

The Ontology of Time and Process

Third Interdisciplinary School on Applied Ontology, Bozen-Bolzano, June 27–July 1, 2016

Antony Galton, University of Exeter, UK, apgalton@ex.ac.uk

1 Continuants and Occurrents

Ontology is the science of *being*. In the first instance this means *what exists*. But the world consists of more than just what exists: it includes what *happens* as well. Nowadays almost all ontologies include things that happen as well as things that exist: events, processes, and activities, as well as people, buildings, and furniture.

What is the *difference* between existing and happening? And how are they related to each other? Here are some possible answers:

1. There is no difference. This view leads to fully-fledged *four-dimensionalism* [15, 23, 26]¹. On this view we cannot meaningfully distinguish between objects (as we usually understand them) and processes: instead, the world comprises four-dimensional entities which may present both object-like and process-like aspects depending on whether they are viewed *across* or *along* the time axis.
2. Things that happen are *properties* of things that exist. For example:

Mary

- ... is 160 cm tall
- ... has red hair
- ... wears glasses
- ... cycles to work every morning
- ... is wearing a blue dress
- ... is painting a picture
- ... bought a new bicycle in June

It is implicit here that properties are in some sense “second-class citizens”—they are not *entities*, and to speak of them as if they were (e.g., “Mary’s hair colour” or “Mary’s purchasing a bicycle”) is merely a *façon de parler*.

3. Things that happen are entities, but they are ontologically dependent on things that exist (objects). If Mary does not exist, then her bicycle-purchasing cannot happen (the bicycle must also exist for this as well). This is the view adopted in Basic Formal Ontology (BFO)[2] which we will examine in more detail below.
4. Things that exist are ontologically dependent on things that happen. This is (one form of) *processism* [22]. On this view, Mary only exists by virtue of many things that happen, e.g., her internal bodily processes and her interactions with the outside world, and ultimately what primarily exists are things that happen.
5. Things that exist and things that happen are mutually dependent. This is the view expressed in [12]: an object is dependent on its *internal* processes, and its *external* processes (in which it *participates*) are dependent on it. The internal processes of an object are themselves external processes of *parts* of the object, so we have a chain of dependencies.²

¹Within four-dimensionalism there are variants such as the “worm” view according to which ordinary objects are 4D “worms” and the “stage” view according to which an object has a stage at each moment of its existence, the separate stages being related as counterparts rather than identical. Further exploration of these variants would take us too far afield here.

²This mutual dependence is also hinted at in [17]: “[I]f we admit that every object has a life, it is hard to exclude a mutual specific constant dependence between the two”.

The creators of BFO made use of a distinction between two sub-ontologies called SNAP and SPAN [14]. (Note: The terminology SNAP vs SPAN is not referred to in the latest incarnation of BFO: but what it stands for is still very much alive.)

- SNAP is the ontology of what exists (at a moment of time — a SNAPshot).
- SPAN is the ontology of what happens (SPANning a period of time).

The inhabitants of the SNAP and SPAN components of the ontology are called *continuants* and *occurents* respectively (terms from philosophy). It is usual to state that

- A continuant exists wholly at each moment of its existence. It *endures* through time, possibly gaining or losing parts, and changing with respect to some of its properties. It may have spatial parts but not temporal parts.
- An occurrent unfolds over time (it “*perdures*”), and has temporal (and possibly also spatial) parts.

What does all this mean? Consider a *house* as an example of a continuant. The history of the house:

1. At t_1 , the foundations are laid.
2. By t_2 , the walls and roof are all in place but the windows, doors, etc, have not yet been installed.
3. At t_3 the house is at last ready to live in.
4. At t_4 , an extension is built on one side of the house.
5. At t_5 , the house is deserted and begins to decay: roof tiles fall off, the windows are broken, the brickwork starts crumbling.
6. At t_6 all that is left standing is a ruined shell, with incomplete walls and no roof.
7. At t_7 all that remains is a pile of rubble on the ground.

What does it mean to say that the house “exists wholly at each moment of its existence”? At t_6 it is clearly an *incomplete* house: so how can we say it exists as a whole at that time? Well, what exists wholly at t_6 is precisely what exists at t_6 ; and that is all of the house that there is at t_6 . It can only be described as incomplete with reference to what it could be (and was) at some other time: but at t_6 it is a *complete incomplete house* — that is, as a house, it is incomplete, but as an incomplete house, it is complete. Its completeness consists in the fact that at each time a certain collection of parts together make up something to which a certain *identity* is ascribed; and we track that identity over time by tracking the particular parts that belong to it at each time. Parts may come and go, but it is by being linked to this particular identity that we can say that this collection of parts at this time belong to the same continuant as that collection at that time.



Figure 1: A person is a SNAP entity (continuant) existing in a series of snapshots

To illustrate the difference between SNAP and SPAN, consider first the sequence of snapshots in Figure 1. Each image shows a continuant entity existing in a particular dated SNAP ontology; and it is the *same* entity in each case: me (AG), as it happens. That entity has various parts, and they are spatial parts: for example, my nose, which appears in each of the snapshots. The whole of my nose exists in each of the SNAP ontologies (though of course the photograph only shows the outer surface of the nose — a photographic snapshot is not the same as a SNAP entity).

The SPAN entity corresponding to this sequence of snapshots is my life (or at least that part of it which has elapsed so far). This is illustrated in Figure 2. This entity — my life — is extended




Figure 2: A person's life is a SPAN entity (occurrent)

through time, and has temporal parts, some of which are shown in the illustration. Each of these parts is an occurrent — some are of short duration (my birth, my wedding), others take years to elapse (my schooldays, my academic career). Each is a SPAN entity in its own right. I myself *participate* in each of these occurrents, and each of them is ontologically dependent on me: necessarily, they could not have happened if I did not exist.

On the four-dimensionalist view, there is no such thing as a continuant. Everything is an occurrent. The sequential snapshots are just temporal parts of my life; unlike on the SNAP view, AG in 1972 is not the same entity as AG in 1975. Rather there is just one four-dimensional entity, AG, and the 1972 and 1975 snapshots show distinct parts of that. There are differences between the parts — the 1972 part was bearded, the 1975 part was clean-shaven, but this is not change as we normally understand it, i.e., it is not the case that there is some one entity which changed from being bearded in 1972 to being beardless in 1975; rather, there is a single four-dimensional entity of which the part which overlaps 1972 is bearded and the part which overlaps 1975 is not. The four-dimensionalist view is probably coherent (and has been argued for in considerable detail, e.g., by Sider, and others), but you have to make many adjustments to ordinary language, or at least to how it is understood, to talk coherently about it. This is not always appreciated, and I have the impression that many people who are drawn to four-dimensionalism are not fully aware of just what a wrenching disruption to our everyday world-view it entails.

