# What Should the World-Wide Mind Believe? Knowledge and Uncertainty at a Global Scale

David Poole

Department of Computer Science,
University of British Columbia

May 2010

For when I am presented with a false theorem, I do not need to examine or even to know the demonstration, since I shall discover its falsity *a posteriori* by means of an easy experiment, that is, by a calculation, costing no more than paper and ink, which will show the error no matter how small it is. . .

And if someone would doubt my results, I should say to him: "Let us calculate, Sir," and thus by taking to pen and ink, we should soon settle the question.

—Gottfried Wilhelm Leibniz [1677]

# Questions considered

- [Grad student:] All of the easy problems have been solved; how do I choose a thesis?

# Questions considered

- [Grad student:] All of the easy problems have been solved; how do I choose a thesis?
- Search engines give me the results of my query, but why should I believe these answers?

# Questions considered

- [Grad student:] All of the easy problems have been solved; how do I choose a thesis?
- Search engines give me the results of my query, but why should I believe these answers?
- The semantic web is supposed to make human knowledge accessible to computers, but how can we evaluate that knowledge and go beyond the sum of human knowledge?

# Questions considered

- [Grad student:] All of the easy problems have been solved; how do I choose a thesis?
- Search engines give me the results of my query, but why should I believe these answers?
- The semantic web is supposed to make human knowledge accessible to computers, but how can we evaluate that knowledge and go beyond the sum of human knowledge?
- There seems to be two branches of AI (Machine learning/uncertainty and the KR/logic/ontology); why do we have to choose one?

# Questions considered

- [Grad student:] All of the easy problems have been solved; how do I choose a thesis?
- Search engines give me the results of my query, but why should I believe these answers?
- The semantic web is supposed to make human knowledge accessible to computers, but how can we evaluate that knowledge and go beyond the sum of human knowledge?
- There seems to be two branches of AI (Machine learning/uncertainty and the KR/logic/ontology); why do we have to choose one?
- What will AI and the web look like in 2025?

# History of AI — a perspective from 2025

- Semantic web has evolved into the world-wide mind (WWM) — a distributed repository of all knowledge, backed up by the best science available.
- The world-wide mind doesn't just accept new knowledge but critically evaluates it and generates new knowledge.
- Scientists freed from mundane data analysis, develop new hypotheses, interesting questions, and observational data.
- World-wide mind is the expert on all questions of truth and makes the best predictions. (Using hypotheses provided by a mix of humans and machine learning).
- Public discourse on values (utilities) to determine the best course of actions for individuals, organizations and society.

# Finding information; e.g. diagnosis from symptoms

| 2010 | 2025 |
|---|---|
| • need to guess keywords; re-guess until exhaustion | • keywords + context + ontologies → unambiguous query |

# Finding information; e.g. diagnosis from symptoms

| 2010 | 2025 |
| --- | --- |
| • need to guess keywords; re-guess until exhaustion<br><br>• what information found is based on popularity and/or appeal to authority | • keywords + context + ontologies → unambiguous query |

# Finding information; e.g. diagnosis from symptoms

| 2010 | 2025 |
|---|---|
| • need to guess keywords; re-guess until exhaustion | • keywords + context + ontologies → unambiguous query |
| • what information found is based on popularity and/or appeal to authority | • information based on best evidence available in world |

# Finding information; e.g. diagnosis from symptoms

| 2010 | 2025 |
| --- | --- |
| • need to guess keywords; re-guess until exhaustion | • keywords + context + ontologies → unambiguous query |
| • what information found is based on popularity and/or appeal to authority | • information based on best evidence available in world |
| • verify information based on other sites (with different wording) | • information justified by presenting the evidence for and against it |

# Finding information; e.g. diagnosis from symptoms

| 2010 | 2025 |
|---|---|
| • need to guess keywords; re-guess until exhaustion | • keywords + context + ontologies → unambiguous query |
| • what information found is based on popularity and/or appeal to authority | • information based on best evidence available in world |
| • verify information based on other sites (with different wording) | • information justified by presenting the evidence for and against it |
| • extract information from text and graphics to make decisions | • decisions based on evidence and utilities |

# Believing information

| 2010 | 2025 |
| --- | --- |
| • skeptics throw doubt on science and scientists say “trust us” | • data is available for all to view; all alternative hypotheses can be evaluated |

# Believing information

| 2010 | 2025 |
|---|---|
| • skeptics throw doubt on science and scientists say "trust us" | • data is available for all to view; all alternative hypotheses can be evaluated |
| • politicians campaign on what is true and what they will do | • politicians campaign on their values |

# Believing information

| 2010 | 2025 |
| --- | --- |
| • skeptics throw doubt on science and scientists say "trust us" | • data is available for all to view; all alternative hypotheses can be evaluated |
| • politicians campaign on what is true and what they will do | • politicians campaign on their values |
| • food shopping is based on price and brands | • food shopping based on optimizing health and well-being (users goals and values, and known risks) |

# AI Research

| 2010 | 2025 |
| --- | --- |
| • separation of uncertainty and KR issues<br>— ML ignores ontologies<br>— rich representations ignore uncertainty | • uncertainty and ontologies are integral parts of world-wide mind |
| • semantic web in its infancy | • world wide mind being used |
| • relational representations starting to be used in ML | • rich representations with uncertainty ubiquitous |
| • learning based on one or few homogeneous data sets | • learning from all data in world |
| • data sets usable only by specialists | • data sets published, available, persistent and interoperable |

# Outline

# Notational Minefield

- Variable (probability, logic, programming languages)
- Model (science, probability, logic, fashion)
- Parameter (mathematics, statistics)
- Domain (science, logic, probability, mathematics)
- Object/class (object-oriented programming, ontologies)
- $=$ (probability, logic)
- First-order (logic, dynamical systems)

# Science is the foundation of belief

- If a KR system makes a prediction, we should ask: what evidence is there? The system should be able to provide such evidence.
- A knowledge-based system should believe based on evidence. Not all beliefs are equally valid.
- The mechanism that has been developed for judging knowledge is called science. We trust scientific conclusions because they are based on evidence.

