

Taking a constructional turn to radically enrich a top-level ontology's foundation: A case history

Chris Partridge^{1,2,*}, Andrew Mitchell^{1,2,*}, Sergio de Cesare², Andreas Cola³, Mesbah Khan^{2,4}, Justin Price^{1,2} and Alexander Hierl⁵

¹ BORO Solutions Limited, United Kingdom

² University of Westminster, United Kingdom

³ Telicent Limited, United Kingdom

⁴ OntoLedgy Limited, United Kingdom

⁵ Shell Deutschland GmbH, Germany

Abstract

We aim to establish, in *practice*, that there is at least one role for constructionalism in applied ontology by giving the case history of an early example where the foundations of a top-level ontology are constructionally refactored. What we have called 'taking the constructional turn'. The example is the BORO foundational ontology which has, over the last decade, been taking this turn. The paper starts by providing a case history of the turn: a chronological profile of the constructional turn and the radical enrichment it delivered. It then speculates on what the practical benefits are – providing an economic justification for the turn. To get traction, this speculation looks at a wider context, the evolution of computing. It argues that there is evolutionary evidence that the turn shapes the form of data to help enable the levels of data fidelity required for the interoperability of computer systems.

Keywords

BORO Foundation ontology, constructionalism, constructional turn, information evolution, cultural evolution, Lamarckian variation, top-level ontologies

1. Introduction

FOUST VIII has encouraged papers on what roles constructionalism could play in applied ontology. Asking what it means, in *practice*, to have a constructional ontology and, crucially, what *practical advantages* adopting such an approach brings. In applied ontology, constructionalism is an emerging practice, with few examples. There is, as far as we know, currently only one top-level constructional ontology, the BORO Foundational Ontology. This has, over the last decade, shifted toward a constructional foundation – what we have called here

Proceedings of the Joint Ontology Workshops (JOWO) - Episode X: The Tukker Zomer of Ontology, and satellite events co-located with the 14th International Conference on Formal Ontology in Information Systems (FOIS 2024), July 15-19, 2024, Enschede, The Netherlands

* Corresponding author.

✉ partridgec@borogroup.co.uk (C. Partridge); mitchella@borogroup.co.uk (A. Mitchell); s.decesare@westminster.ac.uk (S. de Cesare); andreas.cola@telicent.io (A. Cola); khanm@ontoledgy.io (M. Khan); pricej@borogroup.co.uk (J. Price); Alexander.Hierl@shell.com (A. Hierl)

🆔 0000-0003-2631-1627 (C. Partridge); 0000-0001-9131-722X (A. Mitchell); 0000-0002-2559-0567 (S. de Cesare); 0000-0002-1327-6263 (M. Khan); 0000-0003-3545-4355 (J. Price)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

the ‘constructional turn’. This shift introduces radical changes that have reinforced rather than altered the basic form of the ontology – yielding substantial benefits.

The first two sections provide a background context. In the third section, we take advantage of having the BORO Foundational Ontology as an example to come up with a pragmatic, empirical picture of the *practice* of this ‘turn’. This naturally leads us to adopt a case history approach, where we illustrate the turn using historical, publicly available documents.

Within this case history approach, we have reported on the benefits identified in the documents. However, these tend towards formal – rather than practical – benefits. So, in the fourth section of the paper, we suggest what the *practical* benefits are – providing an economic justification for the turn. To get traction, we look at a wider context, the evolution of computing. Within this context, we argue that there is evolutionary evidence that the levels of data fidelity required for computer systems interoperability will rely upon the formal improvements that top-level ontologies in general, and constructional top-level ontologies in particular, provide.

2. Background

To provide context for the main sections, we start with two background sub-sections. The first provides an overview of constructionalism. The second gives a context for BORO’s transformation by providing a sketch of BORO prior to its shift to constructionalism.

2.1. Background – Constructionalism

Constructionalism is not new. It has a longish history in philosophy in the twentieth century, which we outline in this section. Many philosophical and logical views in the 20th century have a constructional flavor or are outright examples of constructional approaches to ontology. Well-known examples include:

- Carnap 1928, *Der Logische Aufbau der Welt* [1]
- Goodman and Quine 1947, *Steps Toward a Constructive Nominalism* [2]
- Goodman 1956, *A World of Individuals* [3]
- Goodman 1958, *On Relations that Generate* [4]
- Gödel 1964, *What is Cantor's Continuum Problem* [5]
- Boolos 1971, *The iterative conception of set* [6]

In [6], Boolos neatly exemplifies the generative power of the set constructor by generating the whole set hierarchy from nothing.