Is the standard 3+1-dimensional view coherent? (Some four-dimensionalists think not.) It is not without difficulties. How can one and the same man be both bearded and not bearded? Well, he was bearded in 1972 and beardless in 1975. How should we construe this?

1. We could read it as *AG-in-1972 is bearded, but AG-in-1975 is not bearded*. This looks like the “stage” version of the 4D view: the paradox is resolved because it is not the same entity that is bearded and not bearded.
2. Alternatively, we could read it *AG is bearded-in-1972, but AG is not bearded-in-1975*. Now the paradox is resolved because it is not the same property that is both asserted and denied of AG: beardedness-in-1972 is a different property from beardedness-in-1975.

But neither of these two resolutions seems to capture exactly what we mean. In a predicate-and-argument notation we might write them as

1. $Bearded(AG_{1972}) \wedge \neg Bearded(AG_{1975})$
2. $Bearded_{1972}(AG) \wedge \neg Bearded_{1975}(AG)$

In some logical analyses one writes instead

3. $Bearded(AG, 1972) \wedge \neg Bearded(AG, 1975)$

Thus, properties (such as beardedness) are ascribed, not to continuants *tout court*, but to continuant-time pairs. But this too does not seem to capture exactly the 3+1-dimensional view. The idea of change is still missing, since again it is different things (different continuant-time pairs) of which beardedness is both asserted and denied.

Another possibility is to *reify* the ascription of a property to a continuant and write something like

4. $True(bearded(AG), 1972) \wedge \neg True(bearded(AG), 1975)$

where the focus is not the change of the continuant AG with respect to the property of beardedness but the change of the property-ascription beardedness-of-AG with respect to property of truth. But it is still the case that the predicate here, *True*, is being asserted of and withheld from different things, different property-ascription/time pairs, with *difference* offered as a surrogate for genuine change.

2 Processes and Events

In the literature one finds much confusion over the meanings of the terms ‘process’ and ‘event’. As Worboys [30] puts it, ‘One person’s process is another’s event, and vice versa’. Here are some examples (taken from [8]) to bear this out:

- Mourelatos [19]: *situation* is divided into *state* and *occurrence*, and the latter is divided into *process* and *event*.
- Allen [1]: *state* (called by him ‘property’), *process* and *event* are categories on an equal footing.
- Pustejovsky [21]: *event* is divided into *state*, *process* and *transition*.
- Moens and Steedman [18]: *event* is subdivided into various categories, including *process*; *state* is kept separate.
- Sowa [25]: *process* is subdivided into *continuous process* and *discrete process*; the former is divided into *initiation*, *continuation* and *cessation*, and the latter into *event* and *state*.

There is ample evidence that, as these terms are ordinarily used in everyday language (i.e., when we’re not doing ontology), ‘process’ and ‘event’ mean very different things. An easy way to collect such evidence is to Google the phrase “Not an event but a process” and see what a wealth of different results you find. When I did this recently, I found that all of the following have been claimed to be processes rather than events:

- | | | |
|------------------------|-----------------------------|----------------|
| 1. Change | 7. Earning investor loyalty | 13. Devolution |
| 2. Learning | 8. Network Node Shutdown | 14. Revolution |
| 3. Marketing | 9. Retirement | 15. Genocide |
| 4. Capacity building | 10. Brain Injury | 16. Success |
| 5. Sales training | 11. Recovery from addiction | 17. Creativity |
| 6. Performance Reviews | 12. ‘Brexit’ | 18. Peace |

The common understanding here is that

- An event is a “one-off” occurrence with a definite beginning and end, it has a completion, typically lasts a short time (relative to the temporal scale of the context in which it is described), and can even be instantaneous.
- A process is open-ended, continues indefinitely, need never reach a state of completion, may be extended over a long period, possibly involving a variety of different activities.

From the web-sites listed above it may seem as though *anything that happens* could be claimed to be “a process rather than an event”. Indeed this, or similar phrases, seems to have become something of a cliché, or mantra, that people trot out with an air of apparent wisdom, as if they were providing a startling new insight that has passed lesser mortals by. So what is going on here?

Let’s consider an example: Making an apple pie. Can we say that

making an apple pie is not an event, but a process?

- The Event case. Making an apple pie has a definite beginning, it has a completion (when the pie is ready), it doesn’t take more than an hour or so.
- The Process case. Making an apple pie involves a succession of different activities: cutting up the apples, preparing the pastry, assembling the pie, heating the oven, cooking the pie; it is by no means instantaneous.

This may not be the most convincing example, but it does serve to highlight a couple of important points:

1. What happens or goes on can partake of both event-like and process-like characteristics;
2. Or, to put this a little differently, what happens can be *described* using either “event language” or “process language”, representing different points of view.

In [8] I called these points of view the *historical* and *experiential* aspects of what happen (the terms are borrowed from [16]).

In language, these different points of view show up as the *aspects* of the verb: *perfective* and *imperfective*. Different languages have different ways of conveying aspectual distinctions; in English it is typically by means of the simple *vs* continuous tenses:

- What did you do yesterday? I made an apple pie.
- What were you doing when I called? I was making an apple pie.

Here the first sentence encapsulates the whole pie-making as a completed integral event, whereas the second sentence presents it as an on-going process (which may or may not have reached completion). One way to capture the difference between events and processes on this picture is to note that, at a sufficiently coarse granularity, an event can be conceptualised as a punctual occurrence (point), and at a sufficiently fine granularity a process can be conceptualised as a state. This shows how one and the same portion of reality can be viewed either as a state or as an event.

These observations have implications for the ontology of occurrents. If the target of your ontology is to characterise the nature of reality as it is in itself, independently of the manner in which it is viewed or presented, then you should regard the difference between processes and events as outside the scope of the ontology. But if the target of your ontology is to characterise elements of reality from different points of view, so that one and the same entity may be classified in different ways in different contexts, then the difference between processes and events should be taken seriously as something to be enshrined in the ontology itself. In the next few sections we shall consider different approaches to handling the event/process dichotomy, and will see examples of both these approaches.

