Formal Ontology in Information Systems
S. Borgo et al. (Eds.)

doi:10.3233/978-1-61499-910-2-33

The Treatment of Time in Upper Ontologies

Antony GALTON
Department of Computer Science, University of Exeter, Exeter, UK

Abstract. It is generally agreed that as a fundamental feature of the world, time merits treatment within an upper ontology, that is, an ontology that is designed to capture those categories which are sufficiently generic to transcend the specific subject matters of any particular domain ontologies. In this paper I examine how time is handled within three well-known upper ontologies (BFO, DOLCE, and GFO), and follow this with a discussion of three key issues emerging from the survey, namely dimensionality (the treatment of instants and intervals), frame-dependence (as required by the Theory of Relativity), and indexicality (the status of past, present and future). The overall conclusion is that while existing upper ontologies tend to adopt some kind of compromise between a supposedly objective, scientific account of physical time, and a more subjective, cognitive account of time as we experience it, the fundamental philosophical and scientific questions concerning the nature of time have scarcely been addressed by any of them.

Keywords. time; upper ontology; BFO; DOLCE; GFO; instants and intervals; relativity; specious present; past, present and future

1. Introduction

The purpose of an upper ontology is to provide a systematic account of those categories—whether these be construed as categories of thought or categories of existence—which are sufficiently generic to transcend the specific subject-matter of any particular domain ontologies. The categories of an upper ontology should, in principle, be applicable to whatever domain one is considering, and therefore it is thought that a properly constructed upper ontology should be able to provide a basis, or starting point, for the development of applied ontologies in any domain.

Amongst the topics generally regarded as within the purview of upper ontologies, time and temporal phenomena have always enjoyed a particular prominence. Except in the unchanging world of pure mathematics, time features in discourse about any subject matter because everything is liable to undergo change. For this reason, there has been much theorising about time both within applied ontology and also in various fields impinging on this, such as philosophy, physics, mathematics, linguistics, and psychology.

In this paper I examine how time is handled within three well-known upper ontologies (BFO, DOLCE, and GFO), and follow this with a discussion of a number of key issues emerging from this survey. It is not my aim in this paper to present a new specific theory of time, nor to review all the existing approaches, but rather to provide indications of areas where a rethinking of some widely-accepted tenets of temporal ontology may be worth considering.

2. Time in BFO, DOLCE, and GFO

2.1. Time in BFO

Basic Formal Ontology (BFO) [1,16] posits a top-level division of entities into two broad categories, designated as *continuant* and *occurrent*. Temporal entities are handled under the category *occurrent*, the subcategories of which include both *spatiotemporal region* and *temporal region*.

2.1.1. Temporal and spatiotemporal regions in BFO

A spatiotemporal region, according to Arp *et al.* [1], is "an occurrent entity at or in which occurrent entities can be located". They go on to say that "one can think of each process as a temporally extended continuum, a spacetime *worm*, stretched out in and through the single unified container that is the entirety of spacetime".

A temporal region, on the other hand, is "an occurrent entity that is a part of time (of the whole of time). Temporal regions differ from spatiotemporal regions in that they are extended or serve as boundaries only along the temporal dimension. A temporal region is the result of projecting a spatiotemporal region onto this temporal dimension".

Although on the surface all this might seem clear enough, on closer examination it appears deeply puzzling. This puzzlement is only compounded when we turn to some of the axioms the authors propose for regulating these concepts, notably:

- Every material entity exists at some temporal interval.
- Every occurrent occupies some spatiotemporal region.
- Every spatiotemporal region occupies some temporal region.

These axioms use the notions of "existence at a time", "occupying a spatiotemporal region" and "occupying a temporal region". Regarding the third of these, it is also stated that every temporal region *occupies itself*.

So what can a temporal region be, that a spatiotemporal region can occupy it? We are told that a temporal region is the result of projecting a spatiotemporal region onto the temporal dimension; this does not help, however, unless we know how the temporal dimension is related to spacetime, and just what is meant, in physical terms, by this operation of projection. I do not think it would be helpful here to refer to the mathematical operation of projecting a four-dimensional space onto a one-dimensional space—although this is obviously the inspiration for what Arp *et al.* say—since these operations are defined on abstract sets whereas what we are interested in here is the nature of space, time and spacetime as elements of physical reality.

One possibility that might be considered is that temporal regions are *parts* of spacetime. There seem to be just two ways in which we might take this.

1. Perhaps a temporal region is a *slice* of spacetime, as shown in Figure 1(a). On this view, although a "one-dimensional" temporal region is strictly four-dimensional, it is uniquely determined by fixing its temporal coordinates, allowing just one degree of freedom. In this case the relation 'occupies temporal region' can simply be understood as parthood restricted to pairs in which the second element is a temporal region.

2. Alternatively, a temporal region could be a one-dimensional subspace of spacetime whose points differ only with respect to their time-coordinate: see Figure 1(b). In this case, for a spatiotemporal region R to occupy a temporal region is for it to have that temporal region as a part which is maximal in the sense that it is not a proper part of any other temporal region that is part of R. Although this is a possible interpretation of a BFO temporal region (since the axiom does not say that every spatiotemporal region occupies *only one* temporal region), it is probably not what is intended by that notion.

