

Towards a Logic of Feature-Based Semantic Science Theories

David Poole

Dept of Computer Science,
University of British Columbia
Vancouver, BC, Canada
<http://cs.ubc.ca/~poole/>

Abstract

The aim of semantic science is to allow for the publications of ontologies, observation data, and hypotheses/theories. Hypotheses make predictions on data and on new cases. Those hypotheses that fit the available evidence are called theories. This paper considers how theories can be used for predictions in new cases. Theories are typically very narrow and not all of the inputs to a theory are observed, so to make predictions on a particular case, many theories need to be used. Without any global design, the available theories do not necessarily fit together nicely. This paper explains how theories can be combined into theory ensembles to make predictions on a particular case. This is needed to evaluate theories, and to make useful predictions. We motivate and give desiderata for theory ensembles for level 1, feature-based, semantic science, which assumes that the data and the theories can be described in terms of features (random variables).

Introduction

If a KR system makes a prediction, it is reasonable for someone to ask: what evidence is there for that prediction? The system should be able to provide such evidence. If a knowledge-based system is to believe something, it should believe it based on evidence, as not all beliefs are equally valid. The mechanism that has been developed for judging knowledge is called *science*. We have used to term *semantic science*, in an analogous way to the *semantic web*, because the computer should understand the hypotheses and data which form the foundation of science itself. It is not meant to just apply to the traditional scientific disciplines; the scientific method is applicable to any domain.

The semantic web (Berners-Lee, Hendler, and Lassila 2001) is an endeavor to make all of the world's knowledge accessible to computers. One of the central planks of the semantic web is how to trust the information given. Trust in the truth of some information, or what (Gil and Artz 2007) call *content trust*, has been cast in terms of social trust relationships and search engines such as Google base their ranking on popularity, but often return authoritative sites. If you are a scientist, popularity and appeal to authority are not the basis for determining what is true. Science determines truth based on empirical evidence: what does all of the available

evidence lead us to conclude? It is this notion of trust that semantic science deals with; we trust scientific conclusions because they are based on evidence. The basis of semantic science (Poole, Smyth, and Sharma 2008) is:

- Information is published using well defined ontologies (Smith 2003) to allow semantic interoperability.
- Data is published (Fox et al. 2006; McGuinness et al. 2007) about observations of the world described using the vocabulary specified by the ontologies.
- Scientists (and others) publish hypotheses that make predictions on data. These hypotheses make reference to ontologies. Hypotheses that fit the data are called theories.
- New data can be used to evaluate the hypotheses that make predictions on that data.
- Descriptions of competing theories can be used to devise experiments that will distinguish the theories.
- To make a prediction for a new case (e.g., predict the effect of treatment of a patient in a diagnostic setting, or predict where a landslide may occur), many theories may need to be used together to make a prediction.
- There is no central authority to vet which theories are legitimate. Each of us can choose to make decisions based on whichever theories we want.
- We expect semantic science search engines to be developed. Given a hypothesis, a search engine would be able to find data that can be used to evaluate or tune the hypothesis. Given data or a new case, a search engine would be able to find the best theories that make predictions on the data or new case. The aim is that many possible hypotheses can be published and the search engine will return theories, those hypotheses that best fit the available evidence.

The relationship amongst ontologies, data and theories is given in Figure 1. The data depends on the world and the ontology. The theories depend on the ontology, indirectly on the world (via a human who is designing the theory), and directly on some of the data (as we would expect that the best theories would be based on as much data as possible). All theories for the foreseeable future will be generated by humans designing the model structures and using machine learning to fit the models to data. Given a new case, theories




Figure 1: Role of Ontologies, Data and Theories in Semantic Science

make predictions about the case that can be used for decision making. The ontologies, data sets and theories, evolve in time.