# Science is the foundation of belief

- If a KR system makes a prediction, we should ask: what evidence is there? The system should be able to provide such evidence.
- A knowledge-based system should believe based on evidence. Not all beliefs are equally valid.
- The mechanism that has been developed for judging knowledge is called science. We trust scientific conclusions because they are based on evidence.
- The semantic web is an endeavor to make all of the world's knowledge accessible to computers.
- We have used to term semantic science, in an anaolgous way to the *semantic web*.
- Claim: semantic science will form the foundation of the world-wide mind.

# Science as the foundation of world-wide mind

I mean *science* in the broadest sense:

- where and when landslides occur
- where to find gold
- what errors students make
- disease symptoms, prognosis and treatment
- what companies will be good to invest in
- what apartment Mary would like
- which celebrities are having affairs

# Semantic Science

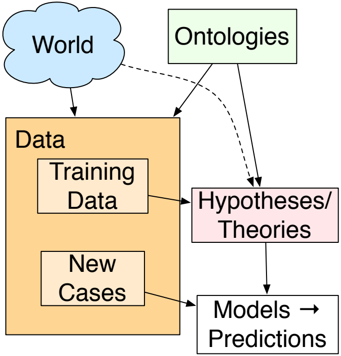


- Ontologies represent the meaning of symbols.
- Data that adheres to an ontology is published.
- Hypotheses that make (probabilistic) predictions on data are published.
- Data used to evaluate hypotheses; the best hypotheses are theories.
- Theories form models for predictions on new cases.
- All evolve in time.

# Outline

1. Semantic Science Overview
   - Ontologies
   - Data
   - Hypotheses and Theories
   - Models
2. Levels of Semantic Science
   - Feature-based Theories
   - Relational Domains
   - Probabilities with Ontologies
   - Existence and Identity Uncertainty

# Ontologies

- In philosophy, ontology the study of existence.
- In CS, an ontology is a (formal) specification of the meaning of the vocabulary used in an information system.
- Ontologies are needed so that information sources can inter-operate at a semantic level.

# Ontologies

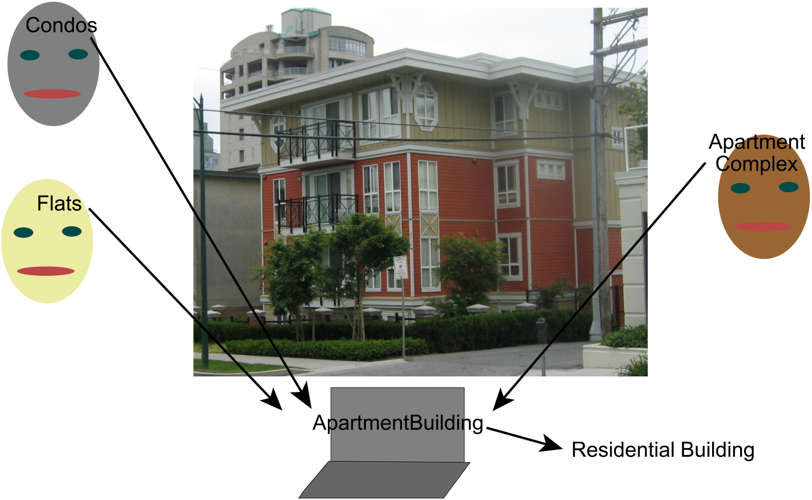
Condos
Flats
Apartment Complex
ApartmentBuilding
Residential Building

# Main Components of an Ontology

- Individuals: the objects in the world (not usually specified as part of the ontology)
- Classes: sets of (potential) individuals
- Properties: between individuals and their values

$\langle Individual, Property, Value \rangle$ triples are universal representations of relations.

# Semantic Web Ontology Languages

- URI — universal resource identifier; everything is a resource
- RDF — language for triples in XML
- RDF Schema — define resources in terms of each other: class, type, subClassOf, subPropertyOf, collections. . .
- OWL — defines vocabulary for equality, restricting domains and ranges of properties, transitivity, cardinality. . .
- OWL-Lite, OWL-DL, OWL-Full

# Aristotelian definitions

Aristotle [350 B.C.] suggested the definition if a class $C$ in terms of:

- Genus: the super-class
- Differentia: the attributes that make members of the class $C$ different from other members of the super-class

*"If genera are different and co-ordinate, their differentiae are themselves different in kind. Take as an instance the genus 'animal' and the genus 'knowledge'. 'With feet', 'two-footed', 'winged', 'aquatic', are differentiae of 'animal'; the species of knowledge are not distinguished by the same differentiae. One species of knowledge does not differ from another in being 'two-footed'."*

Aristotle, *Categories*, 350 B.C.

# An Aristotelian definition

- An **apartment building** is a **residential building** with **multiple units** and **units are rented**.

$$
\begin{aligned}
ApartmentBuilding &\equiv ResidentialBuilding \& \\
&\quad NumUnits = many \& \\
&\quad Ownership = rental
\end{aligned}
$$

  *NumUnits* is a property with domain *ResidentialBuilding* and range $\{one, two, many\}$
  *Ownership* is a property with domain *Building* and range $\{owned, rental, coop\}$.

- All classes are defined in terms of properties.

# Outline

# Data

Real data is messy!

- Multiple levels of abstraction
- Multiple levels of detail
- Uses the vocabulary from many ontologies: rocks, minerals, top-level ontology,. . .
- Rich meta-data:
  - Who collected each datum? (identity and credentials)
  - Who transcribed the information?
  - What was the protocol used to collect the data? (Chosen at random or chosen because interesting?)
  - What were the controls — what was manipulated, when?
  - What sensors were used? What is their reliability and operating range?