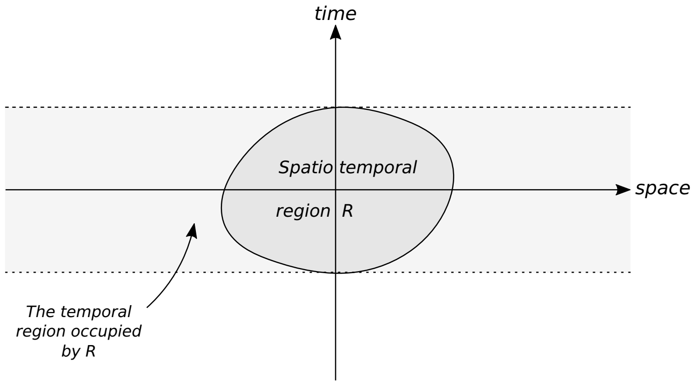

(a) Temporal regions as four-dimensional slices of spacetime.

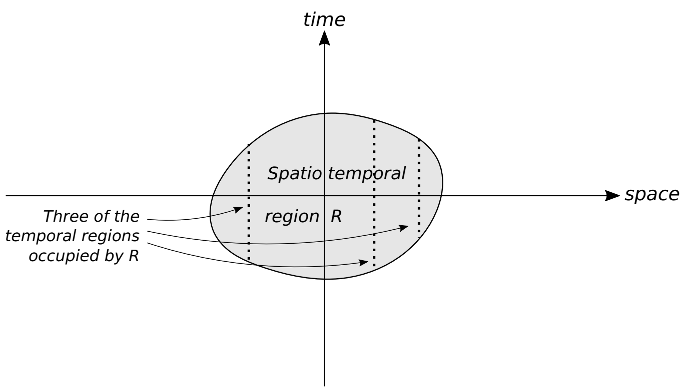

(b) Temporal regions as one-dimensional parts of spacetime.

Figure 1. Two possible interpretations of 'temporal region' in BFO.

Apart from these two pictures, there does not seem to be any other way to understand the idea that spatiotemporal regions occupy temporal regions, if 'occupy' is to be understood as connoting some form of co-location.

Any alternative conception, to allow for the existence of temporal regions as entities within the ontology, must require time to be a one-dimensional extent quite separate from space-time. Mathematically, the obvious way to "create" individual times from space-time is to treat them as equivalence classes under a relation of simultaneity ("simultaneity classes"), ordered in the obvious way. This conception is different from the "slice" conception because it makes times *abstract* entities rather than physical parts of space-time. This may be problematic for BFO, which does not appear to recognise a category of abstract entities. If this line of thought is followed, then a rather more complicated account of what it is for a spatiotemporal region to occupy a temporal one has to be given; and occupancy ceases to be a form of co-location.

Both the "slice" conception and the "simultaneity class" conception presuppose a fixed spatiotemporal frame; specifically, they assume a simultaneity relation that links space-time points at arbitrarily large spatial separations. This is in conflict with the Special Theory of Relativity (STR), in which such a relation can only be defined relative to some inertial frame, so that observers in relative motion to one another will disagree about which sets of events are simultaneous. Arp *et al.* acknowledge frame-relativity in connection with spatial regions [1, pp.115f], and later state that temporal regions also require reference to a frame; but nothing further is made of this. How far should an upper ontology go to incorporate the results of current scientific theories, some of which may be, at least on the surface, incompatible with our everyday "commonsense" ideas about the world? I shall return to this question in §3.2, with specific reference to STR.

2.1.2. Dimensionality of temporal regions in BFO

Turning now to the matter of dimensionality, the category of *temporal region* in BFO is divided into *zero-dimensional temporal region* and *one-dimensional temporal region*. Instances of the former category are called *temporal instants*. Of these, Arp *et al.* [1] say:

> Zero-dimensional temporal regions are the temporal regions that process boundaries are located in. Examples include right now, the moment at which a finger is detached in an industrial accident, the moment at which a child is born, the moment of someone's death, and the turn of the nineteenth century.

Some of these examples, but perhaps not all, involve process boundaries. The detachment of the finger is a process which takes time, albeit possibly very short; the moment referred to must be the time when that process comes to completion, when the last connection between the finger and the rest of the hand is severed. This final severance is a process boundary because it marks the termination of the process of severance. Clearly we do not have the technological means, even in principle, to assign to that event a precise numerical value to designate the instant at which it occurs: any measurement or observation of the event can at best assign it to some short interval which we are capable of designating within whatever time-measurement system we use. The supposition that there is nonetheless, in the nature of things, an instantaneous event occurring at a durationless temporal point, is therefore an idealisation which can only be justified to the extent that it plays a useful role in the battery of tools and techniques we use for describing the world. This is an issue that I shall return to in §3.1; for now, I merely note it.

