Towards an abductive foundation of semantic science

David Poole
Department of Computer Science,
University of British Columbia,
2366 Main Mall, Vancouver, B.C., Canada V6T 1Z4
poole@cs.ubc.ca
http://cs.ubc.ca/~poole/

Abstract

The aim of semantic science is to have scientific data and scientific theories represented and published in machine understandable form. There is much work on developing scientific ontologies and representing scientific data in terms of these ontologies. The next step is to publish scientific theories that can make predictions on the published data and can be used for prediction on new cases. This can be used to advance the development of science and to provide useful predictions that can be evaluated according to all available data.

To make a prediction for a particular case, we need to use an ensemble of theories that fit together (are consistent) and make a prediction on a particular case. We argue that this is a form of abduction, that has similarities and differences to the standard definitions of abduction. This is preliminary work, presenting pre-theoretic foundations of the field.

Introduction

The basic idea of semantic science [Poole et al., 2008] is:

- Information is published using well defined ontologies [Smith, 2003b] to allow semantic interoperability.
- People publish data [Fox et al., 2006; McGuinness et al., 2007] described using the vocabulary specified by the ontologies. Part of this data includes metadata about what the data is about and how it was generated. Data repositories include the Community Data Portal (http://cdp.ucar.edu/) and the Virtual Solar-Terrestrial Observatory (http://vsto.hao.ucar.edu/index.php).
- Scientists publish theories that make predictions on data. These theories make reference to ontologies. These predictions can be tested on the published data. As part of each theory is information about what data this theory is prepared to make predictions about.
- New data can be used to evaluate, and perhaps update, the theories that make predictions on this data. Predictions on new data can be used to judge the theories as well as find outliers in the data, which can be statistical anomalies, fraudulent data or some new, little understood phenomenon.
- The descriptions of competing theories can be used to devise experiments that will distinguish the theories.
- If someone wants to make a prediction for a new case (e.g., a patient in a diagnostic setting, or predicting a landslide), they can use the best theories to make the prediction. They would either use the best theory or theories, or average over all theories weighted by they ability to predict this phenomenon of interest. The use will be able to ask for what evidence there is for the theory.
- There is no central authority to vet as to what counts as legitimate scientific theories. Each of us can choose to make decisions based on the whichever theories we want. We will be able to judge theories by their predictions on unseen data and other criteria.
- We expect semantic science search engines to be developed. Given a theory, a search engine would be able to find data that can be used to evaluate or tune the theory. Given data, a search engine would be able to find the theories that make predictions on the data.

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 (if a human 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). Given a new case, theories make predictions about that case that can be used for decision making. The ontologies, data sets and theories, evolve in time.

The term "science" is meant to be as broad as possible. We can have scientific theories about any natural or artificial phenomenon. We could have scientific theories about traditional disciplines such as earth sciences, physics, chemistry, biology, medicine and psychology but we would also imagine theories as diverse as predicting which companies will be most profitable, predicting where the best parties are, or predicting who will win football games. The only criteria is that a scientific theory must put itself at risk by making predictions about observable phenomenon.

Semantic science has no prior prejudice about the source or the inspiration of theories; as long as the theories are prepared

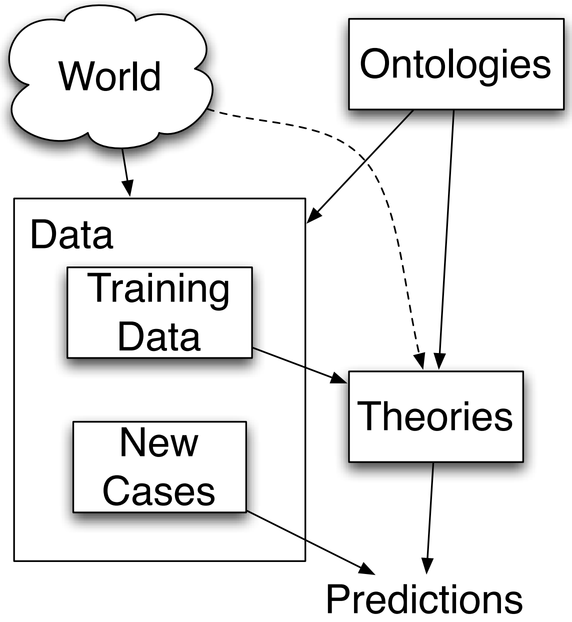

Figure 1: Role of Ontologies, Data and Theories in Semantic Science

to make predictions about unseen data, they can be included. We are not, a priori, excluding religion, astrology, or other areas that make claim to the truth; if they are prepared to make predictions about what will be observed, we can test how well their predictions fit the available data, and use their predictions for other data.