3 Processes and Events in BFO and DOLCE

3.1 BFO

BFO (Basic Formal Ontology) [24] defines what it means for an entity p to be a *process* as

p is an occurrent that has temporal proper parts and for some time t , p s-depends-on some material entity at t .

Here ‘s-depends-on’ stands for ‘specifically depends on’, and is defined so that an entity x s-depends-on y just if x cannot exist unless y exists too. The material entities which a process p s-depends-on are the *participants* of p .

The BFO documentation gives as examples of processes “the life of an organism, a process of sleeping, a process of cell-division, a beating of the heart, a process of meiosis, the course of a disease, the flight of a bird, your process of aging”.

Note carefully that these processes are *particulars*. Each of them occupies a particular spatio-temporal region. The life of an organism, for example, occupies a long thin spatio-temporal “worm”; for each time t within the temporal extent of that worm, the (three-dimensional) cross-section of the worm at t is the spatial region occupied by the organism at t . A process of sleeping is a proper part of the life of some organism: it is a maximal section of the “worm” for which, at each moment during its temporal extent, the organism is asleep. In the BFO list, it is not clear whether a beating of the heart refers to a single heart-beat, i.e., a single cycle of contraction and relaxation of the heart muscle, or to the extended heartbeat process that lasts for the entire lifetime of the organism — both of these are good examples of BFO processes, the latter being made up of a large number of individuals of the former type.

In addition to processes, BFO includes a type *process-boundary*, defined as a temporal part of a process which has no proper temporal parts itself. Thus a process-boundary is not itself a process. Process-boundaries include the beginnings and endings of processes as well as any minimal temporal cross-section within the interior of the process. The diagram in figure 3 shows a process (as a 4D




Figure 3: A BFO process and one of its temporal parts

“sausage”) together with one of its temporal proper parts, the intersection of the process with a certain period of time, delimited by process boundaries at either end.

It should be noted that BFO does not include a separate category of *event*: for BFO, there is essentially no difference between events and processes.

When someone makes an apple pie, this is a BFO process. It has various temporal proper parts that we can describe, for example the cutting up of the apples, which in turn has temporal proper parts which are the cuttings up of individual apples, and the individual knife-strokes; preparing the pastry, with its temporal proper parts such as pouring out the flour, adding butter and water, mixing them, rolling out the dough, lining the pie-tin with it, and so on.

An important feature of BFO is that *occurents do not have qualities*. There is no such entity as, for example, *the speed of this movement*, *the direction of this movement*, or *the loudness of this music*. This means that *occurents cannot change*, since change in an entity is always with respect some quality or other of that entity, which takes different values at different times.

Instead of saying that, for example, the movement is getting faster, the BFO specification document instructs us to say something like:

Motion m is an instance of *motion-at-speed- v_1* at t_1 , motion m is an instance of *motion-at-speed- v_2* at t_2 , where $t_1 < t_2$ and $v_1 < v_2$.

But it is not made clear, I think, what it means to say that one motion instance can be an instance of different motion universals at different times — which would seem to imply that a motion instance *can* change, despite being an occurrent.³ I think the correct way to refer to the changing speed of a motion in BFO must be as follows:

Motion m has temporal proper parts m_1 and m_2 , existing at t_1 and t_2 respectively, such that m_1 is an instance of *motion-at-speed- v_1* and m_2 is an instance of *motion-at-speed- v_2* , where $t_1 < t_2$ and $v_1 < v_2$.

Note that for m_1 and m_2 to be processes, they must themselves have temporal proper parts — hence they cannot exactly occupy instants t_1 and t_2 but must occupy intervals containing those instants. If we want to talk about the instantaneous speed at t_1 , we have to understand this as a statement about the limit of average speeds of a sequence of temporal proper parts of the process occupying ever shorter intervals converging on t , as in the standard definition of the derivative in the differential calculus, where “average speed over an interval” is logically prior to “speed at an instant”. This is in keeping with a “scientific” analysis of time and motion and hence in the spirit of BFO’s claim to be an ontology to support scientific endeavours. It is not so friendly from the point of view of more everyday human purposes.

³Compare the case of continuants: at t_1 , John is an instance of the universal *child*, and at t_2 he is an instance of the universal *adult*. This is fine, because for BFO continuants *can* change — in particular with respect to which universals they instantiate.

3.2 DOLCE

DOLCE (Descriptive Ontology for Linguistic and Cognitive Engineering) [17] embodies a very different philosophy from that of BFO. Whereas the latter purports to provide a set of tools for directly describing the mind-independent reality that is the focus of scientific investigation, DOLCE has the aim of systematising the categories employed in human conceptualisations of the world. A BFO category stands or falls according to whether it succeeds in picking out a class of real-world entities united by the possession of some shared set of objective properties, and it is the task of the scientist to determine whether or not this is the case through empirical investigation; a DOLCE category, by contrast, is intended to represent some concept under which humans might group entities in accordance with some cognitively determined principle of classification which may or may not correspond to any objectively verifiable feature of the underlying reality.

A consequence of this difference in orientation is that DOLCE’s classification of occurrents looks very different from that of BFO. Where the two systems agree is in the top-level division into continuants and occurrents — although in the terminology used by DOLCE these are called *endurants* and *perdurants* respectively. Perdurants are also called *occurrences*; DOLCE reserves the term *occurent* to apply to universals whose instances are occurrences.

Whereas in BFO every entity is either a continuant or an occurrent, for DOLCE the division into endurants and perdurants is not exhaustive, since it regards some entities, notably *qualities*, as neither endurants nor perdurants. BFO, on the other hand, treats qualities as continuants that are specifically dependent on other continuants (e.g., a person’s temperature).

The DOLCE documentation makes a number of useful observations concerning the endurant/perdurant distinction, notably:

- The part-whole relation for endurants must be time-indexed (since endurants can gain or lose parts, we have to specify *when* one endurant is part of another), whereas the part-whole relation for perdurants is not (one perdurant is part of another *tout court*; it does not make sense to speak of a perdurant gaining or losing parts). This distinction is common to BFO and DOLCE (in BFO the two relations are ‘*x* *continuant-part-of* *y* at *t* and *x* *occurent-part-of* *y*’).
- Endurants are primarily located in space, and inherit their temporal locations from the perdurants in which they participate; perdurants are primarily located in time, and inherit their spatial locations from the endurants which participate in them.