To make this project manageable, we can define four levels of semantic science:

0. **Deterministic semantic science** where all of the theories make definitive predictions. This class includes both propositional and first-order theories. This has been studied under the umbrella of abduction (Kakas and Denecker 2002, see e.g.,).
1. **Feature-based semantic science**, where there are non-deterministic¹ predictions are about feature values of data. Predicting values for features is the most common form of machine learning. Such theories can be specified in terms of random variables that represent the values of features.
2. **Relational semantic science**, where predictions are about the properties of objects and relationships among objects. The values of properties may be meaningless names; the structure of the relationships is used to make predictions. The mix of probabilistic and relational reasoning has been studied under the umbrella of statistical relational learning (Getoor and Taskar 2007; De Raedt et al. 2008).
3. **First-order semantic science**, where the aim is to make predictions about the existence of objects, predictions about universally quantified statements or predictions of equality. This is more challenging as conditioning is not well-defined (Poole 2007), as a theory may refer to the existence of an object filling a role, but observational data does not specify which object fills that role.

This paper will concentrate on level 1, feature-based semantic science. An understanding of this is necessary before solving the more complex levels.

Ontologies

An ontology is a specification of the meaning of the symbols in an information system. Ontologies define the vocabulary used in the data and the theories. Any terminology that

¹Non-deterministic can mean many things. Here we consider just the case where there are probabilistic predictions (Jaynes 2003). But there are many alternatives, such as qualitative predictions, probability ranges or fuzzy predictions. Different users may be happy with different predictions.

needs to be shared among data sets and/or theories needs to be defined in an ontology.

We assume that ontologies are written in terms of Aristotelian definitions (Berg 1982; Poole, Smyth, and Sharma 2009), where each class is defined in terms of a super-class, the *genus*, and attributes that distinguish this class from other subclasses of the genus, the *differentia*. Thus we assume that all subclasses are defined in terms of restrictions on property values.

In feature-based representations, we can equate features with properties using property-based ontology languages such as OWL (Hitzler et al. 2009), using some generic individual. Thus we assume that features have domains, which must be true for them to be defined.

Data

Observational data is published referring to the ontologies. As part of each data set, assume the following is specified:

- The **context** in which the data was collected. This is a proposition made up of assignments to features.
- The **control features** that were controlled for in the data (sometimes called the independent variables).
- The **observed features** that this data makes predictions about (sometimes called the dependent variables).
- A database on all of the variables, where each tuple specifies a value for each of the control and observed features.

Theories

Each theory makes predictions about some feature values.

We assume a theory has the following components:

- A **context**, c , which is a proposition that specifies pre-conditions of when the theory can be applied. This is a proposition that implies the domains of the features used in the theory; it must be true for the theory to make sense.
- A set of **input features**, I , about which it does not make predictions.
- A set of **output features**, O , about which it can make a prediction (as a function of the input features). $I \cap O = \{\}$.
- A program P that makes a prediction of values of O given context c and values to the features I . This program can compute $P(O|c, I)$ for given values of I .

One type of theory that is of particular interest is the “null hypothesis”. There is a (maximum likelihood) null hypothesis for each feature which specifies that the feature has a distribution that is independent of the other features.

The context c must entail the domains of the features in I and O , so that each of these is well defined when c is true.

Note that, if a theory makes a prediction on features O , in principle, it can also be used to make predictions on subsets of O . However, it is not always computationally feasible to sum out the variables needed to compute this.

Theory Ensembles

To make a prediction, we need more than the single best theory. We need to use multiple theories that fit together to make a prediction. We call such a collection of theories a **theory ensemble**.

A **query** consists of an observation *obs* and a set of query variables *Q*. The aim is to estimate $P(Q|obs)$.

Given a query, there is typically no theory that can be directly used to make the prediction. Rather, multiple theories are needed to make a prediction. The collection of theories needed to predict *Q* given *obs*, is called a theory ensemble. Some properties of a theory ensemble *T* are:

- Theories do not have to be used in their full generality in a theory ensemble; they can be used in restricted contexts or for only a subset of their output features.
- *T* is coherent: it does not rely on the value of a feature in a context where the feature is not defined (i.e., outside of the domain of the feature). Thus if feature *f* has domain *d*, the feature has to be used in a context where *d* is true.
- *T* is consistent: it does not make different predictions for any feature in any context.
- *T* is predictive: it makes a prediction of *Q* every context that is possible given *obs*.
- *T* is minimal in that it does not include theories that are not required to be predictive.

For level-0 semantic science, this corresponds to the standard definition of abduction (Kakas and Denecker 2002). The predictive condition corresponds to being able to prove the goal. Coherence is needed with ontologies.