Another of the examples, 'right now', is problematic in a different way, in that it suggests that the *present* (and therefore as a corollary the *past* and *future* as well) is an entity of the right kind to populate a temporal ontology with. In view of the avowed realism of BFO this in turn implies (if the example is to be taken seriously) that users of BFO are expected to be realists with regard to tense. This has long been a contentious issue in the philosophy of time, with many philosophers arguing that the distinctions of

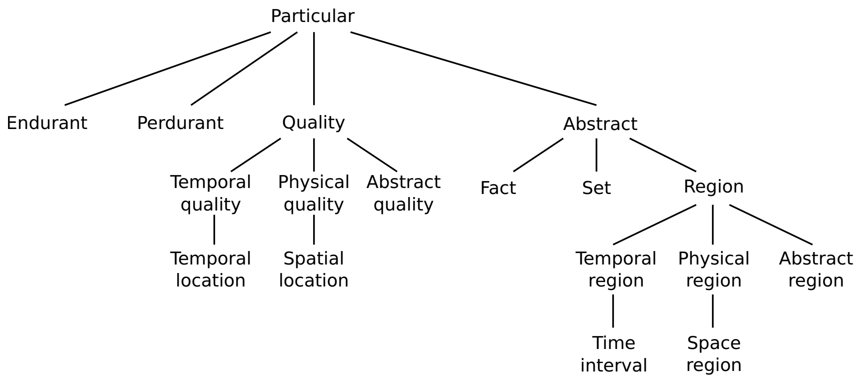

Figure 2. Part of the DOLCE subsumption hierarchy (after [10]). The categories *Endurant* and *Perdurant* correspond approximately to BFO's *Continuant* and *Occurrent*.

past, present, and future[1] have no objective reality and should not be included in any account of the world as it really is. This is another issue that I shall return to in the discussion section, in §3.3.

2.2. *Time in DOLCE*

For this section, my primary source is [10]. Unlike BFO, DOLCE includes a category of *abstract* entities as one of its top-level divisions. The class *abstract* includes a subclass *region*, which in turn subsumes subclasses *temporal region*, *physical region*, and *abstract region*—a slightly confusing terminology in that all three of these classes, not just the last, comprise abstract entities. These categories can be seen in Figure 2, which illustrates the parts of the DOLCE subsumption tree that are relevant to the current discussion.

Masolo *et al.* explain how entities are located in space and time as follows:

> In our ontology, space and time locations are considered as individual qualities like colors, weights, etc. Their corresponding qualia are called *Spatial (temporal) regions*. For example, the spatial location of a physical object belongs to the quality type *space*, and its quale is a region in the geometric space. Similarly for the temporal location of an occurrence, whose quale is a region in the temporal space. This allows an homogeneous approach that remains neutral about the properties of the geometric/temporal space adopted (for instance, one is free to adopt linear, branching, or even circular time). [10, p.18]

Hence their taxonomy includes temporal and spatial locations as subclasses of temporal and physical qualities, themselves subclasses of the high-level category *Quality*, as shown in Figure 2. Thus, for example, DOLCE treats being in such-and-such a location as on a par with having such-and-such a colour. Just as something's colour may change over time, so may its location. But DOLCE extends the same treatment to the *temporal* location of an event: it is a quality of the event, but unlike the spatial location of an object, it does not make sense to say of an event that its temporal location changes over time.

[1]Often, following McTaggart [11], the notions of past, present and future are referred to as *A-series* attributes of time, as opposed to the B-series which handles unchanging temporal relations such as 'before' and 'after'.

This seemingly fundamental disanalogy between space and time seems to be glossed over by the DOLCE taxonomy. This issue cannot arise in this form in BFO because BFO does not allow occurrents such as events to have qualities [16, §3.11.2].

The DOLCE documentation provides a few instances of some of these categories (see [10, Table 1, p.15]):

Temporal quality:	The duration of World War I The starting time of the 2000 Olympics
Temporal region:	The time axis 22 June 2002 One second

By 'the duration of World War I' must be meant here the interval over which the event took place (from 28 July 1914 to 11 November 1918), rather than the more usual use of 'duration' to mean the *length* of that interval. Arguably both are qualia, but whereas the former does indeed belong to the *temporal location* quality space, the relevant quality space for the latter should be something like *temporal duration*.

Regarding the temporal regions, we might wonder, as in BFO, what exactly is meant by the time axis—but at any rate one might say that DOLCE has the advantage of allowing this to be something abstract, whereas BFO would seem to be more committed to regarding it as a part of the physical world. And as for 'one second', this must be understood to mean any one particular second-long interval.

Granted that regions and intervals are abstract entities, we may probe further and ask what exactly, according to DOLCE, abstract entities *are*. This is what we are told:

> The main characteristic of abstract entities is that they do not have spatial nor temporal qualities, and they are not qualities themselves. The only class of abstract entities we consider in the present version of DOLCE is that of *quality regions* .... *Quality spaces* are special kinds of quality regions, being mereological sums of all the regions related to a certain quality type. [10, p.19].