Where DOLCE differs strikingly from BFO is in its classification of perdurants. Whereas BFO uses ‘process’ as a blanket term to cover all occurrents that are extended in time (i.e., have temporal proper parts), DOLCE reserves this term for a more restricted subclass of perdurants. DOLCE’s classification of perdurants is shown in Figure 4.




Figure 4: The classification of perdurants in DOLCE

The terms ‘cumulative’ and ‘homeomeric’ used to provide differentiae in this classification are explained as follows:

- An occurrence-type is cumulative if “it holds of the mereological sum of two of its instances”.
- “An occurrence is homeomeric if and only if all its temporal parts are described *by the very expression* used for the whole occurrence”.

It is important to note that these are not put forward as *intrinsic* properties of occurrences, but rather as properties of occurrences *under a description*. This is explicit in the use of the phrase “by the very

expression” in the definition of homeomerity and by the use of the term “occurrence-type” rather than “occurrence” in the definition of cumulatity. Thus we have:

- Occurrence type T is cumulative so long as, for any two instances x and y of type T , the mereological sum $x + y$ is also an instance of type T .
- Occurrence type T is homeomeric so long as, for each individual instance x of type T , every temporal proper part of x is also an instance of type T .

It follows that the terms “stative”, “eventive”, “state”, “process”, “achievement” and “accomplishment” used in the DOLCE classification of perdurants apply not to occurrences *per se* but to occurrences *under a description* (i.e., as instantiating a certain type).

For example, consider a situation in which an aircraft flies from London to New York. We can describe this situation using either of the clauses

- (a) *The plane flies.*
- (b) *The plane flies from London to New York.*

We shall attempt to determine the categories of perdurant to which DOLCE would assign each of these descriptions.

To test cumulatity, suppose F_1 and F_2 are two distinct occurrences, to each of which both the descriptions (a) and (b) apply. Then (a), but not (b), applies to the mereological sum $F_1 + F_2$. Hence in the DOLCE classification (a) describes a stative occurrence and (b) describes an eventive one.

Taking the case of (a), now, we consider whether or not it is homeomeric. Let F be any occurrence to which description (a) applies, and let F' be a temporal proper part of F . Does (a) apply to F' ? Only if F' is temporally extended! It takes time even to do the smallest amount of flying, and therefore any occurrence that can be correctly described using (a) must be temporally extended. Hence (a) does not apply to temporally atomic parts of F , and therefore is not homeomeric. Thus (a) describes a *process*.⁴

Turning now to (b), we consider whether or not it is atomic. Very clearly, it is not. A flight from London to New York cannot be accomplished in an instant! Hence (b) is not atomic, and therefore describes an *accomplishment*.

A third way of describing our occurrence F uses the sentence

- (c) *The plane is flying.*

Like (a), this is cumulative, but unlike (a) it seems clear that (c) is also homeomeric. A plane can be said to *be* flying even at a instant — “At exactly midnight, the plane was flying”. Hence, in the DOLCE classification, (c) describes a *state*.

DOLCE’s fourth category of perdurants is the achievements. Clearly there is no description under which our occurrence F comes out as an achievement, since F is temporally extended, and therefore not atomic. However, we can consider a closely related description:

- (d) *The plane lands.*

Like (b), this is clearly not cumulative. But at least under one interpretation (e.g., the interpretation by which this sentence refers to the event of first touching down) it can be understood as atomic: it applies to an occurrence which is not extended in time. As such, (d) describes an *achievement*.

A key point to note here is that all this classificatory apparatus is missing from BFO. This is because BFO is concerned with things in themselves rather than things under a description, and as we have seen, the subcategories of *occurrence* provided by DOLCE do not apply to occurrences *tout court* but only to occurrences under some description. This is explicitly acknowledged within DOLCE, as witness the comment attached to the classes “achievement” and “accomplishment” in one of the machine-readable formalisations of DOLCE⁵:

⁴Note that if the example had involved, say, walking rather than flying, we would reach the same conclusion, but in this case it might also be argued—and frequently has been—that homeomerity already fails for temporally extended proper parts so long as their duration is shorter than the intrinsic temporal granularity of the walking process.

⁵www.loa.istc.cnr.it/ontologie/DOLCE-Lite.owl

Eventive occurrences (events) are called achievements if they are atomic, otherwise they are accomplishments. Further developments: being ‘achievement’, ‘accomplishment’, ‘state’, ‘event’, etc. can be also considered ‘aspects’ of processes or of parts of them. For example, the same process ‘rock erosion in the Sinni valley’ can be seen as an accomplishment (what has brought the current state that e.g. we are trying to explain), as an achievement (the erosion process as the result of a previous accomplishment), as a state (collapsing the time interval of the erosion into a time point), as an event (what has changed our focus from a state to another). In the erosion case, we could have good motivations to shift from one aspect to another: a) causation focus, b) effectual focus, c) condensation d) transition (causality).

It would be inappropriate to celebrate one or other of the BFO and DOLCE approaches as “correct” and the other as “incorrect”, since their purposes are avowedly different. In our earlier discussion of events *vs* processes, we noted that we can describe what happens from either a historical or an experiential point of view. These different descriptions apply to the same piece of reality. It could be said that what BFO is concerned with is the ontology of those pieces of reality, whereas DOLCE is concerned with the ontology of the ways in which those pieces of reality can be described. Thus stated, the contrast is stark, but it has to be said that it has frequently been obscured in the controversies that have raged between proponents of these two approaches, which tend to portray them as rivals attempting to cover the same ground rather than potential partners focused on different aspects of the world. There is a lot more that could be said here, but it would take us too far afield from our topic.⁶

Historical Note: The classification into “achievements”, “accomplishments”, “processes” and “states” can be traced back to Vendler [29], who used the same set of terms, except that for “process” he used “activity”. The criteria used by Vendler to differentiate these categories are different from, though clearly related to, the ones used by DOLCE. Vendler’s work has been especially influential in the areas of linguistics, philosophy of language, and knowledge representation.

4 Two Alternative Approaches

4.1 The ‘Waterfall’ Model: Processes as continuant-like

A number of researchers have in recent years advocated an approach according to which processes are viewed as, at least in some respects, more like continuants than occurrents [27, 28, 7, 8, 12]. Clearly such a view is incompatible with the treatment of processes in BFO and most other ontologies that have been proposed. In this section, as one of the authors who has proposed such an approach, I will set out its main features. For brevity, I shall refer to this as the “Waterfall” approach, after the title of one of the papers [12] in which it has been advocated.