Semantic science is trying to be broad and bottom-up. It should serve to democratize science in allowing not just the elite to create data and theories. Like scientists themselves, it should be skeptical of all of the information it is presented with.

We anticipate that the most useful theories will make probabilistic predictions, however theories can make whatever predictions they like. Users of the theories can choose to adopt theories based on whatever criteria they like, e.g., some combination of fit to the existing data and simplicity or prior plausibility. Users can also choose to ignore theories that don't make the sort of predictions they like.

This is like the machine learning vision, where the data sets are heterogeneous and published with respect to formal ontologies. The theories also persist and can be compared when new data arrives. We expect the highest standards to be used in evaluation of the theories. For the foreseeable future virtually all theories will be a mix of human generated and machine learned; humans define the structure and parameter space and the machines optimizes these with respect to fit to data and learning biases. Semantic science provides a mechanism for Bayesian inference, where we need to condition on *all* relevant information that was not part of building the model. A semantic science search engine should allow us to find all of the relevant data on which to condition.

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 abductive logic programming [Kakas and Denecker, 2002].
1. Feature-based semantic science, where there are non-deterministic[1] predictions are about feature values of data. This 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 the predictions are about the properties of objects and relationships among objects. In this case, the values of properties may be meaningless names; the structure of the relationships is used to make predictions. This is what has been studied in 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 or predictions about universally quantified statements. This is more challenging as conditioning is not well-defined [Poole, 2007]. We may not know which object in the world the theory is making a prediction about, as the theory may refer to the existence of object filling a role, but we may know which object fills the role.

In the rest of this paper, we only consider the first two of these. We will describe these in terms of features. Features generalize propositions, as a proposition is a Boolean feature. Features can also be seen as properties of a single individual under consideration. There can also be global features that are not about any individual. Seeing the features as properties allow for a correspondence with the work on ontologies. An attribute is a feature-value (or property-value) pair, for example that rock's age is 50 million years is an attribute of the rock.

Predictions from Theories

Scientific theories are typically narrow; they don't make predictions on arbitrary sets of data. For example, someone may develop a theory for the prognosis of a particular type of lung cancer. To use this theory for a prediction of a particular patient, we first predict whether the patient has this form of lung cancer, then use this theory to predict the prognosis. We need other theories about the prognosis for the possibility that the patient has a different form of lung cancer, or doesn't have lung cancer.

A set of theories that fit together to make a prediction for a particular case is called a *theory ensemble*. The theories in a theory ensemble must be consistent, in a way we describe below, and must be enough to make a prediction in a particular case.

There is a correspondence between theory ensembles and explanations in abduction. There will be a direct match for

[1]Non-deterministic can mean many things. Here we consider just the case where there are probabilistic predictions. But there are many alternatives, such as qualitative predictions, probability ranges or fuzzy predictions.

level 0 semantic science. The other levels will be more complicated when the theories make probabilistic predictions, and when we can't even be sure that a theory makes semantic sense for a particular data set.

We don't assume that theories are explicitly specified, in the sense that theories can can be arbitrarily complex and use arbitrary computation to make predictions.

Ontologies

In AI, an ontology [Smith, 2003b; Noy and Hafner, 1997; Gómez-Pérez et al., 2004] is a specification of the meaning of vocabulary used by an information system. Ontologies form the backbone of the Semantic Web [Berners-Lee et al., 2001]. There has recently been much work in standardizing ontologies, such as using the Web Ontology Language OWL [McGuinness and van Harmelen, 2004]. Science is one of the areas where ontology development and deployment is well under way [Smith et al., 2007].

Ontologies can be very complicated, as would be expected in a world where language has evolved to be useful and new terminology is invented to describe what was not easy to describe using previous terminology.