The general picture here is of *physical endurants* having *physical qualities* whose (time-dependent) values (*qualia*) are *physical regions*. Physical regions are structured by the parthood relation. Similarly *perdurants* have *temporal qualities* whose qualia are *temporal regions*, the latter again being structured by the parthood relation. Thus for example a flower has a physical quality *colour*, whose qualia include all the various shades of red, blue, etc, which go to make up *colour space*. Another physical quality of the flower is its *spatial location*, the qualia of which are *spatial regions*. By analogy with colour, the spatial location of the flower must be located at a point in *region space*, that is, an abstract space whose "points" are regions of ordinary physical space. A *region* of region space would then be a set of regions of physical space.

The situation with temporal location is similar. If the value of the temporal location quality of some event is a certain temporal interval, then the relevant quality space in which this quale is located is not *time* (in the sense of the union of all temporal intervals) but a more abstract "interval space", which unlike time itself is a two-dimensional space, since we need two numbers to specify an interval (e.g., its start point and end point).

Unlike BFO, DOLCE does not seem to make much of the distinction between intervals and points, or between regions and boundaries. DOLCE quality spaces could very well be fashioned in the mould of "pointless" theories of space (e.g., [7]): in this case all we have is regions and intervals, which might perhaps be arbitrarily small but never of size zero. In this way we can "flatten" the two-level structure described above for region and interval spaces: we just say that every quale is a region of its associated quality space, there being no points in that space, only what we may call 'grains', as it were provisional minima, regions too small for us to make any internal distinctions, given our measurement and perception abilities.

2.3. Time in GFO

My main source here is Baumann *et al.* [2]. The authors state (p.181) that they intend to model "abstract phenomenal time", which consists of *chronoids* and *time boundaries*. Every chronoid has exactly two *extremal boundaries* (its first and last time points) and also infinitely many *inner time boundaries* that are extremal boundaries of proper sub-chronoids. Time boundaries are existentially dependent on chronoids. There is a close correspondence between GFO's chronoids and time boundaries and BFO's one- and zero-dimensional temporal regions: both capture the intuition that "times" may be either intervals or instants.

A distinctive feature of GFO's treatment of time, however, is its endorsement of Brentano's idea of *temporal coincidence* between time boundaries. If chronoid c_1 meets chronoid c_2 then the last point of c_1 and the first point of c_2 are treated as distinct time boundaries, namely the later of the two extremal boundaries of c_1 and the earlier of those of c_2, respectively. Although these are distinct time boundaries, they are said to *coincide*, meaning that the temporal distance between them is zero.

The authors of [2] claim a number of benefits of building this notion of temporal coincidence into their model of time. In particular, it furnishes GFO with a way of solving the notorious Dividing Instant Problem: if some proposition ϕ holds over interval c_1 and its negation $\neg\phi$ holds over c_2, where c_1 meets c_2, which of them holds at the instant at which c_1 meets c_2? For GFO, c_1 and c_2 are chronoids, so 'the instant at which c_1 meets c_2' is ambiguous: it can mean either the last extremal boundary of c_1 or the first extremal boundary of c_2. This ambiguity allows us to say that ϕ holds at the former point while $\neg\phi$ holds at the latter. Although these time boundaries are coincident, they are distinct, so no contradiction arises. This solution is ingenious, but to my mind unpersuasive; in particular, it is hard to relate it to any empirical understanding of time.

Chronoids which end together share the same (not merely coincident) last time boundary; and those which begin together share the same first time boundary. The key axioms here are:

(A7) Every time boundary is a boundary of a chronoid.
(A21) Every time boundary coincides with another one.
(A22) At most two distinct time boundaries coincide.

Thus time boundaries come in coincident pairs: each pair contains the shared last extremal boundary of a collection of chronoids which end together, and the shared first extremal boundary of a collection of chronoids which begin together. Thus there are exactly two distinct *kinds* of time boundary, which we may call *endings* and *beginnings*. We can think of endings as "looking back" into the past, and beginnings as "looking forward" into the future. The elements of such a pair together form an equivalence class under temporal coincidence: a "coincidence class". If $[x]$ is the coincidence class to which time boundary x belongs, then coincidence classes can be ordered by the relation

$[x] \prec [y] \equiv$ There is a chronoid whose first extremal boundary belongs to $[x]$ and whose last extremal boundary belongs to $[y]$.

It is shown (Proposition 1) that, with the axioms given in [2], $\prec$ is an unbounded dense linear ordering on the set of coincidence classes—which are therefore, if countable, isomorphic to $(\mathbb{Q}, <)$. In this way a "standard" model of time, in which the primitive components are instants, and there is no notion of temporal coincidence distinct from strict equality, can be reconstructed within the ontologically more complex GFO model.

