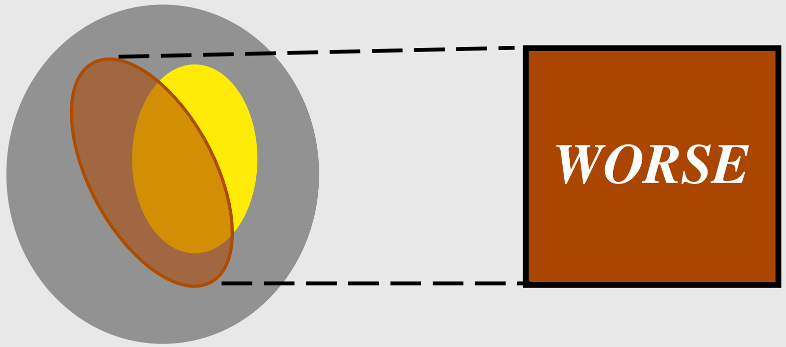

Low precision, limited coverage

Levels of Ontological Precision

Catalog

tennis
football
game
field game
court game
athletic game
outdoor game

Glossary

Taxonomy

game
  athletic game
    court game
      tennis
    outdoor game
      field game
        football

Thesaurus

game
NT athletic game
  NT court game
    RT court
    NT tennis
      RT double fault

DB/OO scheme

Axiomatic theory

$game(x) \rightarrow activity(x)$
$athletic\ game(x) \rightarrow game(x)$
$court\ game(x) \leftrightarrow athletic\ game(x) \wedge \exists y.\ played\_in(x,y) \wedge court(y)$
$tennis(x) \rightarrow court\ game(x)$
$double\ fault(x) \rightarrow fault(x) \wedge \exists y.\ part\_of(x,y) \wedge tennis(y)$