First, it should be noted that those who have claimed that processes are continuant-like are using the term “process” in a narrower sense, comparable to that of DOLCE, and *not* in the broader BFO sense of any occurrent with temporal proper parts. This is the sense in which, for example, flying is a process, but a flight (e.g., from London to New York) is not a process but an event.

The key motivation for the Waterfall approach is the observation that

a process can undergo change.

If we accept that only continuants, and not occurrents, are capable of change, then this immediately leads to the idea that processes might be continuants rather than occurrents.

In order for processes to be capable of change, they must, contrary to BFO, have qualities with respect to which the changes take place. Moreover, these qualities must be applicable to processes at instants of time, so that they can take different values at different instants. Examples of such qualities, for a flying process, are *speed* and *direction*. At any moment during the course of a flight, the plane is flying at a particular speed and in a particular direction. At different moments the speed and direction may be different — indeed, they might change continuously.

⁶But there can be irreconcilable differences here after all. One criticism that has been levelled against the ‘realist’ philosophy of BFO is that since the only purchase we have on ‘reality’ is via our conceptualisations of it, the *only* route to classifying reality is through classifying concepts. BFO can certainly accommodate different ways of classifying reality, corresponding perhaps to different conceptualisations we have, but will still insist that what is being classified is things in reality, not our concepts.

The Waterfall view is that while an event such as a flight from London to New York is, as generally agreed, a temporally extended entity with proper temporal parts, in other words an occurrent, what this occurrent is, so to speak, “made of” is a continuant-like *process*, flying, which exists as a whole at each moment during the duration of the flight, its properties such as speed and direction being evaluable at each moment and possibly taking different values from moment to moment. When a property takes on different values at different moments, we say that the bearer of the property changes: this standardly applies for continuant bearers (e.g., people, buildings), and the Waterfall view extends these to processes, making them, too, similar to continuants.

This does not mean that processes *are* continuants, though, and in my published work advocating the Waterfall approach I have generally only gone so far as to say that processes are continuant-like in certain respects. But there is another respect in which they retain an essential connection to *bona fide* occurrents, which is that they seem to form the “raw material” for events.

A typical event occurrence consists of some process beginning, going on for a while, and then stopping. The event might be described purely in terms of the process of which it consists — e.g., as a flight, a fight, or a walk — or, more often, in terms of some transformation that is accomplished by the process, e.g., the making of an apple pie, a journey from London to New York. The raw material for a journey comprises various locomotion processes (e.g., walking, flying, driving, travelling on a train) and the raw material for the making of an apple pie comprises processes such as cutting apples, mixing dough, rolling out pastry, and baking.

This is closely analogous to the way in which we might say that a table is made of wood — or even that a piece of wood is made of wood. Thus events are made of processes as objects are made of materials. The properties of cumulativity and homomericity apply to the object/material case just as to the event/process case. This being so, just as both objects and the materials they are made of are continuants, it would seem natural to say that processes are of the same top-level ontological kind as the events that are made of them, namely occurrents. This is recognised in the Waterfall approach, which is why that approach only says that processes have continuant-like features, not that they *are* continuants.

In view of this, it might be said that the Waterfall view casts doubt on the viability of the traditional continuant/occurrent distinction. In [8], I proposed to replace this by a distinction between *experiential* and *historical* entities. Experientials are what can be (in principle) experienced directly in the world; they are entities that exist in a snapshot. This includes not only ordinary objects and the matter they are made of (i.e., traditional continuants) but also processes and states of affairs: the view taken here is that a snapshot of the world is inherently dynamic, that there are processes present here and now. But an event does not exist at a moment: it is extended over an interval and has temporal proper parts corresponding to the proper parts of that interval. Where a process can be seen to be in operation at an instant, an event can only be recognised in relation to a longer perspective that takes durations into account: it belongs, not in the instantaneous snapshot, but in the historical record. Events form the historical aspect of reality.

The technical solution to reconciling the continuant-like and occurrent-like aspects of processes is to postulate a “moving window” model whereby a process with temporal grain-size δt may be thought of as occupying a little temporal window $(t - \frac{1}{2}\delta t, t + \frac{1}{2}\delta t)$ existing as a whole (continuant-like) in the “specious present” centred at t . This window moves forward through time, retaining its identity, but possibly changing its properties, as it does so. Just as a quantum wave-packet partakes of the character of both a wave and a particle, so the process window partakes of the character of both an occurrent and a continuant. The Waterfall model postulates four categories of physical phenomena, whose interrelations are portrayed in Figure 5.

Another important feature of the Waterfall model is its handling of the mutual dependence between objects and processes. An object is described as *enacting* those processes in which it participates; these are called the *outer processes* of the object. As in other ontologies, these processes are understood to be dependent on the objects which are its participants. But an object itself depends for its existence on other processes which *sustain* it: for example, the existence of a human being depends on a host of biological processes such as respiration, digestion, the circulation of the blood, etc, which are internal to the organism. These are described as *inner processes* of the object. These inner processes will themselves be outer processes of those parts of the organism which enact them. The mutual dependence of the categories of object and process does not amount to circularity since the processes which depend on an object are distinct from the processes on which the object depends. I have to exist for my eating, singing, walking, or laughing to occur; but my respiration, digestion, blood circulation, and metabolism

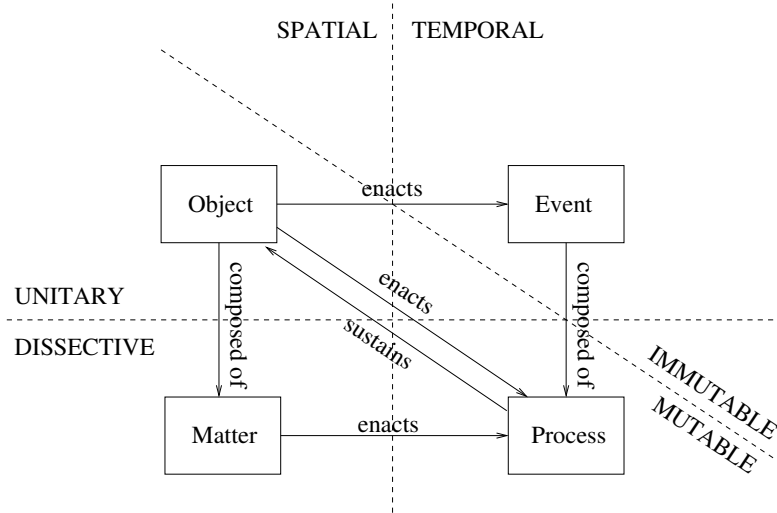


Figure 5: Four categories of physical phenomena, according to the “Waterfall” model (from [12]).

have to occur for me to exist.