We have been advocating a structure for ontologies using what we call Aristotelian definitions [Smith, 2003a; Poole et al., 2009], based on the idea of Aristotle [350 B.C.] that each class should be described in terms of a super-class (the genus) and property values (the differentia) that differentiate this class from other subclasses of the genus. Defining all classes in terms of properties, as opposed to specifying subclass relationships directly, simplifies reasoning as we only need to give the values of properties and the class structure logically follows. It is also a natural way to define concepts in many cases. Simple Aristotelian definitions often give rise to complicated subclass relationships, but simple subclass relationships give simple Aristotelian definitions.

For the rest of this paper, we will thus ignore classes, and consider only features (conflating features and properties as we are only considering feature-based semantic science). Properties, however, have domains; they are only defined in the context where other properties have particular values. Properties are not defined in when their domain does not hold.

Data

We assume that data is published referring to the ontologies used. As part of each data set, for the purpose of this paper, assume the following meta-data is specified:

- The context in which the data was collected. This is a proposition made up of assignments to features. For example, if the data was of people who have a certain type of cancer, the context would be the attributes that define the people and the attributes that define the type of cancer.
- The features that this data makes predictions about (what is often called the dependent variables).
- The features that were controlled for in the data (the independent variables).

To predict such data, a theory needs to predict the values of the dependent variables as a function of the context and the independent variables.

Theories

Each theory makes predictions about some feature values or property values of an individual.

We assume a theory has three components:

- A context in which specifies preconditions of when it can be applied. This is a proposition that must be true for the theory to make sense.
- A set of input features about which it does not make predictions.
- A set of output features about which is can make a prediction (as a function of the input features).

For example, the ideal gas law is a theory that makes predictions about the pressure P, volume V, number of particles n and the temperature in the context of a gas, namely that $PV \propto nT$. It makes predictions that can be judged against data. There are alternative theories are be more accurate for real gasses, e.g., when the pressure is high, and the gas molecules are heterogeneous. This theory is not applicable to rocks or to lung cancer.

Theories are not universally applicable; for example we can't use a theory about the prognosis of people with cancer on rocks. Theories have preconditions that specify what they make predictions about. These preconditions are of three different sorts:

- Conditions which define when the theory makes sense. When these conditions are false, the theory is nonsense. The conditions are the domains of the features used in the theories.
- Conditions which define the intended scope of the theory. These conditions specify what the theory was designed to predict.
- Conditions which specify when the theory will be used in a particular theory ensemble.

For example, a theory that makes predictions of the prognosis of patients with lung cancer may be applicable for arbitrary people. In a particular theory ensemble, it may only be used for the patients with lung cancer who have not had some particular drug, as the theory ensemble may use another theory that makes predictions in that case.

One class of theories that is of particular interest is the "null hypothesis". There is a null hypothesis for each feature. This theory says that the feature has randomly distributed values, with probabilities that are independent of the other features. It is important as it is always applicable, and gives a base case upon which to compare other theories.

Theory Ensembles

To make a prediction, we need more than a single theory. We need to use multiple theories that fit together to make a prediction. We call such a collections of theories a *theory ensemble*. We expect a formal definition of theory ensembles to be quite complex to cover the richness of real theories. However, there does seem some to be properties that we can define independently of any formalism.

A theory ensemble T needs to satisfy the following properties:

- T is coherent: it does not rely on the value of a feature in a context where the features is not defined (i.e., outside of the domain of the feature). Thus if feature f has domain d, it has to be used in a context where d is true. For example, writing $d \wedge f$, which is false if d is false, and has the value of f otherwise, would satisfy coherence.
- T is consistent: it does not make different predictions for any feature in any context.
- T is predictive: it makes a prediction in every context that is possible. Thus if we have a theory that includes $a \rightarrow b$, and we need to make a prediction on b, then we need to have our theory ensemble imply a, or also predict b in the context of $\neg a$.
- 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. The predictive condition corresponds to being able to prove the goal. Coherence is also needed for theories that use ontologies, but if we make the domain of a property as a precondition for the property, coherence is entailed by the other three properties for the deterministic case.