Ontological precision →

Why precision is important

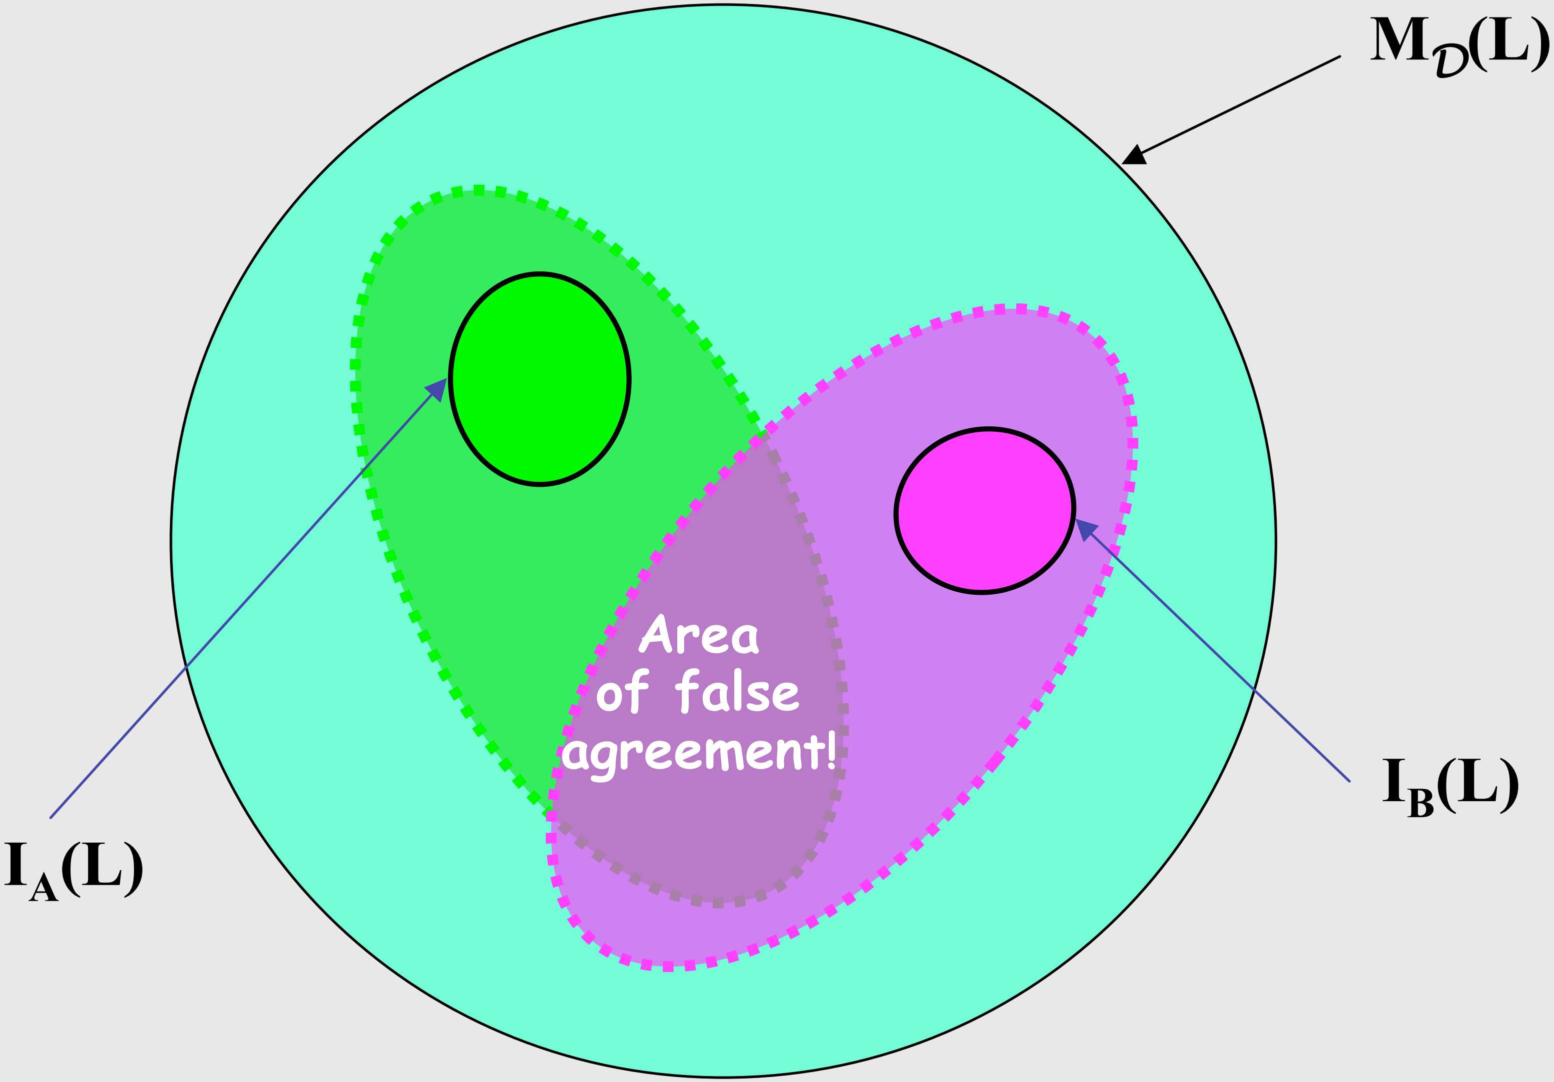

When precision is not enough

Only one binary predicate in the language: on

Only three blocks in the domain: a, b, c.

Axioms (for all x,y,z):

$$on(x,y) \rightarrow \neg on(y,x)$$

$$on(x,y) \rightarrow \neg \exists z\,(on(x,z) \wedge on(z,y))$$

Non-intended *models* are excluded, but the rules for the competent usage of on in different *situations* are not captured.

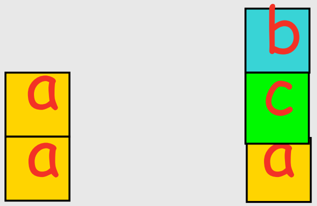

Excluded situations

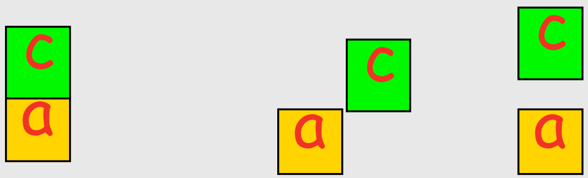

Indistinguishable situations

Precision vs. Accuracy

- In general, a single intended *model* may not discriminate among relevant alternative *situations because of*
  - Lack of *primitives*
  - Lack of *entities*
- Capturing all intended models is not sufficient for a “perfect” ontology

  Precision: non-intended *models* are excluded

  Accuracy: non-intended *situations* are excluded

When is a precise (and well-founded) ontology useful?

1. When *subtle distinctions* are important
2. When *recognizing disagreement* is important
3. When *general abstractions* are important
4. When *careful explanation and justification* of ontological commitment is important
5. When ***mutual understanding*** is more important than interoperability.

5. Ontologies and...

Ontologies and classifications

- Classifications focus on:
  - ***access***, based on pre-determined criteria (encoded by ***syntactic keys***)
- Ontologies focus on:
  - ***Meaning*** of terms
  - ***Nature*** and ***structure*** of a domain

Ontologies and Database Schemas

- Database schemas:
  - Constraints focus on ***data integrity***
  - Relationships and attribute values out of the DoD
- Ontologies:
  - Constraints focus on ***intended meaning***
  - Relationships and attribute values first class citizens

Ontologies vs. Knowledge Bases

- Knowledge base
  - Assertional component
    - **reflects** ***specific (epistemic) states of affairs***
    - designed for ***problem-solving***
    - corresponding to ***episodic memory***
  - Terminological component (*ontology*)
    - ***independent*** of particular ***states of affairs***
    - Designed to support ***terminological services***
    - corresponding to ***semantic memory***

Ontological formulas are (assumed to be)
invariant, necessary information

A single, imperialistic ontology?

- An ontology is first of all ***for understanding each other***
  - ...among people, first of all!
  - not necessarily for thinking in the same way

- A single ontology for multiple applications ***is not necessary***
  - Different applications using different ontologies can co-exist and co-operate (not necessarily inter-operate)
  - ...if linked (and compared) together **by means of a general enough basic categories and relations (*primitives*)**.

- If basic assumptions are not made explicit, any imposed, common ontology risks to be
  - seriously mis-used or misunderstood
  - opaque with respect to other ontologies

Which primitives?
The role of ontological analysis

- Theory of Essence and Identity
- Theory of Parts (Mereology)
- Theory of Wholes
- Theory of Dependence
- Theory of Composition and Constitution
- Theory of Properties and Qualities

The basis for a common ontology vocabulary

Idea of Chris Welty, IBM Watson Research Centre, while visiting our lab in 2000

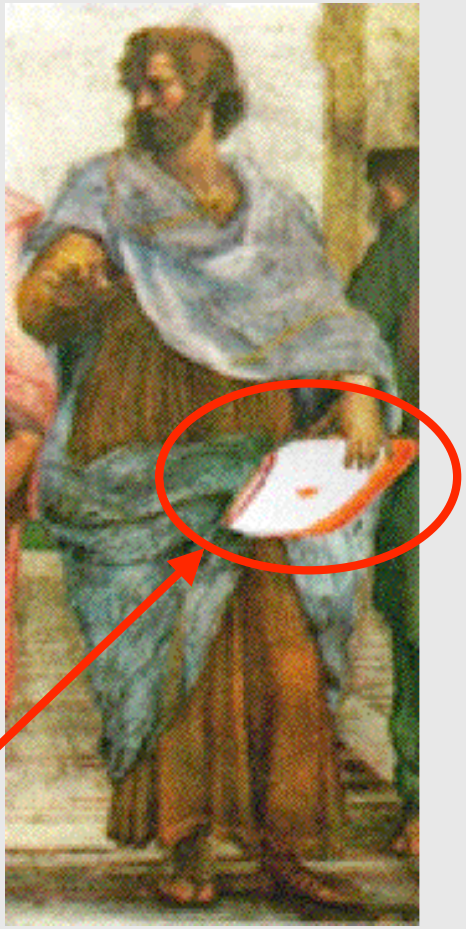

The Ontological Level

(Guarino 94)