3. Discussion of Key Issues

3.1. Dimensionality: Instants vs Intervals

Depending on one's starting point, the notion of a time instant can seem more or less problematic. Most of us, I suppose, are thoroughly familiar with the standard mathematical representation of time as a set of instants isomorphic, with respect to cardinality and ordering, to the real line $(\mathbb{R}, <)$. This form of representation has arisen from our practices of measuring time by means of numbers. Measurement alone, however, will not deliver to us the real numbers: no measurement practice can distinguish, for example, between rational and irrational numbers, so for practical purposes, time measurements can, and always are, given as rational numbers. The reason why, despite this, the real numbers are standardly invoked in representations of time is because only they will allow us to capture the intuition that the time line is continuous. The world seems to present us with physical continua, and it has come to seem that the most natural way to model these is by means of mathematical continua, almost always modelled on the real line with its key property of Dedekind completeness.

If we accept the idea that time instants exist and that the totality of such instants can be adequately modelled as isomorphic to the set of real numbers, then it seems to be a short step from there to the idea that time, as a whole (and therefore any part of time, such as an interval) is *made* of instants: that time instants somehow constitute the very fabric of time. But this is deeply mysterious, since it leaves open what we might call the *problem of duration*. An instant has no duration—or, we might say, its duration is zero. If time intervals are made of instants, where do their durations come from? Not from the instants, since however many zeros we add up we will never arrive at a non-zero duration. It seems that we have to regard duration as a primitive property of parts of time, not to be explained by constructing them out of instants. Thus time is first and foremost something extended; if there are unextended elements of time—instants—they must arise as a result of picking them out from time in some way.[2] The usual way is by reference to *instantaneous events*: these are the process boundaries of BFO. Both BFO and GFO agree that instants (in their terminologies, zero-dimensional temporal regions and time boundaries respectively) are ontologically dependent on intervals (one-dimensional temporal regions and chronoids), and this is fully in accord with the picture I have put forward here.

The question I want to address here is just what status we should accord to these instants: in particular, should an upper ontology be countenancing them at all? Perhaps

[2]This was Aristotle's view—cf. the clear and extremely succinct summary in [6, p.13].

the main purpose of including instants in one's ontology is so that they can serve as the times of occurrence of instantaneous events: "Zero-dimensional temporal regions are the temporal regions that process boundaries are located in" [1]—for example, *In the eclipse, the start of totality will be at 11h22m34s*. In any real-world case where an event is located at an instant, however carefully we specify the event and its time of occurrence, we can never attain infinite precision, that is, we cannot give answers to infinitely many decimal places; yet just this is what is implied by using the real numbers—or even just the rational numbers—as our model for time and other physical magnitudes. This shows that the real-number model is an idealisation—which means that it belongs primarily to the world of thought rather than to the physical world that is the target of thought.

In GFO, entities existing in space and time are classified into:

- *Continuants*, which persist through time and have lifetimes which are chronoids;
- *Processes*, which happen in time and have temporal extension; and
- *Presentials*, which are entities "exhibited" by a continuant at each time point of its lifetime; a presential (but not the continuant which exhibits it) is wholly present at the unique time boundary of its existence.

What is the relationship between a continuant and the presentials it exhibits? For Baumann *et al.* [2], continuants are "cognitive creations of the mind, constructed on the basis of presentials"; they are therefore existentially dependent on presentials. This should be set alongside the equally explicit claim that "time boundaries depend existentially on chronoids". The direction of dependency for times goes the other way from the dependency for entities which exist at those times: continuants exist on chronoids and presentials exist at time boundaries, but whereas continuants are dependent on presentials, time boundaries depend on chronoids. Although time boundaries have only a secondary sort of existence, being dependent on the chronoids whose boundaries they are, they are none the less required to supply the unique times of existence for presentials, which are primary existents on which continuants having chronoids as their lifetimes depend. The picture presented is of many presentials, each of them wholly existing at a unique time boundary, forming the basis for the construction of continuants persisting over chronoids containing those time boundaries. It is hard to see how this picture can be sustained if time boundaries do not themselves have a primary, independent status. Lurking in the background here there seems to be the notion that chronoids are somehow *made* of time boundaries, which is essentially the notion which we have already dismissed as untenable. And indeed, Baumann *et al.* explicitly repudiate such a picture:

> A process is neither the mereological sum of its boundaries, nor can it be identified with the set of its boundaries. There is no way to construct a process from process boundaries, because processes are the more fundamental kind of entity in GFO. Indeed, *and similarly to the relation of time boundaries and chronoids*, process boundaries are (specifically) existentially dependent on their processes.
>
> [2, p.184, my italics]

This would make it very awkward to claim that continuants (existing on chronoids) are *made* of presentials (existing at time boundaries); and indeed they do not claim exactly this—what they say is that continuants are 'constructed on the basis of' presentials. But it does not seem to be made clear what the difference is between making X out of Y and constructing X on the basis of Y.

