Title: Why (and how) to use a metaphysicalist foundational ontology

Abstract

BORO is a metaphysically grounded foundational ontology developed specifically for use with computer systems (a foundational ontology is a system of general domain-independent ontological categories that can form a foundation for domain-specific ontologies; some but not all of these are grounded in metaphysics) and an associated methodology for legacy re-engineering systems. It emerged from a number of system replacement projects that started in the late 1980s. It was developed to mine the ontology-based conceptual models from legacy systems for use in the development of next generation systems.

Once the re-engineering methodology was established in the initial projects, questions arose as to where it could usefully be deployed. To answer this, it would help to understand why it was effective; after all, it would be hard to find a more abstract and esoteric subject than metaphysics – and one that does not immediately seem related to computing. Furthermore metaphysics is a broad subject, it would be good to understand better what areas of metaphysics are important, why they are important and how they are useful. It would also be good to have a better idea of where in computing metaphysics could play a useful role.

The purpose of this position paper is to sketch out how BORO has, over the decades, developed a view that provides answers to these questions (with no claim that this is the only way to answer them). This view is framed by two related themes. The first is that a new kind of information quality – which we label 'computerate' – is needed for computer systems and the second that metaphysics provides the right apparatus for grounding foundational ontologies that can be used to produce this 'computerate' information.

Introduction

BORO (Partridge, 1996) is a metaphysically grounded foundational ontology developed specifically for use with computer systems (a foundational ontology is a system of general domain-independent ontological categories that can form a foundation for domain-specific ontologies; some but not all of these are grounded in metaphysics) and an associated methodology for legacy re-engineering systems. It emerged from a number of system replacement projects that started in the late 1980s. It was developed to mine ontology-driven conceptual models from legacy systems for use in the development of next generation systems.

Once the re-engineering methodology was established in the initial projects, questions arose as to where it could usefully be deployed. (This led naturally to the related question whether the BORO foundational ontology could be deployed separately from its closely intertwined methodology?) To answer this, it would help to understand why it was effective; after all, it would be hard to find a more abstract and esoteric subject than metaphysics – and one that does not immediately seem related to computing. Furthermore metaphysics is a broad subject, it would be good to understand better which areas of metaphysics are important, why they are important and how they are useful. It would also be good to have a better idea of where in computing metaphysics could play a useful role.

The purpose of this position paper is to sketch out how BORO has, over the decades, developed a view that provides answers to these questions (with no claim that this is the only way to answer them). This view is framed by two related themes. The first is that a new kind of information quality – which we label 'computerate' – is needed for computer systems and the second that metaphysics provides the right apparatus for grounding foundational ontologies that can be used to produce this 'computerate' information.

In the first section below, the first theme is fleshed out using a cognitive perspective of the history of information. In the second section, the notion of a 'metaphysicalist' is introduced and used to characterise the key quality that a certain kind of metaphysical approach brings. The third section can then be regarded, depending upon one's interests, as either explaining BORO in terms of the themes of the first two sections; or, illustrating the two themes using BORO. The fourth section looks at how metaphysicalist ontologies might be deployed now, using BORO's experiences as a guide. The final section provides a brief summary.

What computerate-grade information is and why it is needed

From the start, it was clear that the metaphysical aspect of BORO's approach was very different, and surprisingly more effective, tool than the then mainstream approaches to conceptual and information modelling – which were metaphysics-free. This prompted, as noted earlier, the question why something as esoteric as metaphysics should be of value to computing. What is it about computing that connects it to metaphysics?

At the time, and still to an extent now, the typical approach to understanding the role of computing information is an economic perspective. This contrasts the 'information revolution' with the preceding agricultural and industrial revolutions; what distinguishes it is that information economic activities are becoming more important than industrial economic activities. From this perspective, there is one information revolution and technologies, such as computing, are helping drive its evolution. There is little call for metaphysics in this perspective.

However, since Ong's seminal work (Ong, 1982), there is a different cognitive perspective that offers a framework for understanding the role metaphysics could play. The key insight in this research (Olson, 1994; Ong, 1982) is that information revolutions are not just, or even primarily, technological; they are inextricably intertwined with changes in conceptual structures. This gives information a history that is characterised by a series of revolutionary changes followed by periods of consolidation (a Kuhnian (Kuhn, 1970) structure). In this perspective, computing is the latest in a number of such revolutions - the preceding revolutions including; speaking, writing, paper and printing.