# Example Data, Geology

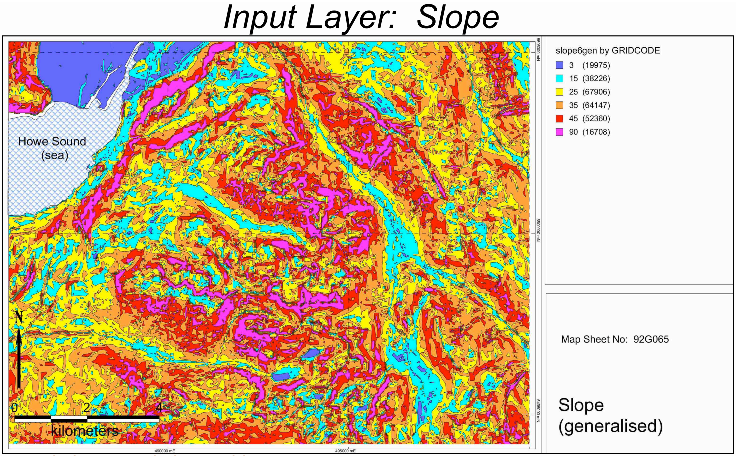

# Example Data, Geology

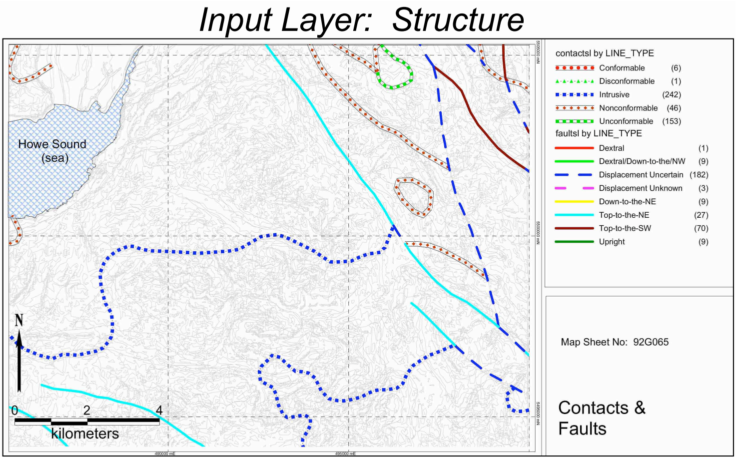

# http://www.vsto.org/

### Welcome to the Virtual Solar Terrestrial Observatory

The Virtual Solar Terrestrial Observatory (VSTO) is a unified semantic environment serving data from diverse data archives in the fields of solar, solar-terrestrial, and space physics (SSTSP), currently:

- Upper atmosphere data from the **CEDAR** (Coupling, Energetics and Dynamics of Atmospheric Regions) archive
- Solar corona data from the **MLSO** (Mauna Loa Solar Observatory) archive

The VSTO portal uses an underlying ontology (i.e. an organized knowledge base of the SSTSP domain) to present a general interface that allows selection and retrieval of products (ascii and binary data files, images, plots) from heterogenous external data services.

**VSTO Data Access**

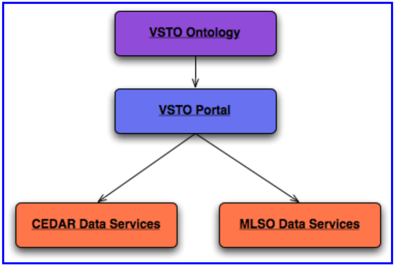

# Data is theory-laden

- Sapir-Whorf Hypothesis [Sapir 1929, Whorf 1940]: people's perception and thought are determined by what can be described in their language. (Controversial in linguistics!)
- A stronger version for information systems:

  *What is stored and communicated by an information system is constrained by the representation and the ontology used by the information system.*

- Ontologies must come logically prior to the data.
- Data can't make distinctions that can't be expressed in the ontology.
- Different ontologies result in different data.

# Outline

# Hypotheses make predictions on data

Hypotheses are procedures that make prediction on data.
Theories are hypotheses that best fit the observational data.

- Hypotheses can make various predictions about data:
  - definitive predictions
  - point probabilities
  - probability ranges
  - ranges with confidence intervals
  - qualitative predictions
- For each prediction type, we need ways to judge predictions on data
- Users can use whatever criteria they like to evaluate theories (e.g., taking into account simplicity and elegance)
- Semantic science search engine: extract theories from published hypotheses.

# Example Prediction from a Theory

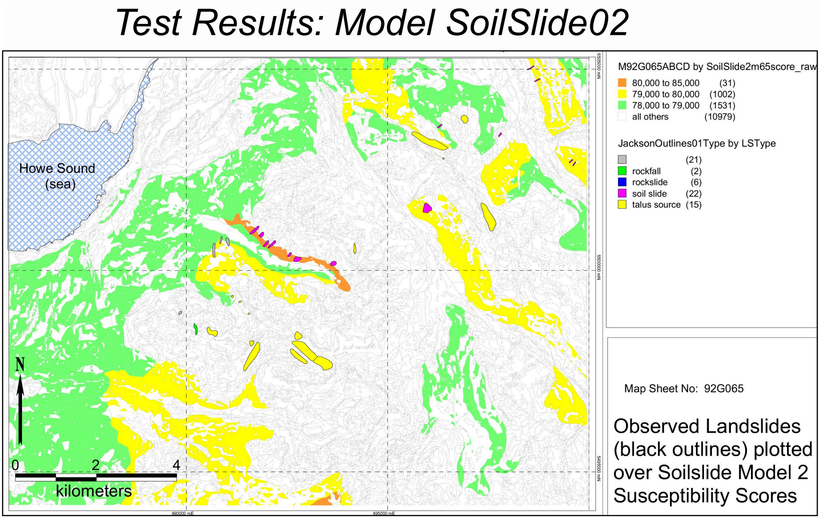

# Applying theories to new cases

- How can we compare theories that differ in their generality?
- Theory $A$ makes predictions about all cancers.
  Theory $B$ makes predictions about lung cancers.
  Should the comparison between $A$ and $B$ take into account $A$'s predictions on non-lung cancer?