More recently, in the late 20th and early 21st century, new ideas by Kit Fine have given a fresh impetus to constructional ontology. These can be found in his papers, including:

- Fine 1991, *The Study of Ontology* [7]
- Fine 2002, *The Limits of Abstraction* [8]
- Fine 2005, *Our Knowledge of Mathematical Objects* [9]
- Fine 2010, *Towards a Theory of Part* [10]

Jonathon Lowe (in *The Oxford Companion to Philosophy*) defined an ontology as “the set of things whose existence is acknowledged by a particular theory or system of thought.” Fine [7]

gives us a new framework for defining ontologies. One where, instead of having to define in advance Lowe's set of things in the ontology, one can define an initial collection of things – the givens – and a collection of constructors. Then, the entire ontology unfolds from repeated constructions – one could say the constructors, in some sense, enact the ontological commitments.

Fine [10] provides a seminal example of a constructional system for his unified theory of parts, which includes constructors for mereology and set theory. These are re-used in the BORO Foundational Ontology. He sees this as “a convincing demonstration of the power and beauty of the method”.

The generative approach is parsimonious in the sense that the definition of the content of the whole ontology is reduced to the givens and constructors – and the generation process. This is a different sense of parsimony from that normally associated with Ockham's razor, one that is captured by Schaffer [11]. Schaffer introduces a distinction between fundamental and derived objects and upgrades Occam's Razor to “do not multiply fundamental objects without necessity”, where the derived objects ‘cost’ less in the ontological accounting.

2.2. Background – BORO

We give here a brief overview firstly of BORO and its foundational ontology and then of its pre-constructional foundation. Over time, under different pressures, the names of the foundational components have changed. For ease of understanding in this paper, we will stick to the latest names: individuals, sets, and tuples.

BORO (an acronym for ‘*Business Object Reference Ontology*’) is one of the earliest top-level information system ontologies. Its development and deployment started in the late 1980s. This early work is described in Partridge's *Business Objects* [12]. BORO's focus was and is on enterprise modelling; more specifically, it aims to provide the tools to salvage the semantics from a range of enterprise systems building a single ontology with a common foundation in a consistent and coherent manner.

BORO was originally developed to address a particular need for a solid legacy re-engineering process. This naturally led to the development of both a top-level ontology (the BORO Foundational Ontology) and a closely intertwined methodology for re-engineering existing information systems (currently named bCLEARer) to align with the ontology. Hence, the term BORO on its own can refer to either of, or both, the ontology and the mining methodology.

Its two components are used to systematically unearth reusable and generalized business patterns from existing data. Most of these patterns have been developed for the enterprise context and have been successfully applied in commercial projects within the financial, defense, and energy industries.

BORO's foundational ontology is grounded in philosophy and has clear meta-ontological choices [13] following paths well-established in twentieth-century philosophy. These choices include perdurantism, extensionalism, and possible worlds [14]. The choices are categorized and compared with other top-level ontologies in Partridge's *A Survey of Top-Level Ontologies* [15] and West [16].

This grounding choice means BORO adopts a closer integration with philosophy than other ontologies in the information systems domain, such as for example, Bunge-Wand-Weber [17] and the Resource Event Agent Enterprise Ontology [18]. Also, unlike them, it emerged from and was developed in commercial projects rather than in academia.

The broad form of the original BORO Foundational Ontology can be seen as embodied by its three main types of objects. These are individuals, sets and tuples. The foundation was built with a clear understanding of the dependence of sets and tuples upon their components (and how this linked to their identity criteria). It was clear that sets had a set theoretic criterion of identity and that individuals were related by a part-whole relation that conformed to General Extensional Mereology (GEM) (this is noted in the survey [15] mentioned above).

In addition, there was an initial understanding that the relations associated with the main types of objects had some similarities. Partridge [12] noted the similar mereological form of the part-whole and sub-super-set relations and described the similar patterns these two relations gave rise to [12 Ch. 10] (also pointed out by Lewis in *Parts of classes* [19]). However, there was no evidence of the close unification between these exposed by the later shift to constructionalism.

3. Case history approach

As this is an early, if not the first, instance of this kind of constructional ‘turn’ for a top-level ontology, we expect this case history to help the community explore the nature of the change. Of course, as this is a single case, one should be careful when generalizing to other uses. In particular, our example of top-level ontology has extensional foundations. This may simplify constructional approaches, so there is work to be done to show, for example, whether the shift will work the same way on a top-level ontology with intensional foundations – particularly whether the grounding will work in such a natural way.

Table 1
Case History Documents

Stage	Year	Publication
Establish 2014-16	2016	<i>BORO as a Foundation to Enterprise Ontology</i> [20]
Deploy 2017-21	2017	<i>Developing an Ontological Sandbox</i> [21]
	2019	<i>Coordinate Systems: Level Ascending Ontological Options</i> [22]
	2020	<i>The Fantastic Combinations and Permutations of Co-ordinate Systems’ Characterising Options:</i> [23]
	2021	<i>A Framework for Composition</i> [24]
Roll out 2020-24	2020	<i>The Approach to Develop the Foundation Data Model for the Information Management Framework</i> [15]
	2022	<i>Core Constructional Ontology</i> [25]
	2024	<i>The ToLO IES5 Report</i> [26]

We use as the source data for the case history a series of eight papers the BORO community published (listed in the table above). We regard these papers as (broadly) evidence illustrating how constructionalism came to play a role in BORO’s top-level ontology. These have the advantage of being reasonably fixed and objective, in the sense that a single unchanging version is publicly available (we include URLs in the bibliography so they can be easily accessed). The papers typically share many common authors, so, for simplicity, in this case history we shall

not distinguish between individual authors and instead talk of BORO (aka the BORO community) as the author.