Earlier we asked how computing and metaphysics might be connected. Well, from this cognitive perspective, the computing revolution should, like previous revolutions, have an intertwined change of conceptual structure. So, maybe, metaphysics – through metaphysically grounded foundational ontologies – is playing a part in this; maybe it is helping to provide the kind of new conceptual structure needed for computing.

Why (and how) to use a metaphysicalist foundational ontology

Why should metaphysics place this role and not some other discipline? As we noted earlier, it is not clear why an apparently unrelated subject such as metaphysics could be useful in computing. The cognitive perspective suggests that it is not so unrelated. Metaphysics has form; it has played a role in a historical revolution. Aristotle's categories (Aristotle, 1963) - a seminal metaphysical work – is a conceptual framework that exploited the opportunities offered by writing technology. One can contrast it with simpler, metaphysics-free lists such as the Ancient Egyptian onomastica (for example the Onomasticon of Amenemipet) which were more transient, much less successful attempts at solving the same problem.

There is also another more immediate connection with metaphysics. One might expect the introduction of new technology to mark the start of the information revolution which then leads to conceptual change. In this picture, the story of computing starts with the first computer. However Olson's (Olson, 1994) detailed analysis of the emergence of printing reveals a different pattern, one where the story starts with the development of a new conceptual framework for text which then enables the new printing technology. We can see a similar pattern in computing, where the metaphysical work of logicians and philosophers, such as Boole, Frege, Peirce, Russell and Turing (to name but a few), provided the foundations for computing technology. Computing has metaphysical roots, and so it should not be so surprising if metaphysics reappears later in its development.

We earlier raised questions of where we could use the new approach, where metaphysics might be useful. If we place computing into the historical context provided by the cognitive perspective we can see more clearly where the opportunities lie. One very simple view maps the revolutions using two rough, broad information distinctions; role and location.

Information has three broad roles; storage, processing and communication. It can also be located inside or outside the body. With these classifications, the revolutions appear as stages in the movement of information from inside to outside the body, from brain-ware to other-ware. The first stage is oral (speaking) where communication is public, outside the body, but storage and processing is private, inside the body. The second, literate (writing) where information storage can be public or private but processing is always private and communication always involves a person - whether reading, writing or speaking (or hearing). Finally, computerate where information storage, processing and communication can be public, outside the body. In this view, the innovation of the computing revolution is enabling public information processing and completely public communication (where the communication is between public systems) - what distinguishes it is the shift of processing and communication into the public domain.

One can refine this picture by using location to broadly characterise the kind of information. This highlights how revolutions create new kinds of information, new ways for information to be. So computing introduces new external kinds of processing and communication and hence new kinds of storage. We often see new technologies replacing existing ones; for example, the introduction of cars replaced horse-drawn carriages. However, at the broad level we are considering, the new kinds of information extend rather than replace existing kinds. So after writing (and printing) emerged, people still spoke to one another - oral information still had a place. Maybe it is too early to be really sure, but the pattern seems to be followed in the latest revolution; the introduction of computing seems unlikely to replace writing or speaking; although we may end up only using computers to write and read, retiring paper. In previous revolutions, it was these new kinds of information that opened up the opportunities for new conceptual structures.

This map helps us to situate where metaphysics, in the shape of foundational ontologies, is likely to play a structure shifting role; with the new kinds of computing information rather than the oral and literate kinds. Also, it is likely, at least initially, that new conceptual structures will emerge in computing systems whose scope is wider than literate systems and so involve processing and communication as well as storage. The current intractable problems with computer semantic interoperability would seem to back this up.

Earlier we raised the question of whether the two intertwined BORO components – the foundational ontology and the legacy reengineering methodology – could usefully be separated. Closer examination shows that there is a third implied component in applications of BORO, the legacy system. In BORO's early legacy reengineering work, we found that this system needed to be both of sufficient quality and size for a sufficiently high quality result.

We can link this finding with a pattern in the earlier writing revolution, which reveals its importance. Central to the exploitation of writing, is a quality improvement process where someone writes something down, she and others review it and compare it with other similar texts - and, if it is wrong, change it and then eventually integrate similar. After a number of cycles of review, the information stabilises into a more reliable resource – literate-grade information. Olson (Olson, 1994) found an interesting way to show this is a key feature of literacy. To be a literate culture, he suggested, meant not only one had writing technology but that one also had established and deployed these quality raising practices. He illustrates this with an example from Luria's (Luria, 1976): unschooled Uzbekistani peasants, representing orality, choose not to accept the premises of a syllogism, whereas school children, representing literacy, do. For us, the right way to read this is to look past the (controversial) details of the specific experiment and see (uncontroversially) that algorithmic processes such as syllogisms depend upon sufficiently high quality premises; where information is generally of poor quality, the premises are unreliable so syllogisms are unreliable. Working syllogisms are one test of writing's quality. It is these practices, these tests, that lead to the creation of literate-grade information. As Olson points out, it is only when there is sufficient literate-grade information that a culture becomes literate.