# Applying theories to new cases

- How can we compare theories that differ in their generality?
- Theory $A$ makes predictions about all cancers.
  Theory $B$ makes predictions about lung cancers.
  Should the comparison between $A$ and $B$ take into account $A$'s predictions on non-lung cancer?
- What about $C$: *if lung cancer, use $B$'s prediction, else use $A$'s prediction*?

# Applying theories to new cases

- How can we compare theories that differ in their generality?
- Theory $A$ makes predictions about all cancers.
  Theory $B$ makes predictions about lung cancers.
  Should the comparison between $A$ and $B$ take into account $A$'s predictions on non-lung cancer?
- What about $C$: *if lung cancer, use $B$'s prediction, else use $A$'s prediction*?
- A model is a set of theories applied to a particular case.
  - Judge theories by how well they fit into models.
  - Models can be judged by simplicity.
  - Theory designers don't need to game the system by manipulating the generality of theories

# Dynamics of Semantic Science

- New data and hypotheses are continually added.
- Anyone can design their own ontologies.
  — People vote with their feet what ontology they use.
  — Need for semantic interoperability leads to ontologies with mappings between them.
- Ontologies evolve with theories:
  A theory hypothesizes unobserved features or useful distinctions
  $\longrightarrow$ add these to an ontology
  $\longrightarrow$ other researchers can refer to them
  $\longrightarrow$ reinterpretation of data
- Ontologies can be judged by the predictions of the theories that use them
  — role of a vocabulary is to describe useful distinctions.

# Outline

1. Semantic Science Overview
   - Ontologies
   - Data
   - Hypotheses and Theories
   - Models

2. Levels of Semantic Science
   - Feature-based Theories
   - Relational Domains
   - Probabilities with Ontologies
   - Existence and Identity Uncertainty

# Levels of Semantic Science

0. Deterministic semantic science where all of the theories make definitive predictions.
1. Feature-based semantic science, with non-deterministic predictions about feature values of data.
2. Relational semantic science, with predictions about the properties of objects and relationships among objects.
3. First-order semantic science, with predictions about the existence of objects, universally quantified statements and relations.

# Feature-Based Semantic Science

- World is described in terms of features and values.
- Random variables / features correspond to properties.
- Random variables / features are not defined in all contexts.
- Aristotelian definitions: each class is defined in terms of
  - genus (superclass) and
  - differentia (property restrictions that distinguish this class).
- Conditioning on a class means observing its differentia are true

# Partial Ontology in OWL

```
DataPropertyDomain(HasLump person)
DataPropertyRange(HasLump xsd:boolean)
EquivalentClasses(lump DataHasValue(hasLump true))
DataPropertyDomain(CancerousLump lump)
DataPropertyRange(CancerousLump xsd:boolean)
SubClassOf(DataHasValue(CancerousLump true)
           personWithCancer)
ObjectPropertyDomain(LumpShape lump)
ObjectPropertyRange(LumpShape ShapesOfLumps)
EquivalentClasses(ShapesOfLumps
      ObjectOneOf(circular oblong irregularShape))
```

# Data

Data is of observations of a world.
Meta-data about observations includes:

- The context in which the data was collected.
- The features that were controlled for
- The features that were observed

# Example Data

person visiting doctor:

| Age | Sex | Coughs | HasLump |
|---|---|---|---|
| 23 | male | true | true |
| ... | ... | ... | ... |

lump for person visiting doctor:

| Location | LumpShape | Colour | CancerousLump |
|---|---|---|---|
| leg | oblong | red | false |
| ... | ... | ... | ... |

person with cancer:

| HasLungCancer | Treatment | Age | Outcome | Months |
|---|---|---|---|---|
| true | chemo | 77 | dies | 7 |
| ... | ... | ... | ... | ... |

# Theories

A theory makes predictions about some feature values.
A theory includes:

- A context $c$ in which specifies when it can be applied.
- A set of input features $\overline{I}$ about which it does not make predictions (can include interventional variables)
- A set of output features to predict (as a function of the input features).
- A program to compute the output from the input.

Represents:

$$P(\overline{O}|c,\overline{I})$$

or perhaps

$$P(\overline{O}|c,\overline{I_{obs}},do(\overline{I_{do}}))$$

# Example

Consider the following theories:

- $T_1$ predicts the prognosis of people with lung cancer.
- $T_2$ predicts the prognosis of people with cancer.
- $T_3$ is the null hypothesis that predicts the prognosis of people in general.
- $T_4$ predicts whether people with cancer have lung cancer, as a function of coughing.
- $T_5$ predicts whether people have cancer.

What should be used to predict the prognosis of a patient with observed coughing?

# Models

To make a prediction, multiple theories need to be used together in a model.
A model consists of multiple theories, where each theory can be used to predict a subset of its output features.
A model $M$ needs to satisfy the following properties:

- $M$ is coherent: it does not rely on the value of a feature in a context where the features is not defined
- $M$ is consistent: it does not make different predictions for any feature in any context.
- $M$ is predictive: it makes a prediction in every context that is possible.
- $M$ is minimal.

# Prototype Feature-based Model

A theory instance is a tuple of the form $\langle t, c, I, O \rangle$ such that:

- $t$ is a theory,
- $c$ is a context in which the theory will be used
- $I$ is a set of inputs used by the theory
- $O$ is a set of outputs the theory will be used to predict.