As Table 1 shows, the reports fall into three broad periods. We look at each of these three periods in their own sub-sections below.

3.1. 2014-16 – establish

BORO's interest in constructionalism was triggered by a chance meeting between Kit Fine and Chris Partridge at the FOIS 2014 (Rio de Janeiro) Conference. They discussed BORO's use of David Lewis' *Parts of classes* [19] and Fine suggested that his approach, especially *Towards a Theory of Part* [10], might provide a better answer. BORO started investigating this almost immediately.

BORO reasonably quickly developed a constructional architecture for its pre-existing ontology. Given BORO's extensional choice, Fine's work *Towards a Theory of Part* [10], on the sum and set constructors to handle mereology and set theory, provided clear guidelines. This emerging BORO architecture is briefly described in *BORO as a Foundation to Enterprise Ontology* [20]. It talks of Fine's [10] providing "a kind of ontogenesis narrative for the objects in the ontology". It describes the BORO "grounding (ontogenesis) narrative starting with a single element, the pluriverse of all possible worlds (a position Schaffer's *Monism: The Priority of the Whole* [27] calls 'priority monism') – this is in Fine's parlance the single *given*. It introduces "the generative operation of decomposition that divides an element into all its parts", noting that "[i]f we apply this to the pluriverse we then have all the elements. This operation exhausts all elements as the pluriverse and its parts are all the elements." In [10] Fine discusses both the use of the pluriverse and decomposition.

In [22] BORO describes constructors for generating sets and tuples (at this stage named 'power type-builder' and 'tuple-builder'). It notes that "BORO's re-engineering methodology is rooted in the philosophical notion of grounding" and that BORO has a "bottom-up, grounded approach" ... "[g]rounded in that the building of types and tuples is always grounded in particulars". This suggests BORO started out with an outlook that is congenial to constructionalism.

Even at this early stage (2016), the two features of the basic constructional architecture noted in the later (2022) *Core Constructional Ontology* [25], are plainly visible: it is foundational and it has a high degree of unification.

Three characteristics enable the architecture to play its foundational role. The first is its categorical completeness. That is, the architecture provides the three basic categories of objects for the top-level ontology: sets, individuals, and tuples, together with their associated hierarchical relations (element-set, sub-super-set, part-whole, and tuple-place). The second characteristic is object completeness. That is, the ontology generates all the base objects needed by the top-level ontology. This could be likened to an 'object factory' which supplies all the base objects that might be needed in any domain. The third characteristic is that the basic categories provide objects in the ontology with appropriate identity criteria. These are broadly extensional, based on the type of constructor and its input. For example, two sets are identical if and only if they are constructed from the same objects. Similarly, two individuals are identical if and only if they have the same parts and two tuples if they have the same places in the same order.

The architecture is unified in the following ways. First, it gives a common development of three key domains: sets, individuals, and tuples. Ontologies that involve sets and sum usually adopt set theory and mereology as separate theories, without an integrated development. We provide a unified treatment of individuals, sets, and tuples as *sui generis* objects. So, these three target domains arise in similar ways through construction. Second, there is a common basis for identity criteria, which are crucial for the foundational role discussed above. Identity criteria for the objects of the basic types are extensional, with differences arising from the way the objects are constructed. Third, the approach offers uniform ways of capturing key commonalities and differences among objects of the basic types. Such commonalities and differences are captured by features of the underlying constructors.

This basic constructional architecture offers several benefits, four of which we now highlight.

1. **Categorical transparency:** where categorical differences are just no more than constructional differences. In constructional approaches, different kinds of objects can be distinguished from one another by the way they are constructed. So categorical differences are explained by constructional differences. Hence, the three constructors lead to three categories.
2. **Dependency:** some objects are built from others and hence “depend on” them. This provides an explanation of ontological dependence and the associated notion of grounding.
3. **Reduction:** the ontology is built out of a single initial object and thus achieves fundamental ontological parsimony in Schaffer’s [11] sense explained above.
4. **Consistency:** construction can be a basis for consistency, avoiding paradoxes such as Russell’s, although care needs to be taken with the construction process. In this architecture’s case, this is achieved by requiring that the construction be ‘bottom-up’ in the sense that the properties of the constructed objects are determined by the properties of the inputs to the construction.