This implies that there needs to be sufficient computerate-grade information for a culture to be computerate. The mere use of computers is not enough. At one level this is obvious; we know that computing systems are sensitive to poor quality information; this is reflected in phrases such as 'Garbage in, Garbage Out'. Interestingly, this idea even appears at the birth of computing: Charles Babbage (Babbage, 1864) writes: "On two occasions I have been asked, "Pray, Mr. Babbage, if you put into the machine wrong figures, will the right answers come out?" ... I am not able rightly to apprehend the kind of confusion of ideas that could provoke such a question."

What is less clear is what good computing data – computerate information – looks like – and how it differs from literate information? In what way does it need to be of a higher level of quality? One can think of two practical criteria. Firstly the data should be of a quality that they can be processed correctly by the computer system over its life - in other words, that the right data goes in and the right data come out. Secondly, the data should be of a quality that it is processed correctly across computer systems. Notice, the first is a processing criterion and the second a communication criterion.

Why (and how) to use a metaphysicalist foundational ontology

What do these criteria tell us about how to create computerate information; about what practices lead to good computing information. What is computing's equivalent of writing's review practices? We know our current practice for building new systems does not meet the first criterion. We do not (yet) seem able to design new systems that work without significant testing and we know from experience that this is only partially adequate, as it does not find all the bugs. So currently, as least, one counterpart to writing's reviews is testing. This needs to be interpreted broadly, as the live running of some kinds of systems is a natural test; particularly systems with fully automated process with reasonably critical outcomes.

If the evolution of computing follows the same pattern as writing, then there is likely to be the need for a number of cycles to bring data up to a reasonable quality. And that if one wants to build a computerate-grade ontology, one needs inputs of sufficient quality. So, choosing the right systems to reengineer is important. Examples of good systems to choose would be payment systems or automated teller machines (ATMs). If the payment system of a retailer (such as Amazon) regularly debited its customers' accounts the wrong amount - or a bank's ATM gave the wrong amount of cash - then they would soon go out of business. Where these systems have been running for a while, they are likely to have data that meets the first criterion. Where the operation of the system does not test the data, then running it will not raise its quality. It is also likely that some kinds of data content are more amenable to computing than others, so easy to raise to a computerate level of quality. The ubiquity of ERP (including accounting) computer systems suggest that their kind of data may be more suitable for computerisation. Ironically, early cuneiform writing (Olson, 1994) focused on similar ERP type functions – so there is another repeating pattern. The lack of computer systems dealing with, for example, 'beauty' and 'happiness' suggests that these may be more difficult to computerise.

The second criterion, processing correctly across computer systems, is much more difficult. There are well-known problems with the semantic integration of systems. This implies that our computing systems are immature and consequently there is an opportunity for improvement.

This framing helped us to develop an understanding of the need for computerate-grade information. The notion of a metaphysicalist, in the next section, provides us with a way of seeing how this need might be met.

What it means to be a metaphysicalist?

BORO includes a metaphysically grounded foundational ontology that has been used to build what are traditionally called conceptual models for computer systems. Much current conceptual modelling takes by default a stance that gives metaphysics no explicit role. We can rationalise these stances - the ways conceptual models can be - as formalist and metaphysicalist. (These are inspired by similar stances in Agustín Rayo's 'The Construction of Logical Space' (Rayo, 2013).)

The formalist stance makes linguistic commitments. It sees each computer system as a system of formal representations that reflect its view of the world - where the only constraint is that the representations are sufficient for the computer to perform in the way that is required. It is only when systems need to communicate that one needs to work out (formally) how to translate from one system of representations to another.

Why (and how) to use a metaphysicalist foundational ontology

The metaphysicalist stance adds metaphysical and linguistic commitments. Firstly the metaphysical (Rayo, 2013 p. 6): "A metaphysicalist believes that there is a ‘metaphysically privileged’ way of carving up reality into its constituent parts." Then the linguistic: there is a certain kind of correspondence (possibly one-to-one) between the (logical) form of the ontological model and the metaphysical structure of reality (i.e. the metaphysically privileged carving of reality into constituent parts). In other words, a metaphysicalist expects the ontological model to reflect the way reality is carved up.