Level	*Primitives*	*Interpretation*	*Main feature*
Logical	Predicates, functions	Arbitrary	Formalization
Epistemological	Structuring relations	Arbitrary	Structure
Ontological	**Ontological relations**	**Constrained** (meaning postulate s)	**Meaning**
Conceptual	Conceptual relations	Subjective	Conceptualization
Linguistic	Linguistic terms	Subjective	Language dependence

The semantic web architecture [Tim Berners Lee 2000]

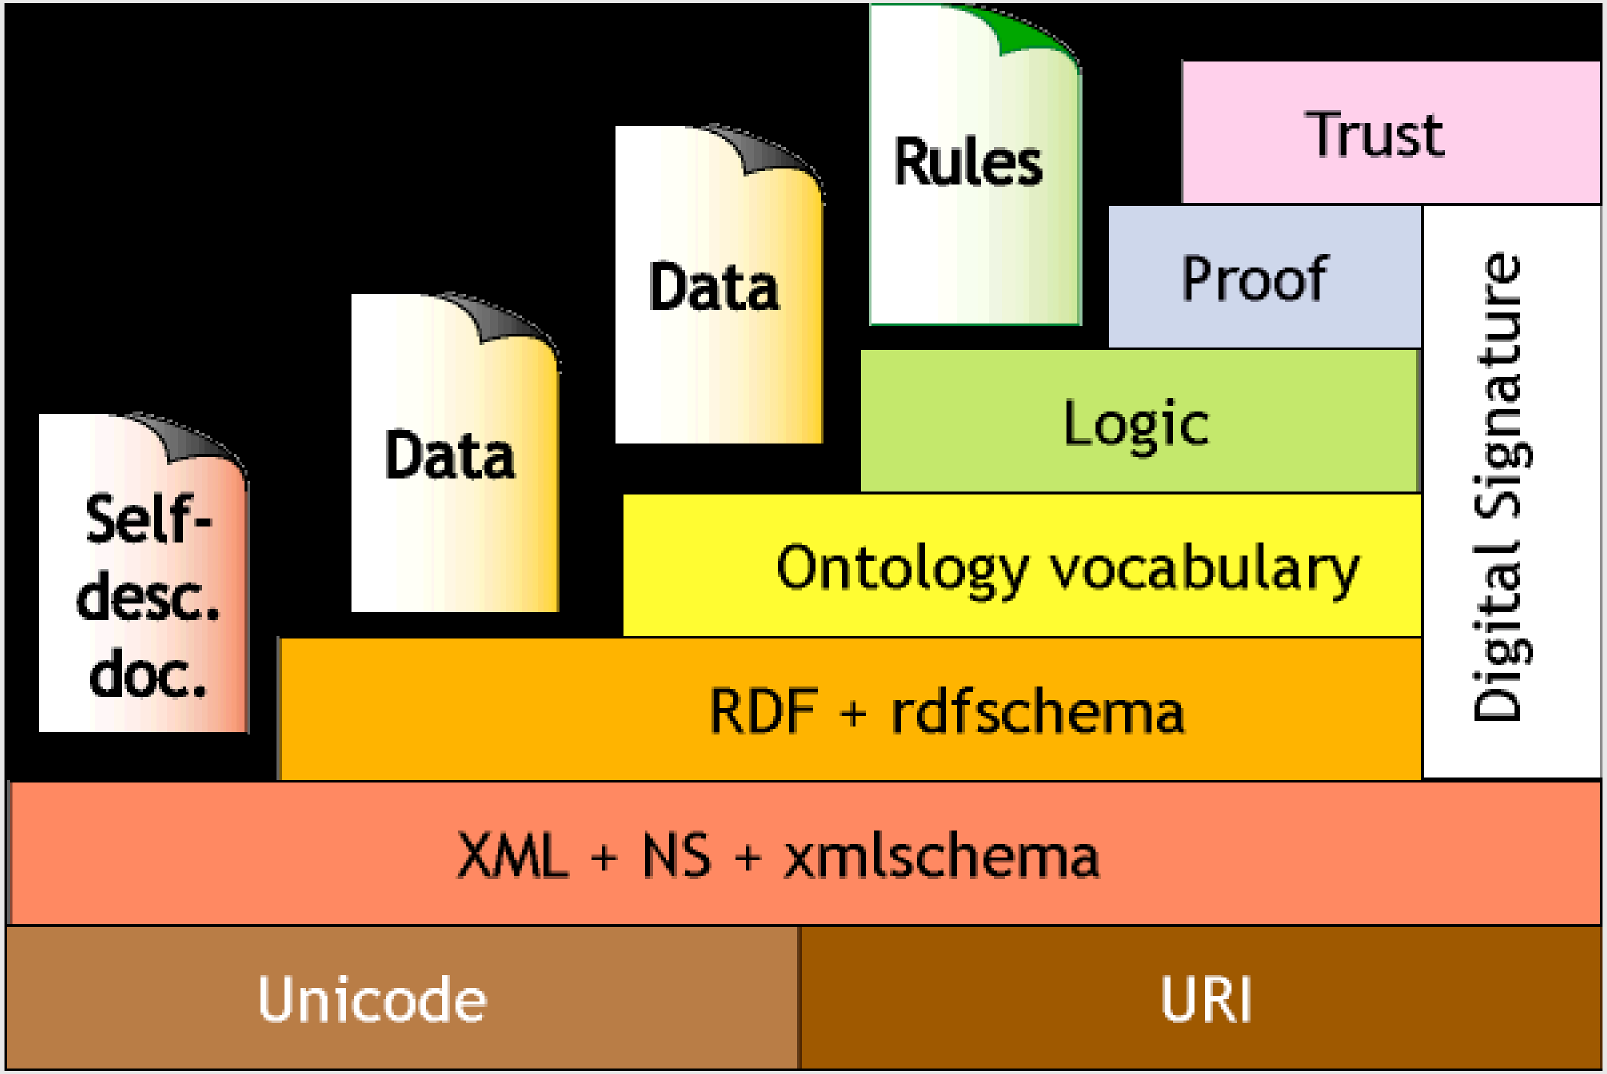

6. The tools of Formal Ontology

Formal Ontology

- Theory of ***formal distinctions and connections*** within:
  - entities of the world, as we perceive it (***particulars***)
  - categories we use to talk about such entities (***universals***)