In a telling remark, Baumann *et al.* characterise the relation between continuants and presentials as follows: “We say that a continuant c *exhibits* a presential p, if p exists at a time boundary t and p corresponds to conceiving of c at t (or as viewing p as a “snapshot” of c at t)”. In itself, this talk of snapshots is perhaps harmless, but it needs handling carefully. A persistent strand in the history of thinking about motion and change has it that the history of the world is nothing but the aggregate of all its instantaneous snapshots. This is the essence of the so-called ‘at-at’ theory of change, according to which all that it means to say that change occurs is that *at* some time t a certain state s obtains, and *at* some later time t' a different state s', incompatible with s, obtains. Each instantaneous snapshot is static, since there is no room for change to occur in an instant, and this has led to the theory being called the *static theory of change*. Against this, the objection has frequently been voiced that the order of explanation here should be reversed: the different states obtain *because* the state of change exists.

In any case, the idea of an instantaneous snapshot is itself problematic. This use of the term ‘snapshot’ is of course a metaphor from photography, but in photography a snapshot always records the world over an interval (the exposure time). The settings on your camera do not include the possibility of zero exposure! But because one can take a snapshot with a sufficiently short exposure that no motion in the scene being photographed will show up, given the spatial granularity of the image, we can fool ourselves into thinking that the dynamic world is somehow built up out of a series of static, instantaneous snapshots. This may be tempting with regard to vision, but if we turn to sound it makes no sense at all: one can look at a still from a video file, but we could not listen to a “still” from an audio file—there would be nothing to hear! It is obvious that sounds are essentially dynamic and require intervals for their realisation; on closer analysis it should be obvious that this is true of visible things as well.

Perhaps therefore we can agree with philosophers such as Whitehead [19], James [9], and Bergson [3] that the instant is a mathematical idealisation not to be found in reality itself. Its utility comes in particular from the fact that by using a mathematical continuum (the real line) as our model for physical continua, we can readily access the mathematical power of the calculus to enable us to explain and predict physical phenomena in the world—typically by means of differential equations. Given this, we can ask whether and in what form time instants should be countenanced by an upper ontology.

This will depend on the purpose of the ontology. In their introductory chapter, Arp *et al.* [1] state that

> we assume that it is uncontroversial that ontologies should be understood as a kind of representational artifact, and that the entities represented are entities in reality—such as cells, molecules, organism, planets, and so forth.

They did not say “— such as zero-dimensional temporal regions, one-dimensional temporal regions, and so forth”; but since these are also terms in the ontology representing entities, use of the BFO ontology in which they occur implies a commitment to regarding these too as “entities in reality”. If it can be conclusively shown that, for example, there is nothing in reality answering to the notion of a zero-dimensional temporal region, then presumably BFO would have to be modified by expunging these entities from it.

DOLCE would appear to be on safer ground here, in that it is not founded on the premise that the terms of the ontology should represent actually existing realities; rather, it is designed as a vehicle enabling the formalisation of *conceptualisations* of reality. A conceptualisation of reality can readily admit entities which are recognised as being in various ways idealisations or even convenient fictions, and the underlying philosophy of DOLCE is quite prepared to countenance this. It is thoroughly in keeping with this that unlike BFO, DOLCE includes a class of abstract entities, and that spatial, temporal, and spatio-temporal regions are located within that class.

Abstraction also features in GFO's approach: "The basic theory of phenomenal time in GFO is abstracted from real-world entities and is inspired by ideas of Franz Brentano" [2, p.181]. By 'phenomenal time' is presumably meant time as manifested in phenomena, that is the manner in which reality presents itself to us as experiencing beings. But there are (at least) two different ways of interpreting this. Does it mean that time as we experience it is already an abstraction from the real-world, which is then faithfully mirrored in the ontology, complete with instantaneous times and process boundaries? Or does it rather mean that ontology itself abstracts these notions from our time experience, and that therefore instantaneous time boundaries and the like are not literally to be found in phenomenal time but only in our cognitive reconstitution of this in the form of an ontology? Let us note here the obvious relevance to these questions of the psychological notion of the "specious present" [8], an interval of subjective time that is experienced as present as a whole; I will return to this at the end of the next section.

3.2. Frame-dependence and Relativity

Can time intervals be defined without reference to space? In STR, the spatiotemporal separation between two points is said to be *timelike* if some slower-than-light causal influence can run from one to the other—this would include, for example, the case of some one *object* being present at both points. A *timeline* can then be defined as any one-dimensional space-time region any two points of which have timelike separation. From within a given reference frame, a time-line may cut across space in the sense that different points along it may be at different spatial positions in that frame—for example, the path traced in space-time by the tip of my nose is a timeline of this kind. We can now say that relative to a frame, a timeline is a *local time interval* if all its points are at the same spatial position; relative to other frames, these points will still form a timeline but will not be a local time interval because the points have differing space-coordinates in those frames. The local time intervals thus defined correspond to the second of the two models of time intervals as parts of space-time given in §2.1.1.

Within a frame we can define a *global time interval* as a minimal spatiotemporal region which includes within it all the *temporally coincident* local time intervals, that is, local time intervals whose beginnings and endings are respectively simultaneous with the beginning and ending of a given such interval. This corresponds to the 'slice' conception described in §2.1.1. Alternatively, we can consider the abstract *set* of all the local time intervals temporally coincident with a given interval. This is a simultaneity class of local time intervals, and if we define a global time interval to be any such simultaneity class, this corresponds to the third conception introduced in §2.1.1. Under either definition, a global time interval is indeed *global* in the sense that it covers all of space, but it is not *universal* in the sense of being frame-independent. And STR tells us that there are no absolute global time-intervals in this sense.

