Why Ontology Evolution Is Essential in Modelling Scientific Discovery *

Alan Bundy
School of Informatics,
University of Edinburgh,
A.Bundy@ed.ac.uk

Abstract

We can model scientific discovery as automated reasoning and learning, e.g., using a logic-based representation of knowledge, which we will here call an *ontology*. Much of what Kuhn calls "normal science" may be modelled as problem solving within the shared ontology of a scientific community (Kuhn 1970). However, to model what Kuhn calls a "paradigm shift", we need mechanisms for *changing* this ontology. This is what W3C call *ontology evolution*[1]. As we will see, ontology evolution can also happen during "normal science". Moreover, ontology evolution requires not just belief revision, but changes to the underlying *signature* of the ontology. For instance, functions might be split or merged; new arguments might be added to or removed from them.

Examples of Ontological Evolution in Revolutionary Science

Scientific revolutions often require not just changes in belief, but changes in the way we view the world. For instance, in 1750 Joseph Black formulated his theory of latent heat. Prior to his work, the concepts of heat and temperature were conflated (Wiser & Carey 1983). This presented Black with a paradox. As a block of ice melted, it gained heat from the environment, but its temperature stayed the same: its heat both changed and did not change. So, before Black could formulate his theory, he had to disentangle these two concepts. One old concept of heat had to become three: the heat measurable via temperature, the latent heat of fusion and their total. Similarly, anomalies in the orbital velocity in the stars in spiral galaxies, led to the conception of a new kind of matter: dark matter. Anomalies in the orbits of the planets, led to the discovery of two new planets, Neptune and Pluto, and the erroneous speculation of a third, Vulcan.

Concepts can also be merged. A well known example is the merging of the concepts of the Morning Star and Evening Star, which were re-understood as two different manifestations of Venus. Other examples are the realisation, in Einstein's equation $E = M.c^2$, that mass and energy were different manifestations of the same thing, and the unifying of the various forms of electromagnetic waves, light, radio, x-rays, etc., as a consequence of Maxwell's equations.

Similarly, there are many discoveries of unexpected dependencies and independencies of concepts on variables. For instance, Galileo's observations that acceleration due to gravity is independent of weight and that the period of a pendulum is independent of amplitude. In the other direction, Roemer showed that the speed of light was finite, so that the travel time was dependent on distance travelled. Milgrom has proposed that the gravitational 'constant' is not constant between objects of low relative acceleration. His MOdified Newtonian Mechanics (MOND) provides an alternative explanation to dark matter for the anomalous orbital velocities of spiral galaxy stars.

Ontological Evolution in Everyday Science

You might conclude that scientific ontology evolution rarely occurs, being restricted to seminal scientific revolutions and the products of extraordinary minds that are well beyond our modelling capabilities. This is wrong. In fact, ontological evolution is also a daily occurrence in normal science.

The world is infinitely rich. Any model can only capture a small finite part of that richness. Scientists and engineers make conscious (and often unconscious) decisions about what to include and exclude. Can a string be regarded as inextensible, and a pendulum as frictionless? Must the solar wind be taken into account during an Apollo moon flight? Can a situation be modelled classically, or must quantum or relativistic effects be taken into account? These modelling decisions are tentative and can be overturned by experimental evidence or changes in specification, e.g., to require more accuracy or better computational tractability. If we regard the model as an ontology, then it will evolve in similar ways to the evolution in scientific revolutions, i.e., concepts will split and merge, dependencies will be included or excluded. The main difference is that this evolutionary process will be informed by known science, so will not create the surprises endemic in revolutionary science. Similar evolutionary mechanisms of fault diagnosis and repair will, however, be at work in both cases.

*The research reported in this paper was supported by EPSRC grants EP/E005713/1 and EP/G000700/1.

[1] http://www.w3.org/TR/webont-req/#goal-evolution

Can Scientific Ontology Evolution be Automated?