- Why ***formal***?
  - Two meanings: ***rigorous*** and ***general***
  - Formal logic: connections between truths - neutral wrt ***truth***
  - Formal ontology: connections between things - neutral wrt ***reality***
- ***NOTE:*** **"represented in a formal language" is not enough for being formal in the above sense!**
- **(*Analytic ontology* may be a better term to avoid this confusion)**

The first steps of ontological analysis

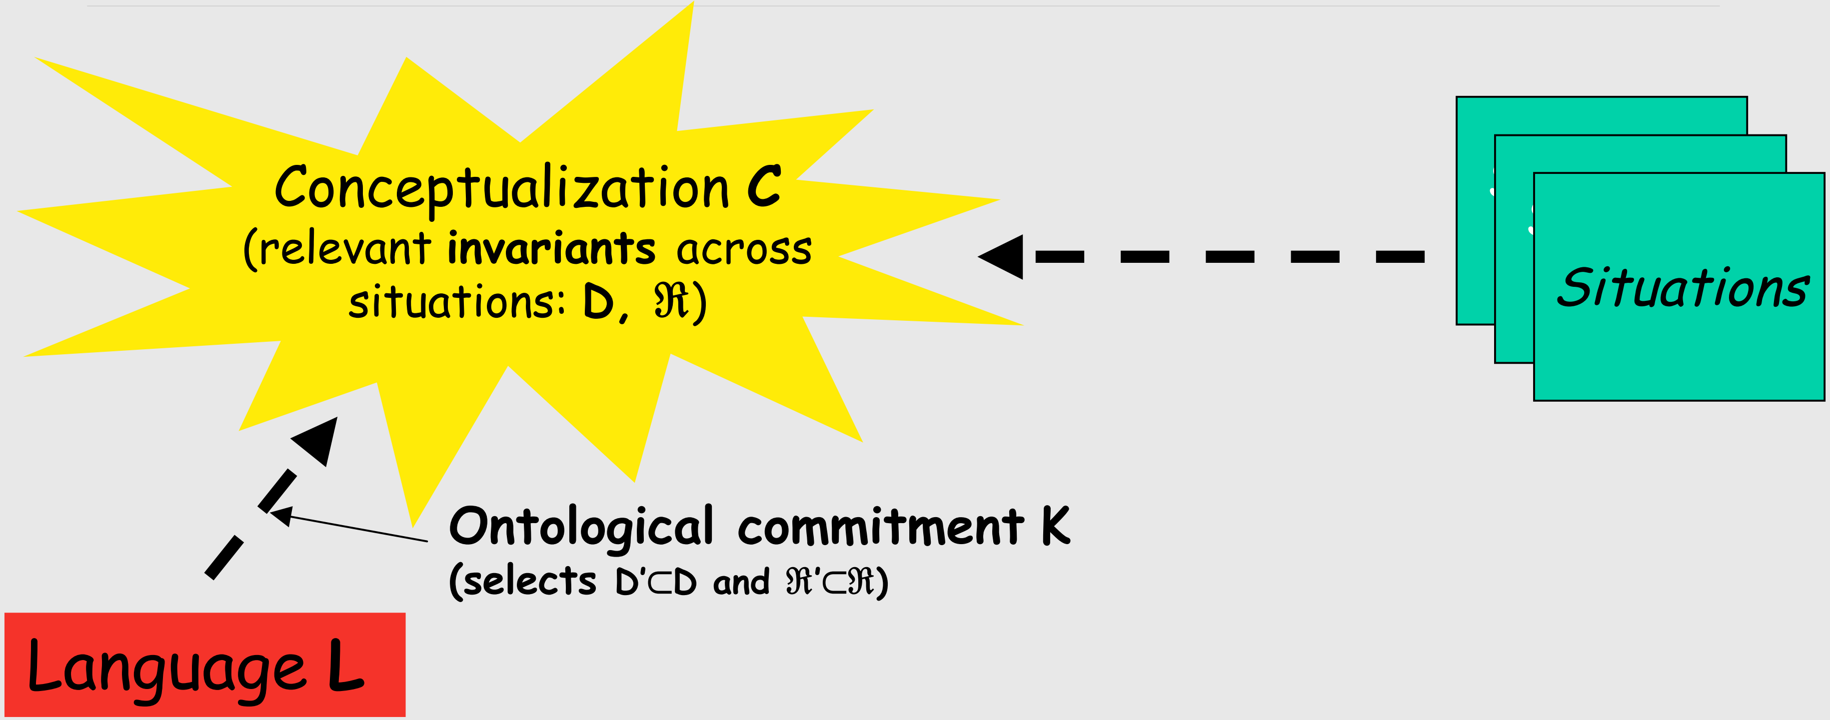

- Be clear about the ***domain of discourse*** (existence...)
- Choose the relevant ***concepts and conceptual relations***
- Choose the ***primive relations***
- Choose meaningful ***names*** for these