It is moreover arguable that even frame-dependent global time-intervals are suspect. At any rate, they are an artefact of a certain way of defining simultaneity in a frame. Einstein asked what it means to say that two spatially separated events are simultaneous; his answer has been accepted as the standard way of defining simultaneity in STR. Suppose I send a radio message to my friend on Mars: call this event E_1. Suppose further that my friend, on receiving my message, immediately replies: call this event E_2. Finally, let E_3 be the event of my receiving my friend's reply here on earth. I know when E_1 and E_3 happened, since in both cases I can consult my clock here on earth as they occur; but when, in terms of *my* clock time, did E_2 happen on Mars? According to Einstein, since the speed of light is the same in all inertial frames, it must have happened exactly halfway between the times of E_1 and E_3. Suppose that at that moment I sneezed (event E_2'): then, the story goes, my sneeze E_2' was simultaneous with E_2, my friend's replying to my message. The problem here is that this judgment of simultaneity depends on my coordinate-frame, with its time-axis defined by my clock readings. Another observer, in motion relative to me, would not judge E_2 and E_2' to be simultaneous. Hence the well-known absence of an absolute, frame-independent definition of simultaneity in STR.

Because of this, it has been argued by Stein [18], among others, that it never makes sense to speak of simultaneity between spatially separated events: the only simultaneous events are ones which happen together at the very same space-time point. As Stein puts it, "an event's present is constituted by itself alone". If this is correct, then the idea of "the world at a time" is unsustainable, even if relativised to a particular observer. In STR, the Newtonian model of an absolute time-line that is independent of space and can therefore support a global simultaneity relation gives way to an integrated Minkowskian space-time in which the only absolutes are the spatio-temporal separations between points. The separation between two points in space-time is either *timelike* (meaning that a slower-than-light causal influence can pass from the earlier to later), *lightlike* (meaning that a light-ray can pass from the earlier to the later), or *spacelike* (meaning that there can be no causal influence either way between the points). Any proposed simultaneity relation apart from identity would have to hold only between points with spacelike separation, which are causally unconnected. This clashes radically with our intuition that "we are all in it together", i.e., that we are as it were moving through time as a cohort of interacting contemporaries. The reason that we can have this intuition is that the present moment of our experience is not, in fact, a durationless instant but rather a somewhat "smeared out" interval, the *specious present* of William James [8]. As pointed out by [4,5], our pace of existence is sufficiently slow relative to the speed of light that within the duration of a single specious present there is time for multiple mutual interactions with others in our vicinity (but not, for example, on the moon or further afield).

In the light of this, an ontology which postulates a time-line that is independent of space is thereby embracing an idealisation: the limit idealisation by which, for sufficiently slowly moving objects within a sufficiently delimited region of space-time, the Minkowskian geometry is effectively indistinguishable from the Newtonian. And of course, for the reason just mentioned, this idealisation is quite serviceable in practice; but it does mean that the relations between the ontology and the reality it purports to be modelling are more complex than is sometimes acknowledged. In particular, the BFO claim to be modelling reality directly would need to be replaced be a more nuanced account in which the idealised nature of the model is recognised. Perhaps this would lead to something like the notion of abstract phenomenal time which is what GFO claims to be modelling. The "phenomenal" aspect here pays due regard to the role of *experience* in forming our temporal concepts; and of course, DOLCE's explicit commitment to conceptual modelling, as opposed to any notion of capturing “reality itself”, could in principle accommodate many different conceptions of time and temporal phenomena.

Russell [14] drew a sharp distinction between what he called *physical time* and *mental time*. The former is whatever is mandated by our best physical theories, where time is integrated with space in a static four-dimensional structure comprising non-denumerably many unextended space-time points; the latter is time as we experience it, a dynamic, “flowing” thing in which a stream of events successively tumble through the narrow “window” of our present awareness. Many of the difficulties that we have in forming an adequate conception of time for the purposes of understanding and reasoning about our world stem from the problems of reconciling these two very different pictures. The three upper ontologies considered in this paper approach this in somewhat different ways, but it is broadly true that, in common with much of our everyday thinking, they none of them really succeed in disentangling the common language and concepts of a subjective view of time from the more objective features of a scientific account.

3.3. Indexicality: Past, Present, and Future

Should an ontology of time include reference to the present moment? As noted above, if ‘right now’ is seriously proposed, as in [1], as an example of a zero-dimensional temporal region, then the implicit answer is that it should. But this brings manifold problems in its wake. In particular: (1) Is the present objective? (2) If it is, what is its nature?