For type 1 semantic science, the situation is more complex, and there is still much more research required to get a satisfactory definition of a theory ensemble. A simplistic notion of a theory ensemble for a particular piece of data (that contains a context, values for its independent variables and values for its dependent variables) consists of a set of $\langle c,t \rangle$ pairs where t is a theory and c is a proposition which implies the domains of the properties used in t. The pair $\langle c,t \rangle$ specifies that theory t will be used for predictions in the context c. The following example shows how this notion of a theory can be used with the properties defined above to give a prediction:

Example 1 Suppose we have data about a person who coughs, and we want to make predictions about their prognosis. We have the following Boolean random variables (we will use the lower case variant as the proposition that the variable is true):

- *Person* is true if the object is a person.
- L is true if the person will live for more than a year (it gives the prognosis of a person).
- *HC* is true if the person has cancer
- *HLC* is true if the cancer the person has is lung cancer
- *Coughs* is true if the person coughs

Suppose the background ontology specifies the *person* is the domain of the properties L, *HC* and *Coughs*. The domain of *HLC* is *hc* (i.e., we can only talk about the value of *HLC* when *HC* is true).

Suppose we have the following theories that have been published:

- T_1 is about the prognosis of people with lung cancer.
- T_2 is about the prognosis of people with cancer.
- T_3 is the null hypothesis that gives the prognosis of people in general.
- T_4 predicts (probabilistically) whether people with cancer have lung cancer, as a function of coughing (i.e., *hc* is the context, *Coughs* is the independent variable and *HLC* is the dependent variable).
- T_5 predicts (probabilistically) whether people have cancer

A possible theory ensemble is $\{\langle person, T_5\rangle, \langle \neg hc, T_3\rangle, \langle hc, T_4\rangle, \langle hlc, T_1\rangle, \langle hc \wedge \neg hlc, T2\rangle\}$

In this ensemble, although T_2 can make predictions for anyone with cancer, it is only used for those without lung cancer. Similarly T_3 is only used when the person does not have cancer.

If T_4 made definitive predictions about lung cancer, only one of $\langle hlc, T_1\rangle$ and $\langle hc \wedge \neg hlc, T2\rangle$ would be in the theory ensemble.

This example has ignored many of the details of real theories. Theories can make predictions about many features and an ensemble may not need to use all of them. We need to treat conditions differently depending of whether they are part of the context, whether they are observed in the data and when they are not observed in the data. We also need to be concerned with how the theories interact with the ontologies.

Frequently Asked Questions

There are a number of questions that have been asked. Some for which I have a reasonable answer are here.

Will this replace peer review?

There is a related question of "What is the role of humans in semantic science?" In some sense the goal of semantic science is to let the computers do what that are good at, and let humans do what they are good at. This is true for evaluation too; we should use computers to evaluate what computers are good at evaluating and let humans evaluate what they are good at evaluating. Computers are (should be) good at evaluating how well theories fit data. But fit to data isn't the only property we want of a scientific theory. We also want insight; computers are not as good at evaluating this. We also want a notion of simplicity and elegance of theories, which may be hard to formally specify for a computer. So semantic science will not replace peer review, but will give extra tools for which to evaluate science.

What is to prevent fraudulent data and theories?

It might seem that the enterprise will break down on fraudulent data, as people publishing fraudulent data can make their theories look good. Suppose someone was to post fraudulent data. First, existing theory ensembles that make predictions on that data will be surprised by the data; they will conclude that the data is very unlikely. Conditioned on the new data, the theory ensembles will become less likely. Next someone could propose a theory that the data is anomalous or fraudulent (or perhaps such a hypothesis could always automatically available). This could have a small prior probability that may depend on the source of the data. The theory ensembles will then split into those that adopt the theory that the data is fraudulent, and those that do not. These new theory ensembles can be evaluated. A theory ensemble that concludes that all of the data is anomalous will not be very likely. A theory ensemble that can account for anomalous data will be become more likely. Thus the general mechanism of evaluating theory ensembles can handle anomalous and fraudulent data.

As theories can make any predictions, it seems as though there could not be fraudulent theories. However, there are two cases that need to be taken into account.

We may expect that theories would declare what data they used to learn from. However a theory could lie. This may be problematic if the theory is trained on the data used to evaluate theories. Such a theory would look good. However, such theories will not continue to look good when tested on brand new data.

The other way a theory can be fraudulent is to use other theories without acknowledgment. If not all theories are open, it is possible to steal parts of other theories. The use of theory ensembles is meant to mitigate any advantage that could be obtained by doing this. Reusing other theories is a legitimate part of semantic science, so there seems to be no advantage of stealing other theories. Theory ensembles also give credit where credit is due.