Essential properties

- For an individual
  - John must have a brain
  - John must be a human
  - John must be alive
- For a type
  - All human beings must have a brain
  - All human beings must be “a whole” (all of a piece)

Essential properties and rigidity

- Certain entities must have some properties in order to *exist*
  - John must have a brain
  - John must be a person.
- Certain properties are essential to ***all*** their instances (***being a person*** vs. ***being hard***).
- These properties are ***rigid*** - Their extension is the same in all possible worlds. If an entity is ever an instance of a rigid property, it must necessarily be such.

- By the way, what's the meaning of *exist*?
  - Being an element of the domain of discourse
  - Being present *at a certain time* (or in a certain world...)

Carrying essential properties

- A property P *carries* an ***informative*** essential property Q (different from P) iff Q is essential to all instances of P, **and yet Q is not rigid**:
  - Every person must have a brain.
- Compare with:
  - Every person must be a mammal.

Carrying an informative essential property implies carrying a (minimal) ***identity criterion***

Identity criteria

- *Classic formulation:*

$$\phi(x) \wedge \phi(y) \rightarrow (\rho(x,y) \leftrightarrow x = y)$$

(ϕ carries the identity criterion ρ)

- *Generalization:*

$$\phi(x,t) \wedge \phi(y,t') \rightarrow (\Gamma(x,y,t,t') \leftrightarrow x = y)$$

(synchronic: $t = t'$; diachronic: $t \neq t'$)

- In most cases, Γ is based on the ***sameness*** of certain ***characteristic features***:

$$\Gamma(x,y,t,t') = \forall z\, (\chi(x,z,t) \wedge \chi(y,z,t'))$$

- **Non-triviality condition:**
  - $\Gamma(x,y, t, t')$ must not contain an identity statement between x and y!

Heuristics for Identity

- Finding necessary *and* sufficient ICs for a given property may be ***very hard***.

- Heuristic 1: ***at least a sufficient IC***.
- Heuristic 2: ***some essential parts or qualities***
- Heuristic 3: ***some essential (non-rigid) properties***

Carrying vs. Supplying Identity

- ***Supplying*** (global) identity (**+O**)
  - Carrying an IC (or relevant essential property) that doesn't hold for *all* directly subsuming properties
- ***Carrying*** identity (**+I**)
  - Not supplying identity, while being subsumed by a property that does.
- ***Common sortal principle***: x=y -> there is a common sortal supplying their identity

- Theorem: only rigid properties supply identity

Identity Disjointness Constraint

ICs impose ***constraints*** on sortals, making their ontological nature explicit:

Properties with incompatible ICs are ***disjoint***

Examples:

- sets vs. ordered sets
- persons and passengers
- amounts of matter vs. assemblies

Example - Identity

- Is *time-interval* a subclass of *time-duration*?
  - Initial answer: yes
- IC for *time-duration*
  - Same-length
- IC for *time-interval*
  - Same start & end

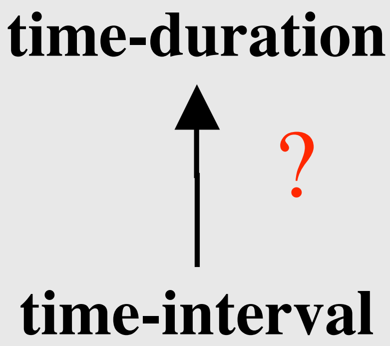

What about our rocks?

- *Igneous rock, metamorphic rock, sedimentary rock* do supply informative essential properties.
- *Large rock, grey rock, pet rock DO NOT!*
- Not all properties are the same...

Sortals and other properties

- **Sortals** (*horse, triangle, amount of matter, person, student...)*
  - Carry (non-trivial) identity conditions
  - Usually correspond to ***nouns***
  - High organizational utility
- **Non-sortals** *(red, big, old, decomposable, dependent...)*
  - No identity
  - Usually correspond to ***adjectives***
  - Span across different sortals
  - Limited organizational utility (but high semantic value)