These dependencies can be illustrated in diagrammatic form in Figure 6, which also draws attention to the idea that the internal processes of an object are specified by *roles* within the specification of the object itself, and these roles are played by parts of the object, or, in some cases, by objects that are external to the object under consideration. Several different examples of this are given in the original paper.

Although the Waterfall view has many attractive features, it has not won widespread acceptance. From an ontological point of view a problematic issue is what kinds of entities are classified as historical or exponential. I earlier referred to experiential and historical *aspects* and this brings the Waterfall approach closer to DOLCE than BFO, since it appears to be classifying not bits of reality in themselves but bits of reality under a description (one might say: conceptualised in such-and-such a way). Historical entities correspond quite well to DOLCE’s eventive occurrences; experiential entities comprise both stative occurrences and endurants.

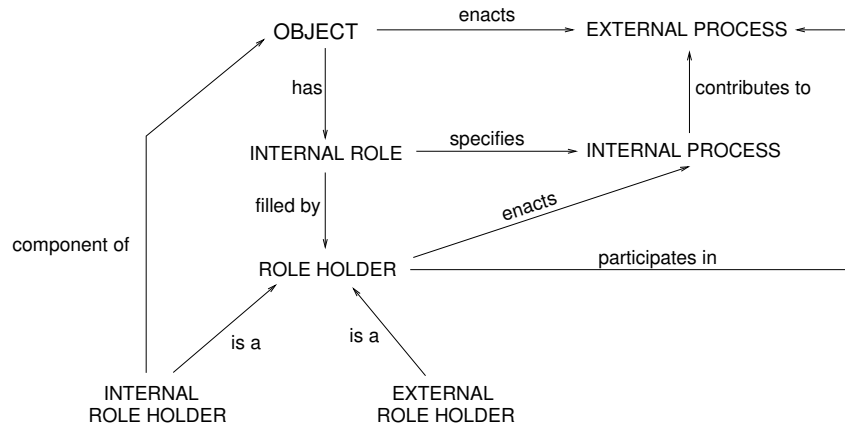


Figure 6: Objects, processes, and their interrelations (from [12])

4.2 Processes as patterns

A different view of processes is suggested by the following quotation from [20]:

Processes are repeatable behaviors whose occurrences cause continuants to undergo changes. It is important to distinguish between processes and process occurrences. Processes are neither endurants (continuants) nor perdurants (occurents), since they do not change and they do not have temporal parts. Process occurrences are perdurants—they may have temporal parts (i.e., sub-occurrences such as changing the coffee filter while making coffee), and they have beginning and end timepoints.

The radical suggestion here is that the notion of “process” should be removed from the realm of continuants and occurents. Processes, on this view, are not themselves occurrences; rather, they are abstract entities whose representatives in the physical world are occurrences which instantiate or realise them.

In this section I will develop this idea in terms of the following general dictum: *A process is a pattern of occurrence*. First we will illustrate this by looking at a simple example, then we will look more closely at the meaning of “pattern”.

Example: Walking. In Vendler’s classification, walking is an *activity*. It is a good example of a (relatively) simple process. It may be characterised as an activity in which a human being⁷ achieves locomotion by means of an alternating sequence of forward movements by the left and right feet ... (the dots here stand for further details required to differentiate walking from e.g., running, skipping, marching). Here walking is specified as a *pattern* of activity, a pattern of alternating footsteps: left, right, left, right, etc.⁸



This pattern is realised concretely in particular occurrences. An occurrence of walking — a walking event — is a temporal part of a person’s life, one which constitutes a realisation of the walking pattern. A particular occurrence of this type might be described in multiple ways; for example, one and the same walking event might be reported using any of the sentences:

Mary walked.

Mary walked for twenty minutes.

Mary walked a mile.

Mary walked to work.

Mary walked along this path.

although in general these sentences will apply to different selections of walking events.

A process, then, shows up in the SPAN world (the world of occurents) as events which realise it. BFO would *call* these events processes, using that term differently from how we are using it in this section. Of course, BFO also has its own category of *realisable* entities; these are dependent continuants such as roles, functions, and dispositions, which are dependent on continuants which realise them by participating in specific processes (for example, the disposition of a doorbell to ring when the button is pressed is realised by those occurrences in which the doorbell does in fact ring as a result of the button being pressed). But I’m not sure how easily our notion of process, as a pattern of activity, can be modelled as a BFO realisable. It would seem to fit more readily into the DOLCE category of *abstracts*, which is separate from both the categories of endurants and perdurants.

The realisation of a pattern of activity shows up historically as an event. Experientially, a pattern of activity is realised as a *state*: when we say that *Mary is walking* we are attributing to her, here and

⁷Here I confine myself to human walking—the details would need to be modified to cover walking by other animals.

⁸The image given here is built up from several copies of an image found at <http://blog.ocad.ca/wordpress/gdes1b26-fw2010-19/category/3-2-project-two/>.




Figure 7: Regular tessellations of hexagons (left) and pentagons (right)

now, the state of being engaged in a realisation of the *walking* activity-pattern. This shows up in the present by virtue of the current spatial configuration and state of motion of her body-parts forming an instantaneous time slice of a continuum of such configurations making up a realisation of walking as an occurrence.

If we accept that processes are patterns of activity then we can distinguish different kinds of process by considering different kinds of pattern. There are spatial patterns and temporal patterns; processes, on this view, are temporal patterns. Where a spatial pattern may be thought of as *a way of filling space*, a process may be thought of as *a way of filling time*.

An important distinction to be made here is between *open patterns* and *closed patterns* [8].⁹ Perhaps the simplest way to illustrate this is with a geometrical example. Compare:

- A tessellation of regular hexagons;
- A tessellation of regular pentagons.

These are illustrated in Figure 7, with hexagons on the left and pentagons on the right.