How does it relate to ensemble learning?

Ensemble learning [Dietterich, 2002] is a common technique for combining multiple learners to get a better prediction. The ensembles are typically a combination of predictions of a target feature based on the input features. Combining predictions for a single feature is definitely allowed as part of a theory ensemble. There are however reasons why this may not become commonplace. First, we don't expect the theories to be developed independently; theories are developed building on previous theories, trying to improve them. Much of the work on ensemble learning is about how to sensibly generate the classifiers that will be combined, whereas here we have predictors that are not designed to be combined. Second, the producers of such theories will want to make the best theories possible. The tools that are available to them include ensemble learning. It doesn't seem that cascaded ensemble learning will work well.

Theory ensembles can be an arbitrary combination of theories. This will range from the linear averaging of bagging to the conditional application of the example given above. A specification language for theory ensembles will allow all such combinations.

Theory ensembles are also related to algorithm portfolios [Xu et al., 2008], which learn which algorithms to run based on features of the problem being solved. We expect that similar learning could be used to choose which theories make the best predictions under which circumstances. Just as the algorithm portfolio learning can use algorithms that were not designed to be used with the portfolio, we expect that scientists will develop theories without needing to be concerned about how they will be used.

How can data, theories, ensembles and ontologies evolve in time?

We expect data to be continually published. Ontologies evolve to accommodate new categories of observations. Theories also improve. Theories can use whatever internal computations they like to make predictions. However, if one theory wants to use the some internal feature of an another theory, or if that feature is added to the data, the vocabulary to describe that feature needs to be added to the ontology. Ontologies can be evaluated by whether the distinctions they describe are useful in making predictions. Thus as theories become more sophisticated, the distinctions they need to make their predictions are added to the vocabulary, and are incorporated into the data.

How do we get there?

There is currently much work on developing scientific ontologies and publishing data with respect to these ontologies [Smith et al., 2007; Fox et al., 2006; McGuinness et al., 2007]. It would seem that publishing ontologies and data will only continue to grow. Scientists want others to use their data, as do their funders. There is growing recognition of the need to develop ontologies to allow for the sharing of such data and other information. Scientists care about the language used to describe their science. Many become involved in developing scientific vocabulary, and the formal representation in ontologies, because they don't want others to define the vocabulary that will become standard.

There has been much less work on developing theories. We have built some systems in geology, for minerals exploration and landslide susceptibility [Jackson, Jr. et al., 2008; Sharma et al., 2009] that represent published theories that make predictions about a limited number of properties. These systems were quite complicated as they reasoned about the probability of existence of individuals that filled roles.

Based on this experience, we recognized that there is still work to be done on feature-based representations before we try to extend it to relational and first-order representations. We need to build future systems on solid foundations.

By first developing level-1 semantic science based on features, we should be able to develop firm foundations for this case, in much the same way that machine learning has been able to develop in this context. We can then move to relations and then to first-order semantic science.

Conclusion

This paper has sketched some pre-theoretic ideas on how theory ensembles work and their relationship to explanations in abduction. I believe that it is important to get the pre-theoretic notions correct before creating a formalism that can be studied as an abstract entity. I also expect that there will be many iterations of getting the definitions of theories and theory ensembles right.

The potential of semantic science seems huge, but there are many technical and social issues that need to be solved before it can come to maturity. The development of ontologies and the publishing of data using those ontologies has advanced greatly in recent years. The main technical issues remaining are to do with the representations of the theories and theory ensembles and the infrastructure to publish and search for data and theories. To bring this vision of semantic science to fruition will require advances in many fields.

References

Aristotle (350 B.C.). *Categories*. Translated by E. M. Edghill, `http://www.classicallibrary.org/Aristotle/categories/`.

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*, pp. 28–37. URL `http://www.sciam.com/article.cfm?id=the-semantic-web`.

De Raedt, L., Frasconi, P., Kersting, K., and Muggleton, S.H. (Eds.) (2008). *Probabilistic Inductive Logic Programming*. Springer.

Dietterich, T.G. (2002). Ensemble learning. In M. Arbib (Ed.), *The Handbook of Brain Theory and Neural Networks,*, pp. 405–408. MIT Press, Cambridge, MA, second edition. URL `http://web.engr.oregonstate.edu/~tgd/publications/hbtnn-ensemble-learning.ps.gz`.