A strong argument for the formalist stance is that it is impossible (or practically impossible) to find a ‘metaphysically privileged’ way of carving. While there may be theoretical ways to undermine the argument, the simplest way to rebut it is to produce a counterexample, such as BORO and the other metaphysicalist foundational ontologies. These also provide a clearer picture of the actual trade-offs involved in the two stances.

These stances may seem too abstract to have any practical relevance, but that is not so. In the current state of computing there is a growing requirement for interoperability - for computer systems to communicate with one another. Various technical solutions such as Enterprise Service Buses are designed to facilitate this. These typically offer a standard technology for passing the messages between systems and mapping data, but do not offer a means for working out how to map the data.

The two stances discussed above suggest different kinds of solutions to this problem. To see this, consider a situation where there is a need to connect a cluster of systems and that the messages that these systems need to send and receive have been identified. Furthermore assume that the messages are sent directly from system to system, a point to point architecture. If this was in a world where all the systems had followed the metaphysicalist approach then semantic interoperability is trivial, as all the systems will have carved reality up in the same way. There will be no requirement for mapping between different ways of carving up the world. If this was a world where all the systems had followed the formalist approach, things would be more difficult. There will be a need to map between the different ways of carving – this is the current situation with most computer systems.

As the number of systems increases, the number of possible mappings increases more, making this practically very difficult for large numbers of systems. Inter-operability is a crunch point for the formalist approach. This suggests a different architecture, where one develops an intermediate interlingua for the messages. Messages from the sending system are translated into the common interlingua and then translated into the ‘language’ of the sending system - sometimes known as a ‘hub and spoke’ architecture. In this architecture, the addition of each new system only needs a mapping to the interlingua, not to all the other systems it needs to communicate with.

This hub and spoke architecture provides a way to leverage the metaphysicalist approach in the real world, where systems are typically metaphysics-free, as it opens up the possibility of developing a metaphysicalist interlingua. With this one has a degree of confidence that, while it might require extension, it will not require substantial revision as the network is extended. The metaphysical framework helps to provide stability. If however, one is a formalist, there is no underlying framework for building the interlingua, and so there is no guarantee that adding a new system will not lead to substantial changes. It is fairer to say that current experience seems to be that substantial changes are almost always guaranteed. It would appear prima facie that if a metaphysicalist approach is possible, then it would be better at supporting interoperability - certainly between large numbers of large systems. Later in the paper I explain why this position needs some qualification.

Why (and how) to use a metaphysicalist foundational ontology

One way not to carve reality

Let’s say we agree to be metaphysicalist. How do we start carving up reality in a ‘metaphysically privileged’ way? One possibility is that our perception of reality is transparent. In other words, that the untutored ‘natural’ way we see or talk about the world reflects reality directly and so the carving is clear. There are problems with this position. A well-known (within philosophy) illustration of this is in Russell’s ‘On Denoting’ (Russell, 1905). There he examines (among other things) the simple sentence “The present King of France is bald”. Taken literally, we cannot easily decide whether it is true or false, as there is a no present king of France to be bald or not bald. One lesson Russell took away from his analysis is that the surface structure of language is not guaranteed to be a good guide to the structure of reality. One needs a process for deconstructing the language to reveal an underlying structure that does reflect reality. This can be interpreted from the cognitive perspective of the previous section, as Russell starting to investigate how the literate structure needs to be improved to support information processing (inference).

More than one way to carve reality - metaontological choices

If ordinary discourse and language are not going to guide the way we should carve up reality, how do we proceed? The metaphysicalist stances was framed as if there was a single ‘metaphysically privileged’ way of carving up reality - an absolute reality. While it is true that many metaphysicians offer their single ‘metaphysically privileged’ way - unfortunately they do not all offer the same way. One could characterise this as a number of competing views - which one has to choose between. However, there is the possibility of a more sophisticated approach. Many philosophers explain the situation with a neutral(-ish) map of the conceptual landscape: one that identifies the possible metaontological choices, along with the possible motivations for the choices. They also explore the dependencies between the choices that make certain combinations more attractive and others less attractive. There is reasonable general agreement on the map - the disagreements being about which choices to make. One can think of this as ontological or reality relativity - relative to these choices.