In most current top-level ontologies, individuals (sums) and sets are treated as separate and different categories if they are part of the ontology. And then set theory is regarded as far more important than mereology. However, the similarity and differences between the categories, individuals (sums) and sets, come down to their constructional differences. Given the broad similarity of their constructors, this implies a broad similarity of these categories. This in turn suggests the novel idea that set theory and mereology are not just remarkably similar in some respects but also could be equally ontologically important.

3.2. 2017-21 – deploy

With the constructional architecture established, BORO started investigating how it could be deployed. Initially looking generally at an ontological sandbox then applying it to the coordinate system domain.

In *The study of ontology* [7], Fine develops what might be called an algebra of constructors and constructional ontologies (while investigating modality in constructional ontology). He explains how one can remove – and add – constructors to create new constructional ontologies.

And how one ontology can be a subontology of another, where it has either a subset of the givens or constructors or both.

BORO's *Developing an Ontological Sandbox* [21] repurposes this 'algebra' for more pragmatic purposes. It introduces the ontological sandbox as a systematic way to investigate the ontological nature and requirements that underlie frameworks and tools. This is a conceptual framework for investigating and comparing multiple variations and components of possible ontologies – without having to commit to any of them – isolated from a full commitment to any foundational ontology. It discusses the sandbox framework as well as walking through an example (from multi-level modelling) showing how it can be used to investigate a simple ontology. The example, despite its simplicity, illustrates how the constructional approach can help to expose and explain the metaphysical structures used in ontologies, and so reveal the underlying nature of multi-level modelling levelling. It also demonstrates how a constructional approach can simplify the structures. The paper also notes that Fine (in *Towards a Theory of Part* [10]) makes similar points. He sees the advantage of his framework is that it naturally reveals the underlying metaphysical structure of reality. He also talks about its power and beauty; its ability to provide a single and elegant account of a variety of structures. BORO [21] notes that one cannot see this in the logical characterization of these structures in other non-constructional top-level ontologies. It also shows how this provides a better tool for the task of investigating the metaphysical structure of possible ontologies as well as the ontological their content.

BORO's *Coordinate Systems: Level Ascending Ontological Options* [22] describes how the constructional approach extends BORO's top-level foundation and how it can be used as an analytic tool at the domain level. The coordinate systems for multi-platform sensor data have a complex range of variations that are not fully captured in a unified structure. One example of increasing complexity is the shift from single to multi-platform sensing. The constructional approach enabled a unified multi-level picture to be built from the coordinate systems' characterizing options. It also revealed the implicit underlying fundamental structure.

BORO's *The Fantastic Combinations and Permutations of Co-ordinate Systems' Characterising Options* [23] builds upon and refines the coordinate systems' characterizing options example in [22]. It focuses on the way simple constructors can build complex frameworks, characterizing this as *radically simplifying by scaling down to scale up*. It makes an analogy with Conway's Game of Life [28], which very neatly illustrates how very simple algorithms give rise to – and so help to explain – the complex 'lives' of cell automata and so perhaps complex natural life. It describes how simple differences in the shape and operation of constructors give rise to different varieties of hierarchies and levels – and the impact this has. It looks at how the constructional approach enables derived constructors to be built from the foundational constructors; ones such as sub-type and powertype. It describes how this framework can reveal and explain the formal levels and hierarchies that underlie the options for characterizing coordinate systems.

In *A Framework for Composition: A Step Towards a Foundation for Assembly* [24], BORO provides an example of how the fundamental constructional structure can be used to derive other constructors that reveal useful features in the ontology. It shows how one can derive a new mereological constructor, based upon the existing one, which creates what it calls the assembly structure. This is the stratified breakdown structure common in engineering and other

disciplines – one that features in modules and components. It also describes how minor changes to the constructor lead to useful breakdown variants.

3.3. 2020-24 – roll out

With the clearly established uses of the constructional foundation and its extension to the domain, the refactored top-level ontology was ready to roll out. This section describes two roll-outs; the UK's National Digital Twin (NDT) project and the IES Top-Level Ontology.

The UK's National Digital Twin project decided to use a constructional ontology – this history is described first – including its selection of a BORO-based ontology. Then we describe the Minimum Viable Product [MVP] project set up to produce an axiomatic version of this foundation and show its consistency, as part of the requirements for ISO/IEC 21838 – Top-level ontologies.

In 2017, the National Infrastructure Commission published *Data for the Public Good* [29] which sets out a number of recommendations including the development of a UK National Digital Twin supported by an Information Management Framework (IMF) of standards for sharing infrastructure data, under the guidance of a Digital Framework Task Group set up by the Centre for Digital Built Britain.