A formal ontology of properties

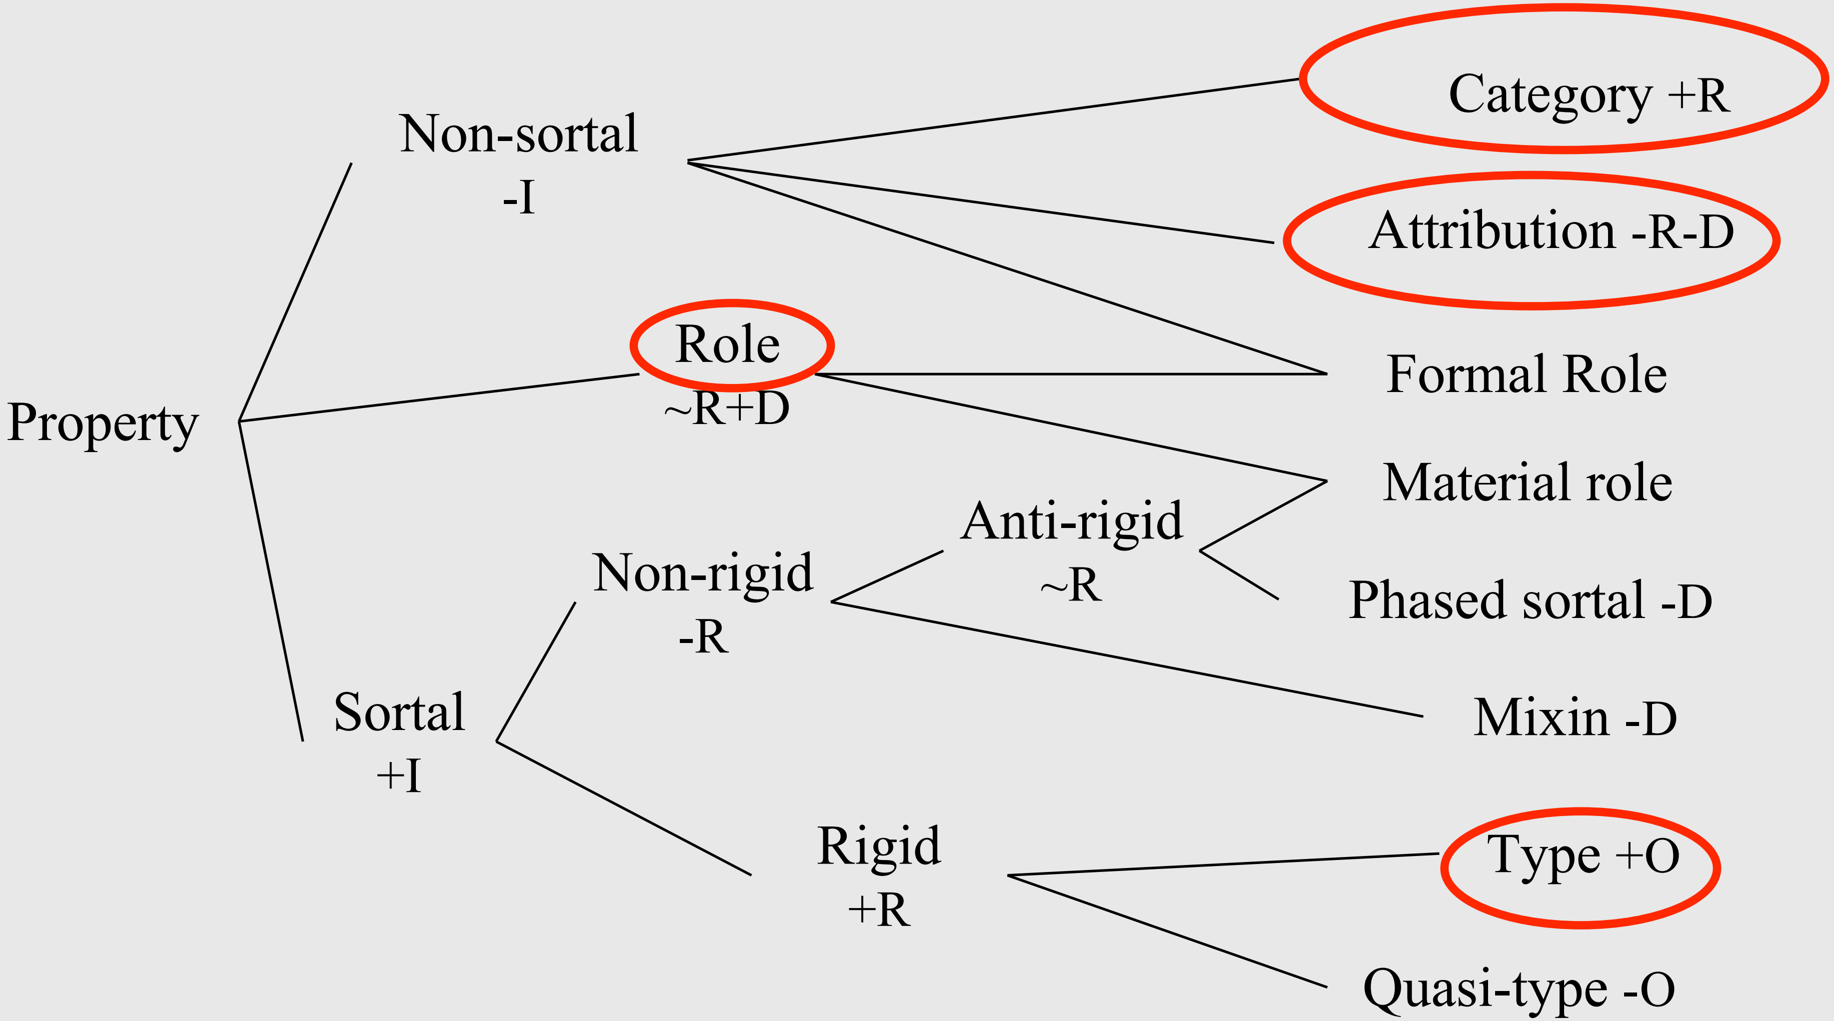

Conclusions

- Not all properties are the same
- Not all relations are the same
- Ontological distinctions *do matter*, and require to be represented *at the suitable level*
- *"...But this is hard!!"*

why should it be EASY?!

Extra slides

The task dependency problem

Representing knowledge for the purpose of solving some problem is strongly affected by the nature of the problem and the inference strategy to be applied to the problem.