The metaphysicalist needs to make sufficient choices to end up with a single ‘metaphysically privileged’ way of carving up reality - so it is important that one is clear about what choices need to be made. Once the choices are made they characterise a canonical reality space, relative to which data can be normalised. While the literature provides a good source for these choices, the experience of working with data also provides a useful perspective, particularly on what choices have a major impact on computer information. During the early development of BORO, while reengineering legacy data, it became apparent which choices were important in tying down the canonical view of reality. This highlights the importance of building an understanding of the practical problems that arise as well as the ‘academic’ metaphysical analyses; the need for empirical as well as rational investigation. The kinds of choices that were salient during the development of BORO are described in detail in (Partridge, 2002) – for information, the choices are listed below.

1. Absolute versus relative space, time and space-time
2. Modally extended versus unextended individuals
3. Materialism and non-materialism
4. Extensionalism versus non-extensionalism – I – Universals
5. Extensionalism versus non-extensionalism – II – Particulars
6. Topology of time – branching or linear.

Why (and how) to use a metaphysicalist foundational ontology

The choices listed in the previous section can seem esoteric; certainly some of them look very abstract and maybe, for some people, obscure. In this situation, it is a good idea to have some explicit criteria for assessing the choices. One (rational) example of criteria is Paul (Paul, 2012), who surveys how philosophy uses various techniques to make what is called an ‘inference to the best explanation’. Another (empirical) example is Kuhn (Kuhn, 1970), who examined successful scientific revolutions to come up with a similar but more extensive list. One can justify the use of scientific criteria by noting that computer systems can be seen, like scientific theories, as theories about the world (Naur, 1985). During BORO’s development, we found that having Kuhn’s criteria - summarised below - helped us to better understand the choices we were making.

1. Generality: where the scope of the improved theory increased.
2. Simplicity: where the improved theory is less complicated (it is typically more ‘deeply simple’ in the complexity theory sense).
3. Explanatory power: the ability of the improved theory to give increased meaning.
4. Fruitfulness: the ability of the improved theory to meet currently unspecified requirements or to be easily extendable to do so.
5. Objectivity: the ability of the improved theory to provide a more objective (shared) understanding of the world.
6. Precision: the ability of the improved theory to give a more precise picture of the world.

So far this section has mainly drawn on work from philosophy. But the narrative has arrived at a situation that would be familiar to business system architects, who recognise that complex requirements, such as the choices and criteria presented here, need architecting,. The architect’s job is to look at the trade-off between the different choices to arrive at a reasonably optimal solution. Underlying this is the recognition that quality systems come from quality architectures. In many ways, building the right foundation from the metaontological choices is a typical business architecting task.

BORO’s experience

This section looks at the BORO approach mainly in the light of the themes of the first two sections - hoping thereby to illustrate them.

BORO was originally developed as a tool for legacy reengineering computer systems but is now being used as a basis for interoperability between systems. Interestingly, these can be seen as similar kinds of problems; legacy system migration as diachronic – over time – interoperability and system messaging as synchronic – at a time - interoperability. Both have similar requirements focus for much improved semantic interoperability.

At one level, BORO can be characterised by its metaontological choices - this is done in the table below. More details on BORO’s choices can be found in (Partridge, 2002).

Metaontological choice	**BORO**
1. Absolute versus relative space, time and space-time	Relative space-time
2. Modally extended versus unextended individuals	Modally unextended individuals
3. Materialism and non-materialism	Materialism
4. Extensionalism versus non-extensionalism – I – Universals	Universal Extensionalism
5. Extensionalism versus non-extensionalism – II – Particulars	Particular Extensionalism
6. Topology of time – branching or linear.	Topology of time – linear.

However there is a better way of summarising the architecture. This is in terms of the metaontological concerns that motivate the choices as this gives a clear picture of the how choices, dependencies and trade-offs are seen to inter-relate. And, in turn, these concerns can often be grouped under the tradition that raised and explored them. One cannot find much discussion of these types of concern, and hence their provenance, within the current computing disciplines.

BORO reconnects with the philosophical tradition that led to computing, with the tradition that started with Frege and Russell, in particular the strand influenced by Quine and his student David Lewis – and so inherits their concerns. There is a key theme running through this tradition - that ordinary everyday ways of thinking need tidying up - which resonates with the idea of the need for an improvement of information quality for computing.

People working in this tradition were well aware of this theme and reflected upon it in ways that help us understand it. One point made repeatedly is that formalising information changes it. Before the introduction of computers, Frege (Frege, 1879) uses the analogy of a microscope and the eye to explain how his concept-script compared with ordinary language, citing in particular the need for sharpness of resolution; Carnap (Carnap, 1928) talks about 'rational reconstruction' (rationale Nachkonstrucktion). Quine (W. V. Quine, 1960) says that one doesn't merely clarify commitments that are already implicit in unregimented language, rather that one often creates new commitments by regimenting. If this correct, then a formalist should expect regimentation to result in different ways of carving up reality.