An investigation (*The Approach to Develop the Foundation Data Model for the Information Management Framework* [16]) recommended that the Foundation Data Model (FDM) seed be founded on a top-level ontology based on the four 4-dimensionalist top-level ontologies (TLOs) that best met the technical requirements of the FDM, underpinned by rigorously established foundations. It further proposed that the work BORO had been doing (mentioned above) to unify top-level ontologies based on a constructional framework developed by Kit Fine should underpin the IMF's TLO, providing a firmer foundation for future development.

As part of the adoption of a constructional ontology, the NDT decided to develop an axiomatic formalization of the constructional foundation to show its consistency. A first project was set up, the results of which are documented in BORO's *Core Constructional Ontology* [25]. The project was faced with a choice on how to formalize the theory. This boiled down to two options: stage theory or procedural postulationism (a theory outlined in Fine's *Our knowledge of mathematical objects* [9]). While the latter appears to be more attuned to constructionalism, as it would have required the development of significant logical machinery, so stage theory was adopted for the time being.

In *The ToLO IES5 Report* [26], BORO provides an example of its Foundational Ontology being used a standard. The Information Exchange Standard (IES) is a standard for information exchange developed within UK Government (<https://github.com/dstl/IES4>). It is a BORO-based standard with an RDF implementation. Recently, work started on version 5, aligning it with the NDT work, including the constructional core, which is contained in the internal report *The ToLO IES5 Report* [26]. This also mentions additional work being undertaken to develop the ways in which modality can be expressed in an extensional 4D ontology.

4. Practical benefits

FOUST VIII's call rightly asks: what are the practical advantages of adopting constructionalism in applied ontology? And correctly labels this question as crucial. In the context of this paper,

the scope of the question is narrowed down to the practical advantages of taking the constructional turn described above.

From the case history perspective, we have identified the benefits mentioned in the documents. These include qualities such as categorical transparency, dependency, reduction and consistency – as well as simplicity and explanatory. However, these are, for the most part, broad theoretical advantages that might impact theory choice – as, for example, described in Kuhn’s *Objectivity, Value Judgment, and Theory Choice* [30]. While undoubtedly in some sense clearly very attractive ‘benefits’ (as noted in Fine [10] and BORO [21]), their immediate practicality is less obvious.

Here, we sketch a way to understand the practical advantages of taking the constructional turn through the lens of evolution. We do this by first establishing, in evolutionary terms, the practical advantages of top-level ontologies – providing a context for the constructional turn. Then, in that context, we explain the advantages the turn brings and why these are different.

4.1. Top-level ontology and semantic interoperability

It is well-recognized that we are not yet able to reliably engineer semantic interoperability with high levels of fidelity between systems or their data at scale. We assume that the practical advantages of semantic interoperability are obvious, so do not dwell upon them here. There is a common claim [31, 32, 33, 34, 35] that a formal ontology can help with this by providing a machine-readable interlingua warranted by its clear representation of the relevant domain. In the case of a formal (philosophical) top-level ontology, the aspiration that it can provide a clear formal representation of most, if not all, domains. The clearness of representation is guaranteed by a simple semantics where one sign represents one object in the domain – the structure of the representation cleanly mirroring the structure of the domain. If one accepts that the top-level ontology, maybe using philosophical research, is able to delineate domains (aka reality) more clearly, then the claim that what is being represented is a ‘good’ picture of reality would seem to have some merit. So, it is no surprise that claims about better representing reality are commonplace in IS modelling discussions.

The claim makes even more sense if situated in a wider evolutionary context. Research into previous major information technology revolutions, such as speaking, writing and printing, suggests that exploitation of the radical changes enabled by an emerging technology is dependent upon the co-evolution of new cultural practices that take advantage of the opportunity to represent some portions of reality better. For example, Ong [36] highlights the new practices that took advantage of the writing and printing technologies. And Olson [37], focus on printing, suggests that new cultural practices are needed to enable the emergence of printing as well as further new literary practices to exploit it. Olson then suggests that these new literary practices are contingent, offering China as a case where the technology of printing emerged, but the associated literary practices did not.

Ong, Olson and others in the community note that among the cultural challenges of an emerging information technology is finding forms that exploit its potential. For example, Olson noted that as botany started to take advantage of printing, strict systems of classification emerged. For the botanist Linnaeus:

“... the fundamental task of natural history was that of “arrangement and designation” ... Descriptions were tied to the visible, the nameable and the depictable features of plants. Each

part of the plant – roots, stem, leaves, flowers, fruits – was seen to be a product of four variables: form, quantity, manner of relation and magnitude.” [37 p. 226]

And this formal analytic process led to a corresponding new form, the formalized binomial nomenclature, which intriguingly is clearly an evolution of Aristotle’s taxonomic scheme.

Hence, it is natural to think that the currently emerging information technology, computing, will create new opportunities for co-evolving – an obvious one being the possibility for machines to talk seamlessly to machines. And also natural to think that one dimension of the co-evolution needs to be developing forms that exploit the new technology. And that, maybe, those forms may build upon ideas drawn from philosophy.