[Bylander & Chandrasekaran 1988]

Making task dependence explicit [Guarino 97]

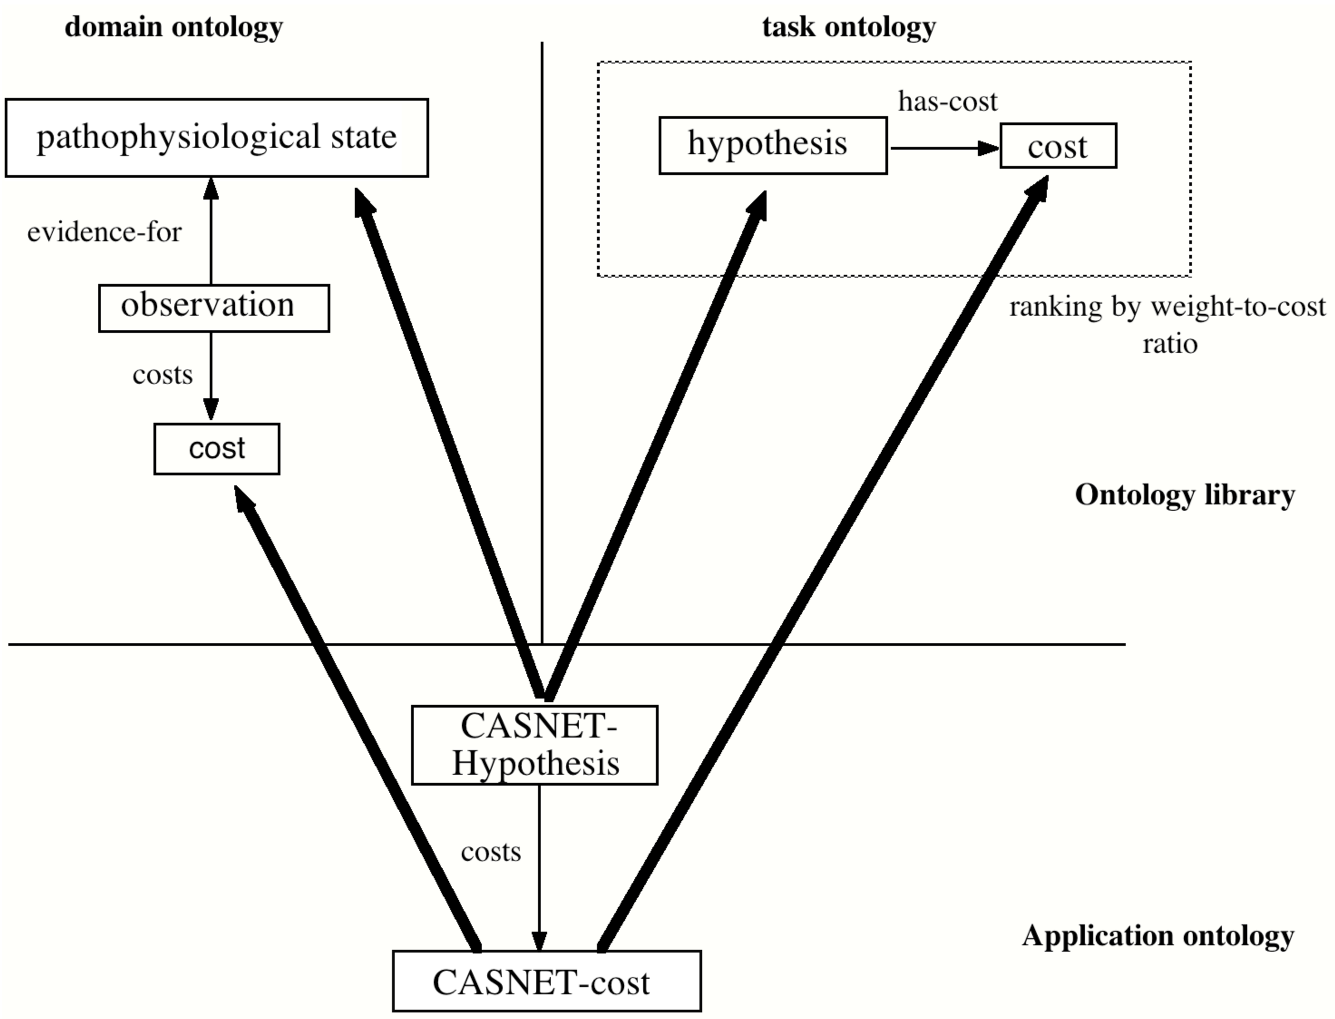

The role of task* ontologies

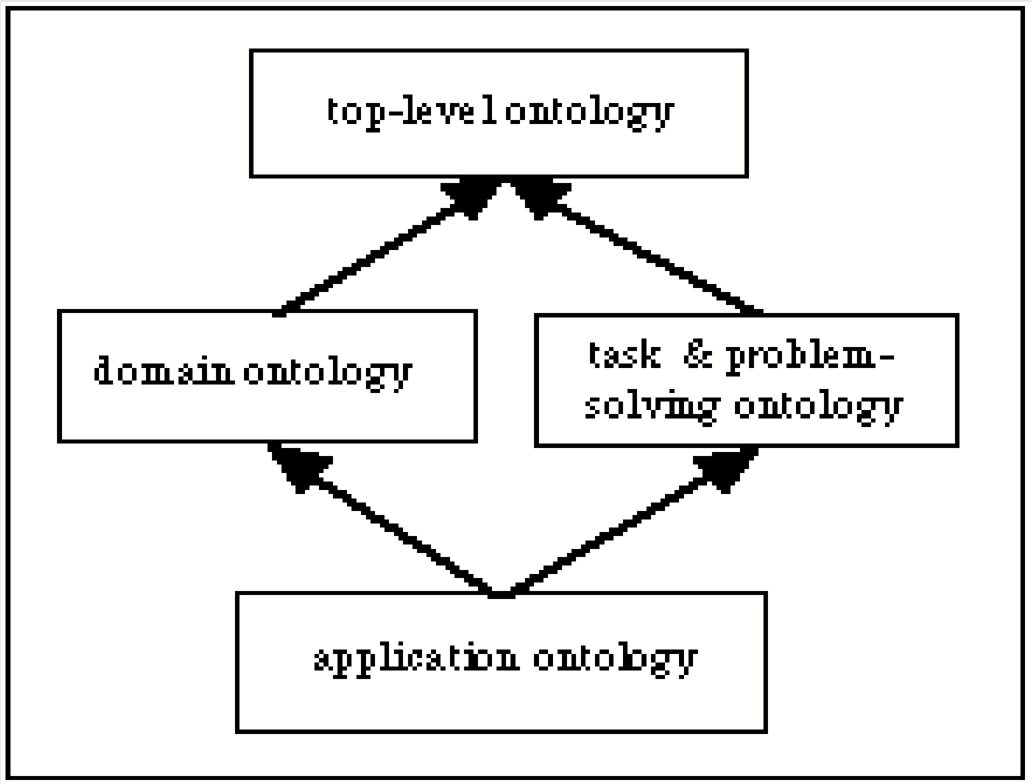

Without explicit domain and task ontologies, semantic interoperability among application ontologies is a myth!

(*) *substitute "task" with "service" if you want to be trendy…*

Unity, Identity, and Essence

- **Unity**: is the collar part of my dog?
  - ***Being a whole*** (of a certain kind) is also a (relevant) essential property

- **Identity**: is this my dog?
  - Essential properties of ***dogs***
  - Essential properties of ***my dog***

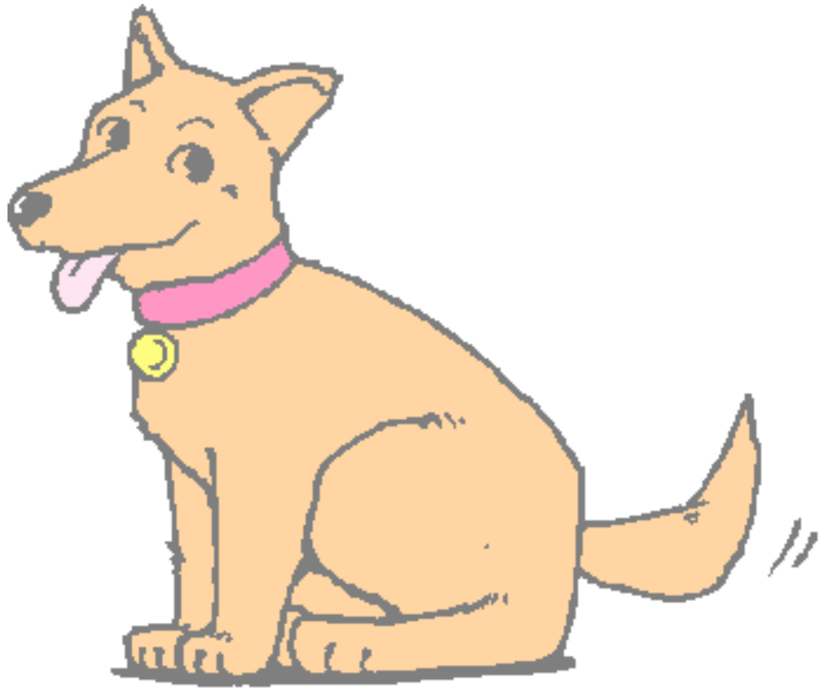

Kinds of Whole

- Depending on the **nature of the *unifying relation***, we can distinguish:
  - ***Topological wholes*** (a piece of coal, a heap of coal)
  - ***Morphological wholes*** (a constellation)
  - ***Functional wholes*** (a hammer, a bikini)
  - ***Social wholes*** (a population)

* a whole can have ***parts that are themselves wholes*** (with a different unifying relation)

Unity and Plurality