Another common point is that if one introduces ontological concerns to the formal concerns, then the regimentation needs to be far stricter and leads to, as Quine notes, more radical 'foreign' differences (W. Quine, 1981, p. 9-10): "Ontological concern is not a correction of lay thought and practice; it is foreign to the lay culture, though an outgrowth of it" Adding "There is room for choice, and one chooses with a view to simplicity in one's overall system of the world." Lewis (D. Lewis, 1986, p. 133-5) following the theme of 'outgrowth', argues that the foreignness is a result of taking the lay common sense seriously, by trying to make it simpler and consistent. This bears a close resemblance to Kuhn's (Kuhn, 1970) picture of scientific revolutions as not only new, simpler and better but also leading to fundamentally different ways of seeing existing scientific data. So a metaphysicalist that follows this tradition would expect to carve up some parts of the world in a radically different, 'foreign' way. This is illustrated below with two examples - identity and parsimony. These show how taking common sense seriously and regimenting it exposes inconsistencies whose resolution leads to new ways of carving the world. These also illustrate the variety of relations between concerns, choices and criteria. The identity example has a simple direct link with the choices 4 and 5 (Extensionalism versus non-extensionalism – Universals and Particulars) which are choices for identity. The parsimony example is general relating to all choices, but closely related to only one criterion – the second in the list above, Simplicity.

Finding Suitable Criteria of Identity

Identity (and difference) have been a topic of research since the beginning of philosophy. In 1884, Frege (Frege, 1980) introduced the notion of a criterion of identity into philosophical terminology. This has been the subject of much debate and analysis since - for example, Wittgenstein (Wittgenstein, 1953). Frege's criterion of identity involved, like the metaphysicalist position both a linguistic and an ontological component. It stated that given two names (the linguistic component) the criterion of identity would, in principle, tell one whether these two names refer to the same thing (the ontological component). For a metaphysicalist, this principle is straightforward. For a formalist, it is less so. She would need to restrict the two names to a single formalisation, as different formalisations reflect different views of the world (this point is stressed by Rayo (Rayo, 2013).

Anyone who has worked in information modelling will appreciate how useful the metaphysicalist approach can be. In modelling teams, a significant amount of time is often spent trying to determine whether two modellers mean the same thing by the same term. Typically recourse is made to each modeller's (private) understanding. With a metaphysicalist criteria of identity, there is a public decidable mechanism for resolving this.

The tricky task is finding a suitable general criterion of identity. For example, one of the problems of identity for physical objects is how to account for identity over time. If I have a different weight and height now from when I was a child, in what sense is me-now identical with me-as-child. Quine (W. V. Quine, 1960) proposed as a solution that we regard physical objects as temporally and well as spatially extended - and that the spatio-temporal extent is sufficient as a criterion of identity. This proposal has been picked up and developed by philosophers such as David Lewis (D. Lewis, 1986), Mark Heller (Heller, 1990) and Theodore Sider (Sider, 2001). BORO has followed in their footsteps.

This manoeuvre also solves a corresponding problem for identity at the type level. We talk of types, such as 'chair', acquiring and losing instances over time - as individual chairs are made and destroyed. However, if the instances of the type are spatio-temporal, then there is no need to index instantiation with a time - the spatio-temporal elements are just instances simpliciter. While this made types immutable, one further manoeuvre was need to make them extensional; this was possible world semantics developed by Lewis (D. K. Lewis, 1968). With this intension was collapsed into extension and types gained an extensional criterion of identity; two types are identical if they have the same extension - in other words, the same instances - across all possible worlds. BORO has followed Lewis's lead.

Ontological Parsimony

Ontological parsimony can be seen as a form of simplicity. It is often met as Ockham's Razor: 'entities are not to be multiplied beyond necessity'. Though one can find versions of it in Aristotle's Posterior Analytics (Aristoteles & Barnes, 1993) and Aquinas (Aquinas, 1997). In modern times, it has been closely associated with the scientific method and, as noted earlier, is the second item in Kuhn's list of criteria above.