A model is a set of theory instances that satisfy the previous conditions.
[Like a Bayesian belief network, but allowing for context-specific independence and avoiding undefined features.]

# Example

- $T_1$ predicts the prognosis of people with lung cancer.
- $T_2$ predicts the prognosis of people with cancer.
- $T_3$ is the null hypothesis that predicts the prognosis of people in general.
- $T_4$ predicts (probabilistically) whether people with cancer have lung cancer, as a function of coughing.
- $T_5$ predicts (probabilistically) whether people have cancer.

A possible model for $P(Lives|person \wedge coughs)$:

- $\langle T_5, person, \{\}, \{HC\}\rangle$,
- $\langle T_3, person \wedge \neg hc, \{\}, \{Lives\}\rangle$,
- $\langle T_4, person \wedge hc, \{Coughs\}, \{HLC\}\rangle$,
- $\langle T_1, person \wedge hlc, \{\}, \{Lives\}\rangle$,
- $\langle T_2, person \wedge hc \wedge \neg hlc, \{\}, \{Lives\}\rangle$.

# Outline

# Relational Learning

- Often the values of properties are not meaningful values but names of individuals.
- It is the properties of these individuals and their relationship to other individuals that needs to be learned.
- Relational learning has been studied under the umbrella of "Inductive Logic Programming" as the representations are often logic programs.

# Example: trading agent

What does Joe like?

| Individual | Property | Value |
|---|---|---|
| *joe* | *likes* | *resort*_14 |
| *joe* | *dislikes* | *resort*_35 |
| ... | ... | ... |
| *resort*_14 | *type* | *resort* |
| *resort*_14 | *near* | *beach*_18 |
| *beach*_18 | *type* | *beach* |
| *beach*_18 | *covered_in* | *ws* |
| *ws* | *type* | *sand* |
| *ws* | *color* | *white* |
| ... | ... | ... |

Values of properties may be meaningless names.

# Example: trading agent

Possible theory that could be learned:

$$
\begin{aligned}
& prop(joe, likes, R) \leftarrow \\
& \qquad prop(R, type, resort) \wedge \\
& \qquad prop(R, near, B) \wedge \\
& \qquad prop(B, type, beach) \wedge \\
& \qquad prop(B, covered\_in, S) \wedge \\
& \qquad prop(S, type, sand).
\end{aligned}
$$

Joe likes resorts that are near sandy beaches.

# Predicting students errors

$$\begin{array}{rccc} & & x_2 & x_1 \\ + & & y_2 & y_1 \\ \hline & z_3 & z_2 & z_1 \end{array}$$

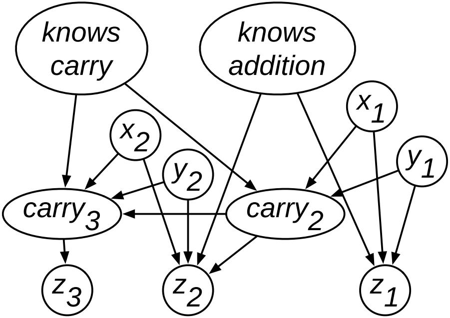

# Predicting students errors

| | | | |
|---|---|---|---|
| | | $x_2$ | $x_1$ |
| $+$ | | $y_2$ | $y_1$ |
| | $z_3$ | $z_2$ | $z_1$ |

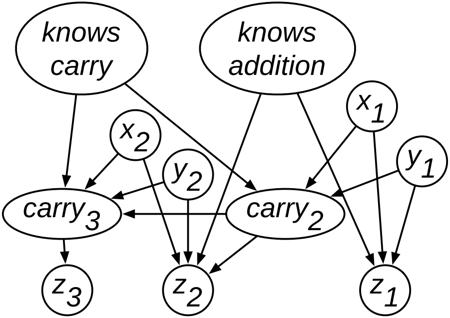


What if there were multiple digits

# Predicting students errors

$$\begin{array}{rrrr} & & x_2 & x_1 \\ + & & y_2 & y_1 \\ \hline & z_3 & z_2 & z_1 \end{array}$$

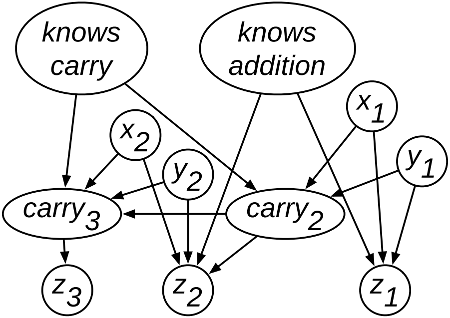


What if there were multiple digits, problems

# Predicting students errors

|  |  | $x_2$ | $x_1$ |
|---|---|---|---|
| + |  | $y_2$ | $y_1$ |
|  | $z_3$ | $z_2$ | $z_1$ |

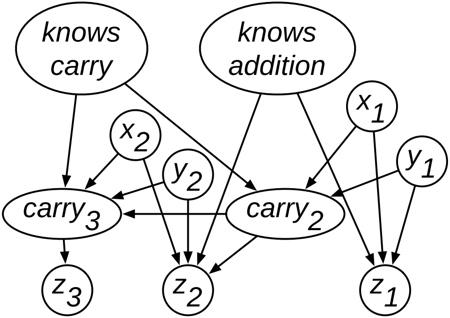


What if there were multiple digits, problems, students

# Predicting students errors

$$\begin{array}{cccc} & & x_2 & x_1 \\ + & & y_2 & y_1 \\ \hline & z_3 & z_2 & z_1 \end{array}$$

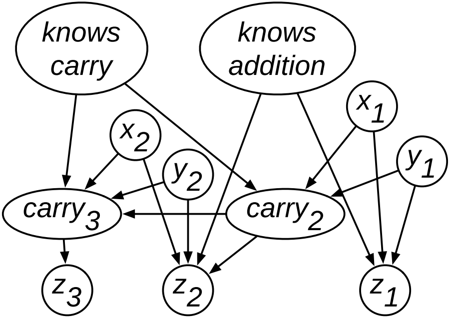


What if there were multiple digits, problems, students, times?