It is worth making a quick detour here and give some context for the use of the term ‘evolution’ here. Biological orthodoxy for much of the 20th century was that all evolution was genetic. However, in the last couple of decades, there has been an emerging recognition that genes are not the only inheritance information system. For example, cultures may persist through their own inheritance information systems – where this includes behaviors as well as symbolic artefacts – see Jablonka’s *Evolution in four dimensions* [38]. Clearly, symbolic cultural inheritance has evolved relatively recently in evolutionary timescales. But one could go further back, taking a macroevolutionary perspective as Szathmáry does in *The major transitions in evolution* [39], and trace the cumulative emergence over time of a multiplicity of different inheritance information systems that co-evolve. One could see the introduction of a new, different, information system as a major evolutionary transition [40, 41, 42]. One could take a further radical step and see the information technology units as biological life [43, 44]. In this richer perspective, one can see patterns of co-evolution into which the current co-evolution of computing technology and cultural practices neatly fits.

It was a dogma of 20th century Darwinian genetic evolution that it only involved natural (random) variation (without a hint of Lamarckian variation) and selection, based upon survival of the fittest. However, as Jablonka [38 Ch. 9] notes, it is clear that cultural evolution often proceeds in a kind of Lamarckian fashion, where the variations are clearly not random. They are often targeted, in the sense that they are educated guesses (in some cases even engineering designs) about what will lead to a fitter outcome. They are sometimes constructed in the sense that there is an intentional choice of what variation to fabricate.

As mentioned above, in the co-evolution of each new information technology and its associated cultural practices, one area ripe for exploitation (Lamarckian variation) is how the information is organized and particularly what form it takes [43, 44]. As also mentioned above, there have typically been evolutionary opportunities for radical innovation in this area that led to leaps in fitness, as described for printing by Olsen [37].

Given this new perspective of the co-evolution information technology, in the case of computing technology, one can see the new form that foundational ontologies introduce as a Lamarckian variation in cultural practice which aims to show its (practical) fitness by exploiting the opportunity of semantic interoperability.

4.2. Notation and the constructional turn

However, this evolutionary explanation of the practical (fitness) of a top-level ontology variation does not (yet) explain the practical (fitness) of the subsequent constructional turn. Recall, that top-level ontology is aiming to ‘crack’ the semantic interoperability challenge by providing a clearer picture of reality. Recall, also, that the case history makes clear that the

constructional turn did not change the basic categories of the BORO Foundational Ontology and so its broad picture of reality. So, if BORO's picture of reality does not (substantially) change, then it suggests that the constructional turn is not directly giving a more accurate picture of reality.

However, there is a way of explaining the way the constructional turn influences the form of the computing data. To do this we turn to the long tradition in mathematics of recognizing the importance of notation – and how some notations are better than others in various contexts. Whitehead in [48, Ch. 5], mentioned in Novaes [45] and developed in Macbeth [46], uses the simple example of Roman and Arabic numerals to make the point that notation can make a substantial difference to the ease with which one does arithmetic calculations. Whitehead says: “By relieving the brain of all unnecessary work, a good notation sets it free to concentrate on more advanced problems, and in effect increases the mental power of the race.” Interestingly, Macbeth clearly makes the point, also implicit in Whitehead, that in this case the notation *shows* rather than *describes* the calculation. Where showing seems to be a characteristic of the constructional approach.

So there seem to be two aspects to improving form. One where it more directly mirrors and reflects the content it is intended to represent. This is what the formal top-level ontology is aiming to work on. Another is what might be called the notational aspect, which has an impact on the ease of calculation and ease of use. This is where the constructional turn has an impact. One could regard the derived constructors, both the derived sub-class constructor [23] and the breakdown constructors [24], as examples of the ease of calculation.

5. Conclusion

FOUST VIII raised the question of what role constructionalism could play in applied ontology. This paper has, in a sense, answered a slightly different question, what role has constructionalism played in applied ontology – where, of course, if the role has been played it necessarily could be played. This different question requires taking a historical rather than theoretical approach.

The paper has provided an example of a particular role: providing a constructional foundation for a top-level IS ontology: what we have called ‘taking a constructional turn’. It has done this in the form of a case history of the BORO Foundational Ontology's constructional turn. Hopefully this is unequivocal evidence that such a role can be played.

This paper points out that while there are a range of very attractive benefits noted in the case history, these tend to the theoretical rather than the practical. To clarify the practical benefits, the paper provides an evolutionary perspective that recognized there is typically a pattern of co-evolution of information technology and cultural practices – where the cultural practices aim to provide a form that exploits the opportunities created by the new technology. It firstly notes that top-level ontologies are an example of a Lamarckian cultural variation in the co-evolution of computing technology – one that aims to provide a more accurate form for reality. It then points out that, at least in the case of BORO's top-level ontology, the constructional turn made no improvement to that aspect of the form. It then notes that the constructional turn did make a practical improvement to another, traditionally important, ‘notational’ aspect of the form.