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)*, volume 4273 of *Lecture Notes in Computer Science*, pp. 792–805. Springer-Verlag. URL `http://www.ksl.stanford.edu/KSL_Abstracts/KSL-06-19.html`.

Getoor, L. and Taskar, B. (Eds.) (2007). *Introduction to Statistical Relational Learning*. MIT Press, Cambridge, MA.

Gómez-Pérez, A., Fernández-López, M., and Corchu, O. (2004). *Ontological Engineering*. Springer.

Jackson, Jr., L.E., Smyth, C.P., and Poole, D. (2008). Hazardmatch: an application of artificial intelligence to landslide susceptibility mapping, Howe Sound area, British Columbia. In *4th Canadian Conference on Geohazards*.

Kakas, A. and Denecker, M. (2002). Abduction in logic programming. In A. Kakas and F. Sadri (Eds.), *Computational Logic: Logic Programming and Beyond*, number 2407 in LNAI, pp. 402–436. Springer Verlag. URL `http://www2.cs.kuleuven.be/cgi-bin/dtai/publ_info.pl?id=39495`.

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 *Proceedings of the Nineteenth Conference on Innovative Applications of Artificial Intelligence (IAAI-07)*. Vancouver, BC, Canada. URL `http://www.ksl.stanford.edu/KSL_Abstracts/KSL-07-01.html`.

McGuinness, D.L. and van Harmelen, F. (2004). OWL web ontology language overview. W3C Recommendation 10 February 2004, W3C. URL `http://www.w3.org/TR/owl-features/`.

Noy, N.F. and Hafner, C.D. (1997). The state of the art in ontology design: A survey and comparative review. *AI Magazine*, 18(3): 53–74. URL `http://www.aaai.org/Library/Magazine/vol18.php\#Fall`.

Poole, D. (2007). Logical generative models for probabilistic reasoning about existence, roles and identity. In *22nd AAAI Conference on AI (AAAI-07)*. URL `http://www.cs.ubc.ca/spider/poole/papers/AAAI07-Poole.pdf`.

Poole, D., Smyth, C., and Sharma, R. (2008). Semantic science: Ontologies, data and probabilistic theories. In P.C. da Costa, C. d'Amato, N. Fanizzi, K.B. Laskey, K. Laskey, T. Lukasiewicz, M. Nickles, and M. Pool (Eds.), *Uncertainty Reasoning for the Semantic Web I*, LNAI/LNCS. Springer. URL `http://www.cs.ubc.ca/spider/poole/papers/SemSciChapter2008.pdf`.

Poole, D., Smyth, C., and Sharma, R. (2009). Ontology design for scientific theories that make probabilistic predictions. *IEEE Intelligent Systems*, pp. 27–36. URL `http://www2.computer.org/portal/web/computingnow/2009/0209/x1poo`.

Sharma, R., Poole, D., and Smyth, C. (2009). A framework for ontologically-grounded probabilistic matching. *International Journal of Approximate Reasoning*, to appear.

Smith, B. (2003a). The logic of biological classification and the foundations of biomedical ontology. In D. Westerståhl (Ed.), *Invited Papers from the 10th International Conference in Logic Methodology and Philosophy of Science*. Elsevier-North-Holland, Oviedo, Spain. URL `http://ontology.buffalo.edu/bio/logic_of_classes.pdf`.

Smith, B. (2003b). Ontology. In L. Floridi (Ed.), *Blackwell Guide to the Philosophy of Computing and Information*, pp. 155—166. Oxford: Blackwell. URL `http://ontology.buffalo.edu/smith/articles/ontology_pic.pdf`.

Smith, B. et al. (2007). The OBO foundry: coordinated evolution of ontologies to support biomedical data integration. *Nature Biotechnology*, 25(11): 1251–1255. URL `http://www.nature.com/nbt/journal/v25/n11/pdf/nbt1346.pdf`.

Xu, L., Hutter, F., Hoos, H.H., and Leyton-Brown, K. (2008). SATzilla: portfolio-based algorithm selection for SAT. *Journal of Artificial Intelligence Research*, 32: 565–606. URL `http://jair.org/papers/paper2490.html`.