# Predicting students errors

$$\begin{array}{rccc} & & x_2 & x_1 \\ + & & y_2 & y_1 \\ \hline & z_3 & z_2 & z_1 \end{array}$$

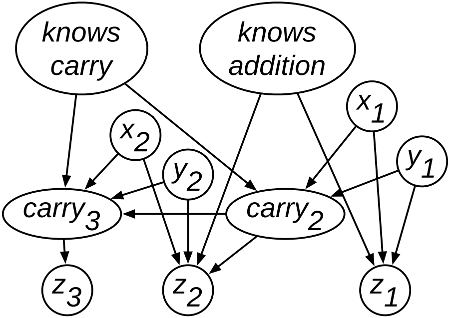


What if there were multiple digits, problems, students, times?
How can we build a model before we know the individuals?

# Multi-digit addition with parametrized BNs / plates

$$\begin{array}{ccccc} & x_{j_x} & \cdots & x_2 & x_1 \\ + & y_{j_z} & \cdots & y_2 & y_1 \\ \hline & z_{j_z} & \cdots & z_2 & z_1 \end{array}$$

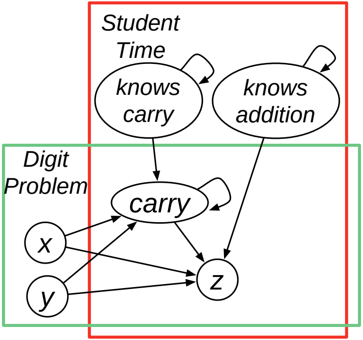


Random Variables: $x(D, P)$, $y(D, P)$, $knowsCarry(S, T)$, $knowsAddition(S, T)$, $carry(D, P, S, T)$, $z(D, P, S, T)$
for each: digit $D$, problem $P$, student $S$, time $T$

# Creating Dependencies: Relational Structure

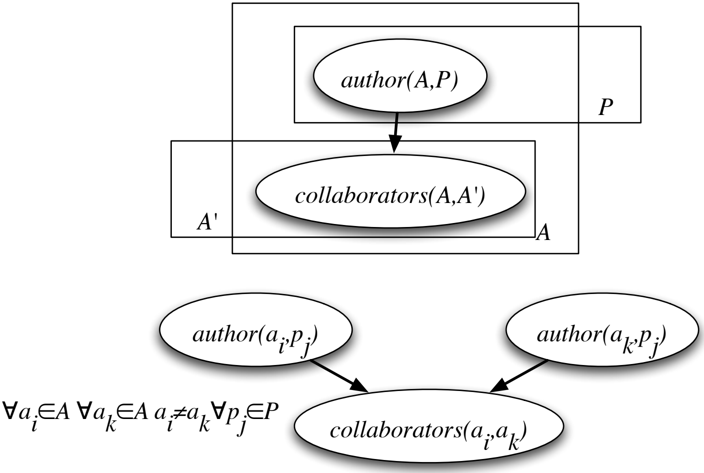

# Independent Choice Logic

- A language for first-order probabilistic models.
- Idea: combine logic and probability, where all uncertainty in handled in terms of Bayesian decision theory, and logic specifies consequences of choices.
- History: parametrized Bayesian networks, abduction and default reasoning $\longrightarrow$ probabilistic Horn abduction (IJCAI-91); richer language (negation as failure + choices by other agents $\longrightarrow$ independent choice logic (AIJ 1997).

# Independent Choice Logic

- An alternative is a set of atomic formula.
  $\mathcal{C}$, the choice space is a set of disjoint alternatives.
- $\mathcal{F}$, the facts is a logic program that gives consequences of choices.
- $P_0$ a probability distribution over alternatives:

$$\forall A \in \mathcal{C} \sum_{a \in A} P_0(a) = 1.$$

# Meaningless Example

$$\mathcal{C} = \{\{c_1, c_2, c_3\}, \{b_1, b_2\}\}$$

$$\begin{aligned}\mathcal{F} = \{\ & f \leftarrow c_1 \wedge b_1, && f \leftarrow c_3 \wedge b_2, \\ & d \leftarrow c_1, && d \leftarrow \neg c_2 \wedge b_1, \\ & e \leftarrow f, && e \leftarrow \neg d\}\end{aligned}$$

$$\begin{array}{lll} P_0(c_1) = 0.5 & P_0(c_2) = 0.3 & P_0(c_3) = 0.2 \\ P_0(b_1) = 0.9 & P_0(b_2) = 0.1 & \end{array}$$

# Semantics of ICL

- Possible world for each selection of one element from each alternative
- What is true in a possible world is given by the logic program
- Alternatives are assumed to be unconditionally independent

# Meaningless Example: Semantics

$$\mathcal{C} = \{\{c_1, c_2, c_3\}, \{b_1, b_2\}\}$$

$$\begin{aligned}\mathcal{F} = \{ \; & f \leftarrow c_1 \wedge b_1, & f \leftarrow c_3 \wedge b_2, \quad & d \leftarrow c_1, \\ & d \leftarrow \neg c_2 \wedge b_1, & e \leftarrow f, \quad & e \leftarrow \neg d\}\end{aligned}$$

$$\begin{array}{lll} P_0(c_1) = 0.5 & P_0(c_2) = 0.3 & P_0(c_3) = 0.2 \\ P_0(b_1) = 0.9 & P_0(b_2) = 0.1 & \end{array}$$

$$\begin{array}{llllllll} w_1 & \models & c_1 & b_1 & f & d & e & P(w_1) = 0.45 \\ w_2 & \models & c_2 & b_1 & \neg f & \neg d & e & P(w_2) = 0.27 \\ w_3 & \models & c_3 & b_1 & \neg f & d & \neg e & P(w_3) = 0.18 \\ w_4 & \models & c_1 & b_2 & \neg f & d & \neg e & P(w_4) = 0.05 \\ w_5 & \models & c_2 & b_2 & \neg f & \neg d & e & P(w_5) = 0.03 \\ w_6 & \models & c_3 & b_2 & f & \neg d & e & P(w_6) = 0.02 \end{array}$$

$$P(e) = 0.45 + 0.27 + 0.03 + 0.02 = 0.77$$

# Belief Networks, Decision trees and ICL rules

- There is a local mapping from belief networks into ICL.
- Rules can represent decision tree representation of conditional probabilities:

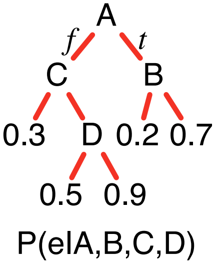


$$
\begin{array}{ll}
e \leftarrow a \wedge b \wedge h_1. & P_0(h_1) = 0.7 \\
e \leftarrow a \wedge \neg b \wedge h_2. & P_0(h_2) = 0.2 \\
e \leftarrow \neg a \wedge c \wedge d \wedge h_3. & P_0(h_3) = 0.9 \\
e \leftarrow \neg a \wedge c \wedge \neg d \wedge h_4. & P_0(h_4) = 0.5 \\
e \leftarrow \neg a \wedge \neg c \wedge h_5. & P_0(h_5) = 0.3
\end{array}
$$

# Example: Multi-digit addition

$$
\begin{array}{rcccc}
 & x_{j_x} & \cdots & x_2 & x_1 \\
+ & y_{j_z} & \cdots & y_2 & y_1 \\
\hline
 & z_{j_z} & \cdots & z_2 & z_1
\end{array}
$$

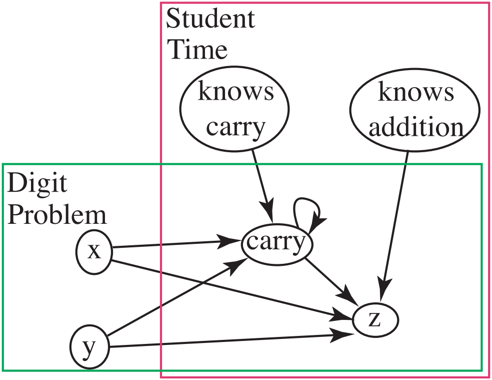

# ICL rules for multi-digit addition

$$
\begin{aligned}
z(D,P,S,T) = V \leftarrow\; & \\
& x(D,P) = Vx \wedge \\
& y(D,P) = Vy \wedge \\
& carry(D,P,S,T) = Vc \wedge \\
& knowsAddition(S,T) \wedge \\
& \neg mistake(D,P,S,T) \wedge \\
& V \text{ is } (Vx + Vy + Vc) \text{ div } 10.
\end{aligned}
$$

$$
\begin{aligned}
z(D,P,S,T) = V \leftarrow\; & \\
& knowsAddition(S,T) \wedge \\
& mistake(D,P,S,T) \wedge \\
& selectDig(D,P,S,T) = V. \\
z(D,P,S,T) = V \leftarrow\; & \\
& \neg knowsAddition(S,T) \wedge \\
& selectDig(D,P,S,T) = V.
\end{aligned}
$$

Alternatives:

$\forall DPST \{noMistake(D,P,S,T), mistake(D,P,S,T)\}$

$\forall DPST \{selectDig(D,P,S,T) = V \mid V \in \{0..9\}\}$

# Outline

1. Semantic Science Overview
   - Ontologies
   - Data
   - Hypotheses and Theories
   - Models

2. Levels of Semantic Science
   - Feature-based Theories
   - Relational Domains
   - Probabilities with Ontologies
   - Existence and Identity Uncertainty

# Random Variables and Triples

- Reconcile:
  - random variables of probability theory
  - individuals, classes, properties of modern ontologies

# Random Variables and Triples

- Reconcile:
  - random variables of probability theory
  - individuals, classes, properties of modern ontologies
- For functional properties:
  random variable for each $\langle individual, property \rangle$ pair, where the domain of the random variable is the range of the property.
- For non-functional properties:
  Boolean random variable for each $\langle individual, property, value \rangle$ triple.

# Triples and Probabilities

- $\langle individual, property, value \rangle$ triples are complete for representing relations
- $\langle individual, property, value, probability \rangle$ quadruples can represent probabilities of relations (or reify again)
- e.g., in addition $P(z(3, prob23, fred, t3) = 4) = 0.43$:

$$\left.\begin{array}{l}\langle z543, type, AdditionZValue \rangle \\ \langle z543, digit, 3 \rangle \\ \langle z543, problem, prob23 \rangle \\ \langle z543, student, fred \rangle \\ \langle z543, time, t3 \rangle\end{array}\right\} \text{defines random variable}$$

$$\left.\begin{array}{l}\langle z543, valueWithProb, 4, 0.43 \rangle \\ \langle z543, valueWithProb, 5, 0.03 \rangle \\ \ldots\end{array}\right\} \text{defines distribution}$$

# Probabilities and Aristotelian Definitions

Aristotelian definition

$$\begin{aligned} ApartmentBuilding \equiv\ & ResidentialBuilding\,\& \\ & NumUnits = many\,\& \\ & Ownership = rental \end{aligned}$$

leads to probability over property values

$$\begin{aligned} & P(\langle A, type, ApartmentBuilding\rangle) \\ & = P(\langle A, type, ResidentialBuilding\rangle) \times \\ & \quad P(\langle A, NumUnits, many\rangle \mid \langle A, type, ResidentialBuilding\rangle) \times \\ & \quad P(\langle A, Ownership, rental\rangle \mid \langle A, NumUnits, many\rangle\,, \\ & \qquad \langle A, type, ResidentialBuilding\rangle) \end{aligned}$$

No need to consider undefined propositions.