The Quine-Lewis tradition has taken a broadly parsimonious (deflationary) path. One interesting minor example, which links to the previous example of identity, is with regard to regions. In ordinary language we talk of things occupying a region. We say 'this cup occupies this space'. It is tempting to reify regions or locations as a separate ontological category from physical objects. However, this has practical problems as, in ordinary language, we also use physical objects as locations: for example, we say "The house is located in the city of London." The connection with identity is that regions would seem to be a counter-example to the claim that spatio-temporal extents are a good criteria of identity - as the spatio-temporal extent of a physical object will be exactly the same as the spatio-temporal extent of the region it occupies. There are two obvious ways out of this. One could amend the criterion to include the ontological category - so to be identical would be to be both the same ontological category and to occupy the same spatio-temporal extent. Or one could deflate the two categories into one - a manoeuvre originally proposed by Leibniz. Under this view, the spatio-temporal extent of a physical object is identical with the spatio-temporal extent of the region it occupies. The latter option is plainly more parsimonious and Lewis adopts it - and BORO follows. (This metaphysical thesis should not be confused with the physics thesis of absolute (substantial) spacetime - see (D. Lewis, 1986, footnote 55 p. 76).

This example neatly shows how one is faced with trade-offs. One could start by trying to rescue the apparent distinction in ordinary language between regions and objects by positing two categories of regions and objects. However, as noted above, this fails to capture many of the ways we talk of location - and further distinctions would be necessary to deal with these. A more radical solution of deflating the distinction has the benefit of dissolving all these problems, but has the downside of being different, 'foreign'. Our view is that it is worth coping with the foreignness to get the simplicity. In part this decision is motivated by the way foreignness, a new way of seeing things, is has been identified as a mark of successful scientific conceptual revolutions.

How are metaphysicalist ontologies going to emerge and evolve?

If we accept the picture so far, this leads to an important question, which can be asked generally or specifically. General: how are metaphysicalist systems going to emerge and evolve? Specific: how should the metaphysicalist ontologies, such as BORO, be deployed? One would expect the answer to these to align, so an answer to the second - which we give below - should give some idea of an answer to the first.

The views on where ontologies can and should be deployed has been shaped by BORO's experiences, particular on the ways in which ontology projects can be designed so that they can compete with traditional projects in terms of cost and risk. These have led to the revision of a number of assumptions in the traditional approaches to conceptual modelling, three of which are described below.

The first assumption is that the design process for new systems starts with a conceptual model of the domain. This approach has had a chequered history and is not often adopted now. Initially, we developed a hypothesis that one reason for the general lack of success was that the conceptual model was not grounded in a foundational ontology. Frustratingly, we found that when we built our models using our foundational ontology, we were not particularly successful - the systems typically did not work without significant revisions. What we came to realise was that there were two challenges, two stages of maturity. If there was no system in place, the first challenge was to formalise the system. This was usually not trivial. The second challenge was to metaphysicalise the system. (Note these reflect the formal and metaphysicalist stances mentioned earlier.) Trying to overcome both challenges at the same time was normally not effective, but taking one challenge at a time worked well. Initially, we would formalise the system, get a working prototype - and when this was reasonably robust, we would then use our legacy re-engineering techniques to mine the metaphysicalist ontology. So, our revised hypothesis is that new systems should be built in two main stages. An initial, maybe conceptual modelling free, formalist stage, whose purpose is to raise the data to a formalist level of quality. Iterative techniques such as the currently fashionable agile (Beck et al., 2001) that build the system in small stages are useful here. Then a second stage that re-engineers the first stage to a metaphysicalist level of quality using a foundational ontology.

The second assumption is that it makes better economic sense for the conceptual model to only provide a high level design of the domain - one that typically describes the major classes/types and leaves out the detail of minor class and all the instances. In our legacy re-engineering, we realised at an early stage that the data did not match well the data schema - and that so we needed to start with the data. We also found that the provenance of the data was important, well-tested data were significantly easy to reengineer than poorly tested data - and the results were of much higher quality. It was also important to get a wide range of data – from a multitude of systems if possible. Given that we typically partially automated the reengineering, we found it feasible to assimilate large volumes of data - indeed, we found this was necessary to test the model. In our typical process, we start with the data and build a full detailed model that included all the classes and all their instances. This changed the traditional development cycle; for example, the data migration that is usually done at the end of the development is done at the beginning. This is, in principle, a more economical way of structuring the design, as it significantly reduces the likelihood of conceptual modelling faults being found at later stages of the development.

The third assumption is that the conceptual model should be produced during the design of a system. This may have been a reasonable assumption in the past when a significant amount of the IT budget was spent on new (greenfield) projects. Nowadays, the majority of IT spend is on maintaining, extending and connecting existing (brownfield) systems or installing new packages. Enterprises are wary of the expensive and risky projects that involve building new systems or replacing existing one. In this environment, the opportunities to build conceptual models to raise the quality of the data come when connecting systems. This could be driven by a desire to build a common picture across systems or the need to migrate data from one system to another. This is currently regarded as a particularly intractable problem, so a principled approach to solving it is attractive. In these brownfield cases, there is no need for the first ‘agile’ stage described above as the data already exists; the project can start with the legacy re-engineering of the systems to be connected. This is not to say that in the longer term, as the conceptual models prove themselves, they could then be used (reused) to build new systems.