For level 1 semantic science, the situation is more complex, as a theory ensemble does not give definitive predictions for the context of the theories of a theory ensemble. For example, if a theory ensemble contains a theory that makes a prediction on *B* when *a* is true, the theory ensemble needs to predict *a* from *obs* and, if this prediction is not definitive (with probability 1), also predict *B* in the context of $\neg a$.

One way to build a feature-based theory ensemble (ignoring coherence) is to construct a Bayesian network from the theories. We can do better than this by allowing different theories in different contexts.

Conclusion

This paper presents some desiderata of a theory ensemble that can be used to make a prediction for a new case, given a set of theories. By giving a syntactic specification of a theory ensemble, we can then search over the theory ensembles looking for the theory ensemble that is best supported by the evidence. It is also possible to be Bayesian and to predict from all theory ensembles based on their posterior probability. Searching efficiently for the appropriate theory ensembles given a large distributed collection of theories is still an open problem.

Defining theory ensembles is just a first step in developing a logic of feature-based semantic science theories. The next step is a mechanism to evaluate theories and theory ensembles based on observational data. We can use such an evaluation to judge predictions.

This paper has ignored interventions, which are central to science. The data and the theories need to distinguish observations from controls. The predictions from observing and controlling a variable can be very different (Pearl 2000).

Once we have good understanding of feature-based semantic science, the next step is relational semantic science, where the random variables are a function of the individuals in the domain. Even more complex is the first-order case, with theories about the existence or identity of individuals.

References

- Berg, J. 1982. Aristotle's theory of definition. In *ATTI del Convegno Internazionale di Storia della Logica*, 19–30.
- Berners-Lee, T.; Hendler, J.; and Lassila, O. 2001. The semantic web: A new form of web content that is meaningful to computers will unleash a revolution of new possibilities. *Scientific American* May:28–37.
- De Raedt, L.; Frasconi, P.; Kersting, K.; and Muggleton, S. H., eds. 2008. *Probabilistic Inductive Logic Programming*. Springer.
- Fox, P.; McGuinness, D.; Middleton, D.; Cinquini, L.; Darnell, J.; Garcia, J.; West, P.; Benedict, J.; and Solomon, S. 2006. Semantically-enabled large-scale science data repositories. In *5th International Semantic Web Conference (ISWC06)*, LNAI 4273, 792–805. Springer-Verlag.
- Getoor, L., and Taskar, B., eds. 2007. *Introduction to Statistical Relational Learning*. Cambridge, MA: MIT Press.
- Gil, Y., and Artz, D. 2007. Towards content trust of web resources. *Journal of Web Semantics* 5(4):227–239.
- Hitzler, P.; Krötzsch, M.; Parsia, B.; Patel-Schneider, P. F.; and Rudolph, S. 2009. *OWL 2 Web Ontology Language Primer*. W3C. <http://www.w3.org/TR/owl2-primer/>
- Jaynes, E. T. 2003. *Probability Theory: The Logic of Science*. Cambridge University Press.
- Kakas, A., and Denecker, M. 2002. Abduction in logic programming. In Kakas, A., and Sadri, F., eds., *Computational Logic: Logic Programming and Beyond*, number 2407 in LNAI. Springer Verlag. 402–436.
- McGuinness, D.; Fox, P.; Cinquini, L.; West, P.; Garcia, J.; Benedict, J. L.; and Middleton, D. 2007. The virtual solar-terrestrial observatory: A deployed semantic web application case study for scientific research. In *IAAI-07*.
- Pearl, J. 2000. *Causality: Models, Reasoning and Inference*. Cambridge University Press.
- Poole, D.; Smyth, C.; and Sharma, R. 2008. Semantic science: Ontologies, data and probabilistic theories. In da Costa et al. eds., *Uncertainty Reasoning for the Semantic Web I*, LNAI/LNCS. Springer.
- Poole, D.; Smyth, C.; and Sharma, R. 2009. Ontology design for scientific theories that make probabilistic predictions. *IEEE Intelligent Systems* 24(1):27–36.
- Poole, D. 2007. Logical generative models for probabilistic reasoning about existence, roles and identity. In *22nd AAAI Conference on AI (AAAI-07)*.
- Smith, B. 2003. Ontology. In Floridi, L., ed., *Blackwell Guide to the Philosophy of Computing and Information*. Oxford: Blackwell. 155–166.