Acknowledgements

We would like to acknowledge the help of Salvatore Florio, Patricia Arcenegui Calvo and Diego Alberto Zabala Betancur in preparing the paper.

References

- [1] R. Carnap, 'Der Logische Aufbau der Welt, Berlin 1928', *Kwartalnik Filozoficzny*, vol. 11, no. 2, pp. 183–188, 1928.
- [2] N. Goodman and W. V. O. Quine, 'Steps Toward a Constructive Nominalism', *The Journal of Symbolic Logic*, vol. 12, no. 4, pp. 105–122, 1947.
- [3] N. Goodman, 'A world of individuals', *The problem of universals*, pp. 13–31, 1956.
- [4] N. Goodman, 'On Relations That Generate', *Philosophical Studies*, vol. 9, no. 5–6, pp. 65–66, 1958, doi: 10.1007/BF00725420.
- [5] K. Gödel, 'What is Cantor's continuum problem?', *The American Mathematical Monthly*, vol. 54, no. 9, pp. 515–525, 1947.
- [6] G. Boolos, 'The Iterative Conception of Set', *Journal of Philosophy*, vol. 68, no. 8, pp. 215–231, 1971, doi: 10.2307/2025204.
- [7] K. Fine, 'The Study of Ontology', *Noûs*, vol. 25, no. 3, p. 263, Jun. 1991, doi: 10.2307/2215504.
- [8] K. Fine, *The limits of abstraction*. Oxford; New York: Oxford University Press, 2002.
- [9] K. Fine, 'Our Knowledge of Mathematical Objects', in *Oxford Studies in Epistemology*, T. S. Gendler and J. Hawthorne, Eds., Oxford University Press, 2005, pp. 89–109.
- [10] K. Fine, 'Towards a Theory of Part', *Journal of Philosophy*, vol. 107, no. 11, pp. 559–589, 2010, doi: 10.5840/jphil20101071139.
- [11] J. Schaffer, 'What Not to Multiply Without Necessity', *Australasian Journal of Philosophy*, vol. 93, no. 4, pp. 644–664, Oct. 2015, doi: 10.1080/00048402.2014.992447.
- [12] C. Partridge, *Business Objects: Re-Engineering for Re-Use*, 1st Edition. Oxford: Butterworth-Heinemann, 1996.
- [13] C. Partridge, 'Note: A couple of meta-ontological choices for ontological architectures', *LADSEB-CNR, Padova, Italy*, 2002,
- [14] C. Partridge, A. Mitchell, and S. de Cesare, 'Guidelines for Developing Ontological Architectures in Modelling and Simulation', in *Ontology, Epistemology, and Teleology for Modeling and Simulation*, First., vol. 44, A. Tolk, Ed., in *Ontology, Epistemology, and Teleology for Modeling and Simulation: Philosophical Foundations for Intelligent M&S Applications*, vol. 44. , Berlin, Heidelberg: Springer, 2013, pp. 27–57.
- [15] C. Partridge, A. Mitchell, A. Cook, J. Sullivan, and M. West, 'A Survey of Top-Level Ontologies - to inform the ontological choices for a Foundation Data Model.', CDBB, 2020.
- [16] M. West, A. Cook, D. Leal, A. Mitchell, C. Partridge, and J. Sullivan, 'The Approach to Develop the Foundation Data Model for the Information Management Framework', CDBB, 2020.
- [17] Y. Wand and R. Weber, 'On the Ontological Expressiveness of Information Systems Analysis and Design Grammars', *Information Systems Journal*, vol. 3, no. 4, pp. 217–237, 1993.
- [18] G. L. Geerts and W. E. McCarthy, 'An ontological analysis of the economic primitives of the extended-REA enterprise information architecture', *International Journal of Accounting Information*, vol. 3(1), pp. 1–16, 2002.
- [19] D. K. Lewis, *Parts of classes*. Oxford, UK; Cambridge, Mass.: Basil Blackwood, 1991.