For modelling time as we experience it, the distinction between past, present and future seems to be fundamental. But if, in time as we experience it, the present has some, albeit brief, duration, it is perhaps not best modelled as an instant. As evidence of this, note that as a matter of common experience, we can sense motion and change directly. A moving object *looks* and *feels* different from a stationary one. If our experience were confined to an instant at a time, then our knowledge of motion and change would have to be deduced from a comparison of present experience with memory of the recent past—this is how we detect slow changes such as the motion of the hour hand of a clock, as opposed to the motion of the second hand which we can see directly. Another piece of evidence comes from the sense of hearing: sounds are by nature extended in time; a world experienced at an instant, if such were conceivable, would perforce be silent.

For the purposes of mathematical analysis, however, it is convenient to treat time as “made of” instants, corresponding to real or rational numbers. This, no doubt, lies at the root of Russell’s distinction between physical and mental time, alluded to above. On this view, if there is an objective present, as distinct from our individual specious presents, perhaps this would indeed be an instant. But it is often claimed that the advent of Relativity disposed once and for all of the idea of an objective present, since that requires a global simultaneity relation, which STR gives no support for. Some authors, such as Rakić [15], have suggested that even though such a relation is not derivable from the causal structure of Minkowski space-time, it is not incompatible with it, leaving room for the existence of such a relation to support the idea of an objective present: past, present and future, she says, are not temporal notions but ontological ones, having to do with the distinction between “realised” and “unrealised” events. It has even been suggested that such a notion is *required* by Quantum Mechanics to account for the collapse of the wave-function. If this is right, then it would be reasonable for a formal ontology that claims to provide an account of objective reality to include pastness, presentness, and futurity as attributes of times, though it is hard to see how to integrate this into the overall temporal framework, and as far as I am aware none of the currently existing formal upper ontologies has attempted to do this. Until such time as this is done, it would be best to steer clear of problematic entities such as ‘right now’ as examples of instants.

4. Concluding Remarks

Although it is natural that an upper ontology should provide an account of time—not least in order to support reasoning about processes and events—it is hard to do this otherwise than by largely endorsing something like a “commonsense” view of time that, while taking on board some issues relating to scientific as well as everyday practices of time-measurement, does not really come to grips either with long-standing philosophical debates concerning the nature of time or with the radical departures from the commonsense view suggested by physical theories. While this is undoubtedly workable for the purposes of most current applications of ontology, it does raise questions about the ultimate status to be accorded to such an ontology—in particular in relation to the much-debated issue of realism vs conceptualism (see for example the exchange of views in [12,17,13]). These questions, for the most part, remain unanswered.

References

[1] Robert Arp, Barry Smith, and Andrew D. Spear. *Building Ontologies with Basic Formal Ontology*. The MIT Press, Cambridge MA and London UK, 2015.

[2] R. Baumann, F. Loebe, and H. Herre. Axiomatic theories of the ontology of time in GFO. *Applied Ontology*, 9:171–215, 2014.

[3] Henri Bergson. *The Creative Mind*. Philosophical Library, New York, 1946.

[4] Jeremy Butterfield. Sensing the present. *Mind*, 93:161–176, 1984.

[5] Craig Callender. The common now. *Philosophical Issues*, 18:339–361, 2008.

[6] Ursula Coope. *Time for Aristotle: Physics IV. 10-14*. Oxford University Press, 2005.

[7] Giangiacomo Gerla. Pointless geometries. In F. Buekenhout, editor, *Handbook of Incidence Geometry*, pages 1015–31. Elsevier Science B.V., 1995.

[8] William James. *The Principles of Psychology*. Henry Holt and Co., New York, 1890.

[9] William James. *A Pluralistic Universe: Hibbert Lectures at Manchester College on the Present Situation in Philosophy*. Longmans, Green, and Co., New York, 1909.

[10] Claudio Masolo, Stefano Borgo, Aldo Gangemi, Nicola Guarino, and Alessandro Oltramari. WonderWeb Deliverable D18: Ontology Library (final). Technical report, LOA-ISTC, CNR, 2003.

[11] J. M. E. McTaggart. The unreality of time. *Mind*, 17:457–74, 1908.

[12] Gary H. Merrill. Ontological realism: Methodology or misdirection? *Applied Ontology*, 5(2):79–108, 2010.

[13] Gary H. Merrill. Realism and reference ontologies: Considerations, reflections and problems. *Applied Ontology*, 5(3-4):189–221, 2010.

[14] Bertrand Russell. On the experience of time. *The Monist*, 25(2):212–233, 1915.

[15] Nataša Rakić. Past, present, future, and special relativity. *British Journal for the Philosophy of Science*, pages 257–280, 1997.

[16] Barry Smith. *Basic Formal Ontology 2.0: Specification and User’s Guide*. Accessible from https://github.com/BFO-ontology/BFO, 2015.

[17] Barry Smith and Werner Ceusters. Ontological realism: A methodology for coordinated evolution of scientific ontologies. *Applied Ontology*, 5(3-4):139–188, 2010.

[18] Howard Stein. On Einstein-Minkowski space-time. *Journal of Philosophy*, 65(1):5–23, 1968.

[19] Alfred North Whitehead. *The Concept of Nature*. Cambridge University Press, Cambridge, UK, 1920.