The hexagonal tessellation is planar: although the illustration only shows ten hexagons, it is clear that any number of additional hexagons could be added to it, following the same pattern, extending indefinitely far over the plane. By contrast, the pentagonal tessellation is non-planar, and moreover it closes in on itself to form a bounded form comprising exactly 12 pentagons that cannot be extended as the hexagonal tessellation can: it is, of course, a dodecahedron.

These patterns are regular; patterns don't have to be regular in this way, but from these examples we can extract two key ideas:

- *Open patterns*, which can be extended indefinitely while still being a realisation of the pattern. Any particular instance of such a pattern will necessarily have a boundary, but the form of the boundary will be imposed from outside and not specified by the pattern itself.
- *Closed patterns*, which are finite, necessarily bounded arrangements of components, the boundary being specified, implicitly or explicitly, as an intrinsic part of the pattern itself.

Examples of the two kinds of pattern in everyday life are *wallpaper patterns* (open) and *dress patterns* (closed). The reason we call these both patterns is because both involve the notion of repetition. But what is repeated is different in the two cases. With an open pattern the repetition is internal to the pattern, e.g., the wallpaper pattern itself consists of an indefinite number of repetitions of some basic motif, and the spatial arrangement of these repetitions is specified as part of the pattern. A closed pattern, on the other hand, just specifies one unit, of which, however, there may be indefinitely many repetitions, e.g., a realisation of a dress pattern is one dress, and there can be many such realisations (many dresses of the same pattern), but the spatial arrangement of these dresses is arbitrary and no part of the dress pattern itself.¹⁰

⁹See also [10].

¹⁰Note that the dodecahedron example we began with is potentially misleading in that we originally introduced it as the result of trying to form a repeating pattern of regular pentagons—which had the mathematical reality been different “might have” produced an open pattern; the closing of the pattern forces us to think of the repeatable unit now as the whole dodecahedron, not the pentagon.

All the above is spatial, but with temporal patterns — processes — we have a closely analogous situation. We have:

- *Open processes.* These are open-ended activities which can in principle be continued in the same way indefinitely. More exactly, the process itself does not specify how its realisations should be bounded. Any actual occurrence of walking must come to an end some time, but the process itself does not determine when this should be, rather this is imposed by external factors (such as the walker arriving at the destination, or simply deciding to stop, or being prevented by an obstacle from continuing). In addition to walking, open processes include other human activities such as eating, reading, knitting, driving, and sleeping, as well as non-human activities such as raining, flowing (of a river), rotting, photosynthesis, the beating of a heart, coastal erosion, and nuclear fusion (e.g., in a star).
- *Closed processes.* These are finite routines for which the start point and end point are specified as part of the process itself and not imposed arbitrarily from outside. Once completed, this realisation of the process cannot continue, although another realisation of the same process could begin. Examples include, with human agency, making a pot of tea, booking a theatre ticket, checking in for a flight, filling in a form, erecting a greenhouse, filling up with petrol, and registering for a course; and in the natural world, cell division, volcanic eruption, maturation (of an organism), the hatching of an egg, a single heartbeat, and a supernova explosion.

Open patterns do not have to have a discrete repeated motif, like a wallpaper pattern. A limiting case of wallpaper patterns is a completely uniform expanse of one particular colour. This pattern is realised concretely in the form of patches of that colour; these patches can be of any shape or size, and can always, in principle, be extended. This is a strictly homogenous spatial pattern. There are strictly homogeneous processes too. Examples include uniform motion in a straight line or in a circle; rolling (of a ball), rotation (of the earth about its axis), sliding, gliding, humming, buzzing, warming, and cooling.

But mostly, homogeneity is relative to the spatial or temporal granularity at which a pattern or process is observed or described: viewed from a distance, a tessellation of red and blue squares looks like a uniform expanse of purple. The humming of the radiator (a homogeneous open process), on closer examination, is revealed to consist of a repeated cycle of back and forth vibrations. We experience the emission of light by a fluorescent tube as homogeneous; but to an insect whose flicker-fusion rate is higher than the frequency of the alternating current powering the tube, it would appear as a sequence of bright flashes interspersed with intervals of darkness.

The simplest kind of closed pattern is just a “chunk” of some open pattern. In the spatial case the closed pattern can be realised by delineating a region of space and filling it with a realisation of the open pattern. A chessboard realises a closed pattern which consists of a square area filled with 64 squares’ worth of the open pattern “tessellation of black-and-white squares”.

In the temporal case, a similarly simple example of a closed process is “walking a mile”. This is realised as a chunk of walking for which the distance covered is one mile. Another example is “walking for an hour”, which is realised as a chunk of walking lasting one hour.

You can specify more complex closed processes as assemblies of simple closed process components. An example is the closed process specified by the following sequence of instructions:

Repeat four times: Walk in a straight line for one mile and then turn 90° clockwise.

This specifies a process of walking once round the perimeter of a square mile. Both “walking” and “turning” are open processes but “walking a mile” and “turning 90° clockwise” are closed processes, specified as chunks of the corresponding open processes; the instructions compose these simple closed processes into a single complex process. This shows how closed processes generally can be specified in terms of open processes.

But open processes can be specified in terms of closed processes, too; this is done through the basic operation of *repetition*. We could specify an open process by means of an open-ended repetition of the closed “walk round a square mile” process. The instruction could be “keep walking round the square mile”. It does not say when to stop, but anyone following the instruction must eventually do so; when this happens is not determined by the instruction (i.e., by the process specification) but by external factors.

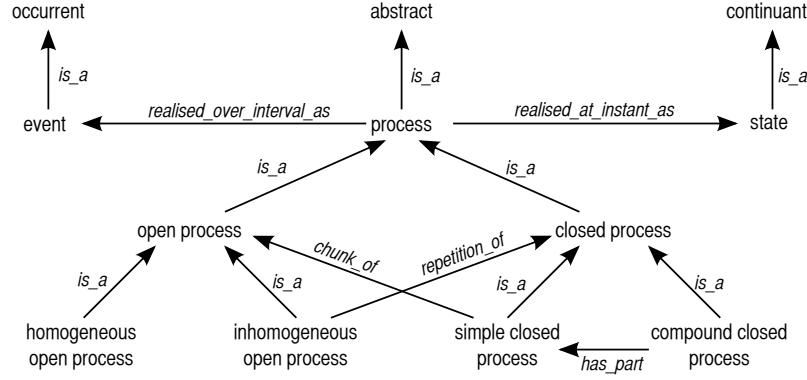


Figure 8: Processes as patterns of activity

The overall picture we have described here can be presented as a little fragment of an ontology, shown in Figure 8.