# Outline

# Existence and Identity

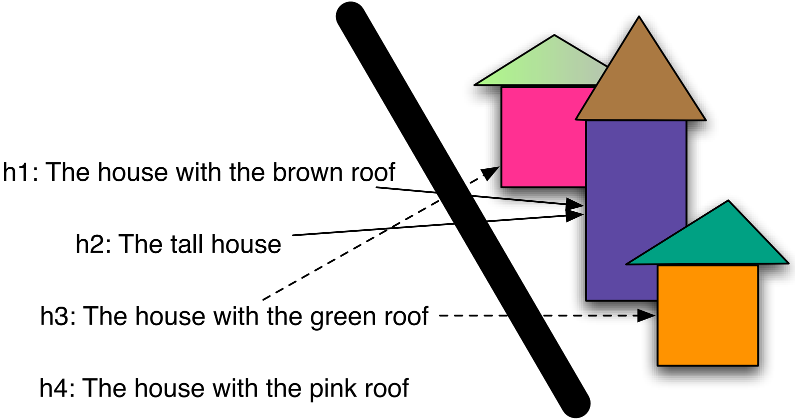

# Clarity Principle

Clarity principle: probabilities must be over well-defined propositions.

- What if an individual doesn't exist?
  - $house(h4) \land roof\_colour(h4, pink) \land \neg exists(h4)$

# Clarity Principle

Clarity principle: probabilities must be over well-defined propositions.

- What if an individual doesn't exist?
  - $house(h4) \wedge roof\_colour(h4, pink) \wedge \neg exists(h4)$
- What if more than one individual exists? Which one are we referring to?
  —In a house with three bedrooms, which is the second bedroom?

# Clarity Principle

Clarity principle: probabilities must be over well-defined propositions.

- What if an individual doesn't exist?
  - $house(h4) \wedge roof\_colour(h4, pink) \wedge \neg exists(h4)$
- What if more than one individual exists? Which one are we referring to?
  —In a house with three bedrooms, which is the second bedroom?
- Reified individuals are special:
  — Non-existence means the relation is false.
  — Well defined what doesn't exist when existence is false.
  — Reified individuals with the same description are the same individual.

# Correspondence Problem

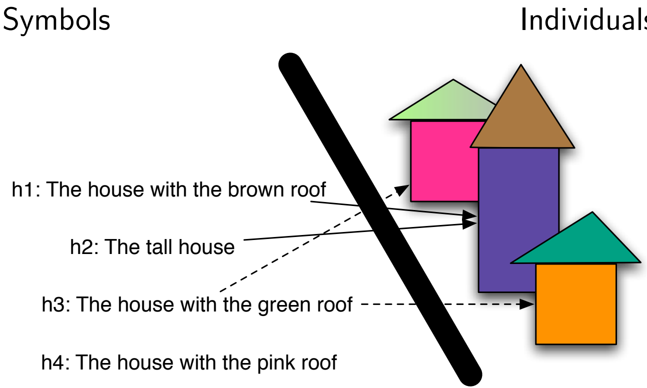


$c$ symbols and $i$ individuals $\longrightarrow c^{i+1}$ correspondences

# Role assignments

Theory about what apartment Mary would like.

Whether Mary likes an apartment depends on:

- Whether there is a bedroom for daughter Sam
- Whether Sam's room is green
- Whether there is a bedroom for Mary
- Whether Mary's room is large
- Whether they share

# Role assignments

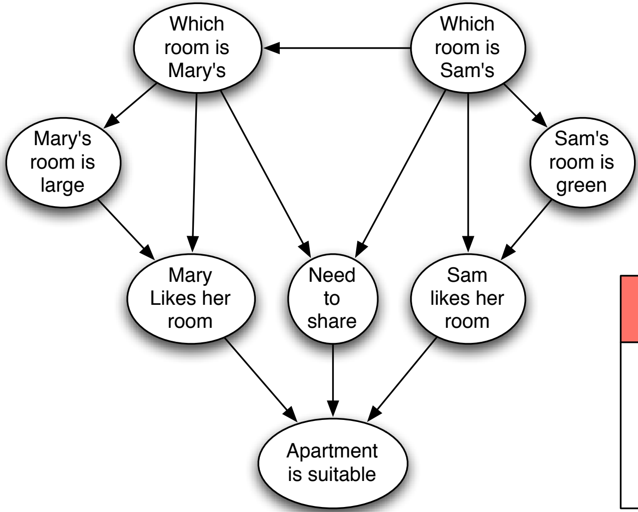


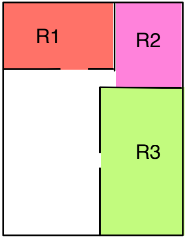

# Conclusion

- Semantic science is a way to develop and deploy knowledge about how the world works.
- Scientists (and others) develop theories that refer to standardized ontologies and predict for new cases.
- Multiple theories—forming models—are needed to make predictions in particular cases.
- For each prediction, we want to be able to ask, what theories it is based on.
- For each theory, we want to be able to ask what evidence it is based on.
- This talk is deliberately pre-theoretic. Many formalisms will be developed and discarded before we converge on useful representations.

# To Do

- Theories of combining theories.
- Representing, reasoning and learning complex (probabilistic) theories.
- Build infrastructure to allow publishing and interaction of ontologies, data, theories, models, evaluation criteria, meta-data.
- Build inverse semantic science web:
  - Given a theory, find relevant data
  - Given data, find models that make predictions on the data
  - Given a new case, find relevant models with explanations
- More complex models, e.g., for relational reinforcement learning where individuals are created and destroyed