- [20] S. de Cesare and C. Partridge, 'BORO as a Foundation to Enterprise Ontology', *Journal of Information Systems*, vol. 30, no. 2 (Summer 2016), pp. 83–112, 2016, doi: 10.2308/isys-51428.
- [21] C. Partridge, S. deCesare, A. Mitchell, F. Gailly, and M. Khan, 'Developing an Ontological Sandbox: Investigating Multi-Level Modelling's Possible Metaphysical Structures', in *Proceedings of ACM/IEEE 20th International Conference on Model Driven Engineering Languages and Systems (MODELS 2017)*, in CEUR Workshop Proceedings, vol. 2019. CEUR Workshop Proceedings, Sep. 2017, pp. 226–234.
- [22] C. Partridge, A. Mitchell, M. Loneragan, H. Atkinson, S. de Cesare, and M. Khan, 'Coordinate Systems: Level Ascending Ontological Options', in *2019 ACM/IEEE 22nd International Conference on Model Driven Engineering Languages and Systems Companion (MODELS-C)*, IEEE, 2019, pp. 78–87.
- [23] C. Partridge, A. Mitchell, M. Loneragan, H. Atkinson, S. de Cesare, and M. Khan, 'The Fantastic Combinations and Permutations of Co-ordinate Systems' Characterising Options: The Game of Constructional Ontology', 2020.
- [24] C. Partridge, A. Mitchell, and P. Grenon, 'A Framework for Composition: A Step Towards a Foundation for Assembly', CDBB, Apr. 2021. doi: 10.17863/CAM.66459.
- [25] S. Florio and Ø. Linnebo, 'Core Constructional Ontology: The Foundation for the Top-Level Ontology of the Information Management Framework', *NDT*, v. 1.91, Apr. 2022.
- [26] C. Partridge, A. Mitchell, O. Xiberta-Soto, and A. Cola, 'The ToLO IES5 Report', 2024.
- [27] J. Schaffer, 'Monism: The Priority of the Whole', *The Philosophical Review*, vol. 119, no. 1, pp. 31–76, Jan. 2010, doi: 10.1215/00318108-2009-025.
- [28] M. Gardner, 'The Fantastic Combinations of John Conway's New Solitaire Game "Life"', *Scientific American*, vol. 223, pp. 120–123, 1970.
- [29] NIC, 'Data for the Public Good', National Infrastructure Commission, 2017.
- [30] T. S. Kuhn, 'Objectivity, Value Judgment, and Theory Choice', in *The Essential Tension*, Chicago: The University of Chicago Press, 1977, pp. 320–339.
- [31] T. Bittner and M. Donnelly, 'Ontology and Semantic Interoperability'.
- [32] L. Obrst, 'Ontologies for semantically interoperable systems', in *Proceedings of the twelfth international conference on Information and knowledge management*, New Orleans LA USA: ACM, Nov. 2003, pp. 366–369. doi: 10.1145/956863.956932.
- [33] C. Partridge, 'The Role of Ontology in Semantic Integration', in *Second International Workshop on Semantics of Enterprise Integration at OOPSLA 2002*, Seattle, 2002.
- [34] C. Partridge, M. Lambert, M. Loneragan, A. Mitchell, and P. Garbacz, 'A Novel Ontological Approach to Semantic Interoperability Between Legacy Air Defence Command and Control Systems', *International Journal of Intelligent Defence Support Systems*, vol. 4, no. 3, pp. 232–262, 2011.
- [35] C. Partridge, 'The role of ontology in integrating semantically heterogeneous databases', *Rapport technique*, vol. 5, no. 02, 2002.
- [36] W. J. Ong, *Orality and Literacy - The Technologizing of the Word*, vol. 1. London and New York: Methuen, 1982.
- [37] D. R. Olson, *The world on paper: the conceptual and cognitive implications of writing and reading*. Cambridge; New York: Cambridge University Press, 1994.
- [38] E. Jablonka and M. J. Lamb, *Evolution in four dimensions, revised edition: Genetic, epigenetic, behavioral, and symbolic variation in the history of life*. MIT press, 2014.
- [39] E. Szathmáry and J. Maynard Smith, *The major transitions in evolution*, 1st ed. WH Freeman Spektrum Oxford, 1995.
- [40] J. Maynard Smith, 'The idea of information in biology', *The Quarterly Review of Biology*, vol. 74, no. 4, pp. 395–400, 1999.

- [41] J. Maynard Smith, 'The concept of information in biology', *Philosophy of science*, vol. 67, no. 2, pp. 177–194, 2000.
- [42] E. Jablonka, M. J. Lamb, and A. Zeligowski, *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*. in Bradford books. A Bradford Book, 2006.
- [43] C. Partridge, 'How an Evolutionary Framework Can Help Us To Understand What A Domain Ontology Is (Or Should Be) And How To Build One', presented at the FOMI 2022, 12th International Workshop on Formal Ontologies Meet Industry, 12-15 September 2022, Tarbes, France, Sep. 12, 2022.
- [44] C. Partridge, 'Why Form, and so Unification of Types, is Important', presented at the King's College London, Workshop on the Unification of Types and Multi-Level Modeling, 13 March 2024, London, UK, Mar. 13, 2024.
- [45] C. Dutilh Novaes, *Formal Languages in Logic: A Philosophical and Cognitive Analysis*. Cambridge University Press, 2012.
- [46] D. Macbeth, 'Seeing How It Goes: Paper-and-Pencil Reasoning in Mathematical Practice', *Philosophia Mathematica*, vol. 20, no. 1, pp. 58–85, Feb. 2012, doi: 10.1093/phimat/nkr006.