In a series of projects, my research group has been exploring mechanisms for the automation of ontology evolution — our latest project being in the area of physics. More details can be found in a submitted paper to this Symposium and in (Bundy & Chan 2008). The key idea is *ontology repair plans*. These perform several functions.

- They contain a trigger formula that signals the need for the kind of ontology evolution implemented within the repair plan. It is matched to two or more ontologies, e.g., representing contradictory predictive theory and experimental evidence.
- The repair plan then controls search both within these ontologies and within the evolutionary mechanism. For instance, adding a new argument to an existing function requires a choice of which values each occurrence of this new argument should take. The repair plan directs this choice.
- The old ontologies are updated to new ones by modification of their signatures and/or their beliefs. The contradiction between them is thus removed.

Our hope is that a wide and diverse collection of case studies of ontology evolution in physics can be accounted for by a small, generic set of ontology repair plans: of the order of a few tens of such plans. We will use case studies, such as those mentioned above, to inform the development of these plans, then seek further cases studies as test examples. We are interested in the physics ontologies being developed by other participants in this symposium as a potential basis for formalising these case studies.

Should Anyone Else Care?

This might seem to be an esoteric project, which is unlikely to lead to practical applications, e.g., a robot scientist, in the near future. As can be seen from the W3C interest, however, ontology evolution is of great practical importance in the world of online, multi-agent communication. Differences in agent ontologies are likely to lead to miscommunication or communication breakdown. Ontology standardisation is an incomplete solution; differences in versions, local customisation, etc. will lead to ontological mismatches. Static ontology matching will not work in the context of vast numbers of agents in a rapidly changing environment. Dynamic, automated ontology evolution is required to produce 'good enough' repairs and matching to enable agent communication to proceed.

In an earlier project, we applied predecessors of the ideas outlined above in the context of cooperating, online agents in the ORS project (McNeill & Bundy 2007). Agent's built plans to achieve their goals by combining the services provided by other agents. Failure was diagnosed as the execution failure of an agent's plan and attributed to an erroneous model of the world, e.g. of the preconditions under which other agents were prepared to provide their services. The planning agent's ontology was then repaired after analysis of the fault and additional communication with the agent's concerned. Agents were assumed *not* to be able to directly access the ontologies of other agents, but could use fault diagnosis techniques to hypothesise ontological differences. The planning agent's ontology was repaired, e.g., by splitting/merging predicates, adding/removing arguments, adding/removing preconditions.

Evaluation of ORS was promising but hampered by the lack of sufficient records of sequences of manual ontology evolutions, especially explanations of the reason for any changes. This was one of the motivations for looking at the physics domain, where detailed historical records are available for the detection of theory *vs* experimental clashes, the resulting ontological changes and the reasoning behind them.

Conclusion

We have argued that ontological evolution is a key part of scientific discovery. It is often at the heart of paradigm shifts, but is also involved in the modelling of normal science. We have outlined a project to automate ontology evolution in physics using ontology repair plans. These repair plans trigger ontology evolution and guide the evolution of old ontologies into new ones, resolving the inherent ambiguities involved ontology refinements in the process. Maintaining a modular representation, consisting of a large number of small, interacting ontologies, is a key aspect of this approach. It is hoped that the ontology evolution techniques developed in the physics domain will transfer to other domains.

References

Bundy, A., and Chan, M. 2008. Towards ontology evolution in physics. In Hodges, W., ed., *Procs. Wollic 2008*. LNCS, Springer-Verlag.

Kuhn, T. 1970. *The Structure of Scientific Revolutions*. University of Chicago Press, enlarged 2nd edition.

McNeill, F., and Bundy, A. 2007. Dynamic, automatic, first-order ontology repair by diagnosis of failed plan execution. *IJSWIS* 3(3):1–35. Special issue on ontology matching.

Wiser, M., and Carey, S. 1983. When heat and temperature were one. In Stevens, A., and Gentner, D., eds., *Mental Models*. Erlbaum. 267–297.