- *Ordinary objects:* ***wholes or sums of wholes***
  - *Singular*: ***no wholes as proper parts***
  - *Plural*: ***sums of wholes***
    - *Plural wholes* (the sum is ***also a whole***)
    - *Collections* (the sum is not a whole)

Mereology as an example of formal ontological analysis

- Primitive: ***proper part-of*** relation (PP)
  - asymmetric
  - transitive
  - **$Pxy =_{def} PPxy \vee x=y$**
  - **$Oxy =_{def} \exists z(Pzx \wedge Pzy)$**
- Axioms:

supplementation: $PPxy \rightarrow \exists z\,(PPzy \wedge \neg Ozx)$

principle of sum: $\exists z\, \forall w\, (Owz \leftrightarrow (Owx \vee Owy))$

extensionality: $x = y \leftrightarrow \forall w (Pwx \leftrightarrow Pwy)$

Excluded models:

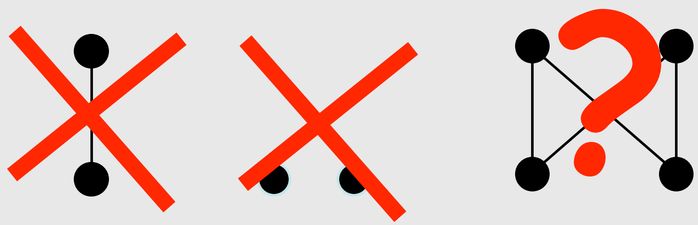

Part, Constitution, and Identity

- ***Parts not enough*** *to make the whole:* structure changes identity
- ***Mereological extensionality*** is lost
- ***Constitution*** links the two entities
- **Constitution is** asymmetric (implies ***dependence***)

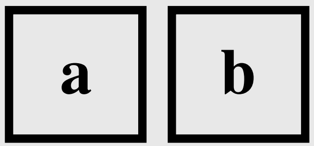

Two blocks

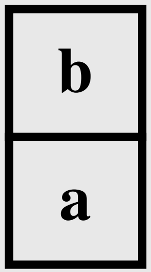

A castle