Overall, the picture that emerges, at the current time, is of an iterative, reengineering approach that focuses on data rather than systems. One where the goal is to establish data that is of the right ontological quality.

Summary

This paper motivates the use of metaphysically grounded foundational ontologies and associated re-engineering methodologies in computing projects. It bases this on two related arguments. It firstly makes an argument for the need for computerate-grade data in the context of information history and the latest computing revolution; data engineered to support computer processing and communication. It then makes an argument that a kind of metaphysically grounded foundational ontologies, metaphysicalist foundational ontologies, are particularly suited to act as the foundation for computerate data as they simplify the challenge of interoperability. It explains the relativity challenge facing the metaphysicalist and offers a solution.

It then briefly describes an example of a metaphysicalist ontology – BORO – illustrating how it deals with the relativity challenge. Finally it offers some suggestions on how metaphysicalist ontologies are going to emerge and evolve based upon BORO’s experiences. It highlights three assumptions in the traditional approaches to conceptual modelling that could usefully be revised to facilitate the adoption of metaphysicalist ontologies.

References

Aquinas, T. (1997). *Basic writings of st. thomas aquinas* Hackett Publishing.

Aristoteles, & Barnes, J. (1993). *Posterior analytics* Clarendon Press.

Aristotle. (1963). *Categories, and de interpretatione* (J. L. Ackrill Trans.). Oxford: Clarendon Press.

Babbage, C. (1864). *Passages from the life of a philosopher* Longman, Green, Longman, Roberts, & Green.

Beck, K., Beedle, M., Van Bennekum, A., Cockburn, A., Cunningham, W., Fowler, M., . . . Jeffries, R. (2001). Manifesto for agile software development.

Carnap, R. (1928). *Der logische aufbau der welt* Berlin-Schlachtensee.

Frege, G. (1879). Begriffsschrift, a formula language, modeled upon that of arithmetic, for pure thought. *From Frege to Gödel: A Source Book in Mathematical Logic, 1931*, 1-82.

Frege, G. (1980). *The foundations of arithmetic; a logico-mathematical enquiry into the concept of number* (2.th ed.). Oxford,: Blackwell.

Heller, M. (1990). *The ontology of physical objects. four dimensional hunks of matter* Cambridge University Press.

Kuhn, T. S. (1970). *The structure of scientific revolutions* (2dth ed.). Chicago,: University of Chicago Press.

Lewis, D. (1986). *On the plurality of worlds*. New York, NY: B. Blackwell.

Lewis, D. K. (1968). Counterpart theory and quantified modal logic. *The Journal of Philosophy,* , 113-126.

Luria, A. R. (1976). *Cognitive development: Its cultural and social foundations* Harvard university press.

Naur, P. (1985). Programming as theory building. *Microprocessing and Microprogramming, 15*(5), 253-261.

Olson, D. R. (1994). *The world on paper: The conceptual and cognitive implications of writing and reading*. Cambridge; New York: Cambridge University Press.

Ong, W. J. (1982). *Orality and literacy: The technologizing of the word*. London and New York: Methuen.

Partridge, C. (1996). *Business objects: Re - engineering for re - use*. Oxford: Butterworth Heinemann.

Partridge, C. (2002). *LADSEB-CNR - technical report 06/02 - note: A couple of meta-ontological choices for ontological architectures.* (). Padova: The BORO Program, LADSEB CNR, Italy.

Paul, L. (2012). Metaphysics as modeling: The handmaiden's tale. *Philosophical Studies, 160*(1), 1-29.

Quine, W. (1981). Things and their place in theories. *Contemporary Materialism,* , 199.

Quine, W. V. (1960). *Word and object*. Cambridge: MIT Press.

Rayo, A. (2013). *The construction of logical space*. Oxford, UK: Oxford University Press.

Russell, B. (1905). ON DENOTING. *Mind, 14*(4), 479-493.

Sider, T. (2001). *Four-dimensionalism : An ontology of persistence and time*. Oxford New York: Clarendon Press ; Oxford University Press.

Wittgenstein, L. (1953). *Philosophical investigations*. Oxford,: B. Blackwell.