I have explained how the experiential realisation of an open process is a state: the state reported in, for example, “Mary is walking”. From the diagram it appears that closed processes should have experiential realisations too. Consider:

“What is Mary doing?” — “She’s walking a mile.”

While the idea of a state of walking seems acceptable, one might well question whether there can be such a thing as a state of walking a mile. If you see someone walking, you see them in a particular state which you can identify as walking on the basis of what you see. But if they happen to be walking a mile, you cannot see this. So the question is, what is it about the present state of affairs that makes their walking a walking-for-a-mile rather than some other species of walking? This is the notoriously difficult problem of specifying the semantics of the progressive aspect for accomplishments. To count as a state of walking-a-mile, a walking state must include more than just the appropriate physical configuration and state of motion of body parts; it must include a mental element as well, in this case Mary’s intention to keep on walking until she has completed the mile. There is, of course, much more to be said here, and it is probably fair to say that we do not yet have a definitive account of how sentences of the same general form as “Mary is walking a mile” are to be interpreted; I shall not pursue this further here.¹¹

5 Modelling Causation

When thinking about causation we should be careful to distinguish between *causal facts* and *causal laws*. Causal facts are about causal relations between particular occurrences, for example

CF *A loud noise made me jump.*

Causal laws are generalities, expressing causal relation between *classes* of occurrence, as in

CL *Loud noises make me jump.*

An obvious question is how causal facts are related to causal laws, and in particular which category is more fundamental. The argument could, in principle, go either way:

1. Causal facts are fundamental; causal laws are *summaries* of them. On this view statement CL is only true by virtue of a sufficient number of particular occasions on which CF is true.
2. Causal laws are fundamental, and provide the only ground for asserting causal facts. Without CL, the most one could say is that there was a loud noise, and I jumped; we need CL to justify the assertion CF.

¹¹For further discussion see, e.g., [4, 5].

From a philosophical point of view the key question about causality is perhaps

What does it mean to say that something causes something else?

However, this will not be the focus of the present section. Instead, we shall look at the *structure of causal facts*. In particular, we shall be interested in different *types* of causal or causal-like relations, and think about what exactly it is that they relate.

It seems that many different kinds of things can play the role of cause. Consider a simple scenario in which a driver on an icy road brakes suddenly and skids. What should we say is the cause of the car's skidding? It could be any of

- An agent — the driver caused the car to skid.
- An object — the ice on the road caused the car to skid.
- A property — the slipperiness of the ice caused the car to skid.
- A state — the road's being slippery caused the car to skid.
- An event — the driver's braking caused the car to skid.
- A fact — the fact that the driver braked when the road was icy caused the car to skid.

Equally, there are many candidates for the role of effect. Consider what things are caused in the above scenario. Any of the above causes might be said to be the cause of:

- An event — the car starting to skid
- A process — the skidding of the car
- A fact — the fact that the car skidded
- A property — the car's being out of control
- An object — a wreck by the roadside.

Instead of referring all these to a single causal relation “causes”, Mike Worboys and I [13] proposed a collection of *different* causal and causal-like relations, distinguished by different classes of relata. These ideas were developed further in [9], and led to some practical work in methods of detecting different kinds of causal relations in real data [3, 11]. In this section I shall present an outline of this theory, and illustrate it with some case studies.

Following [6] we will use the word ‘eventuality’ here as a general term to cover whatever goes on in time. Eventualities can be classified along two dimensions: *continuous* vs *non-continuous* and *dynamic* vs *non-dynamic*. A continuous non-dynamic eventuality is a *state*; a continuous dynamic eventuality is a *process*; and a non-continuous dynamic eventuality is an *event*.

Topic for discussion: Are there any non-continuous non-dynamic eventualities? What would they look like?

We now distinguish three different causal relations, as follows:

- An event can *cause* another event.
- A process can *perpetuate* another process.
- A process or state can *maintain* another state.

In addition we distinguish three causal-like *constitutive* relations:

- A process can *result in* an event.
- An event can *initiate* or *terminate* a process or state.

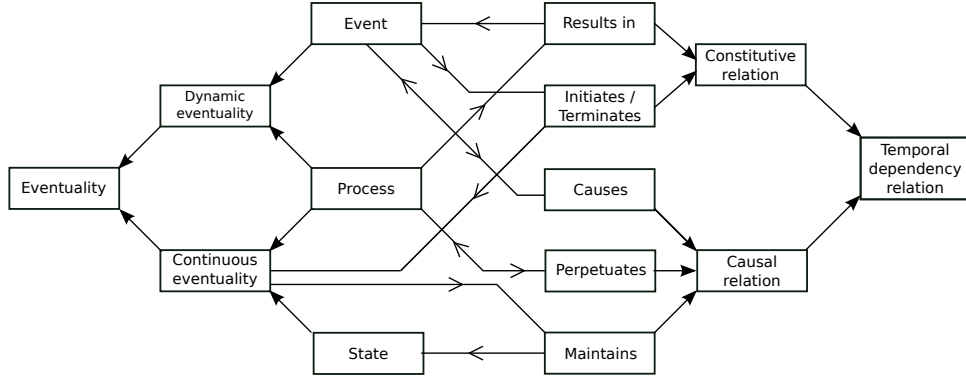


Figure 9: Taxonomy of Eventualities and their Relations

These constitutive relations are non-causal temporal dependency relations which reflect logically necessary connections between certain eventualities. The relation ‘results in’ is used of those cases where it follows as a matter of logical necessity that the operation of a process, so long as it is not interrupted, will have a certain event as an outcome. For example ‘The train is approaching the station’ can result in ‘The train arrives at the station’. It would be very odd to say that the approach *causes* the arrival; we need a different term for this. Again, an increase in the level of liquid in a container can result in the container’s becoming full. These examples also furnish examples of our second constitutive relations. The train’s arriving at the station (an event) initiates the train’s *being* at the station (a state). Again, it would be odd to speak of causation here, since essentially the train’s arrival at the station is nothing over and above the initiation of its being there, so the connection here has a flavour of logical necessity that is missing from strict causation. Similarly, the container’s *becoming* full (an event) initiates its *being* full. The third constitutive relation, termination, is just the opposite of initiation. The example, the train’s coming to a halt terminates its motion.

It is important to stress that all these relations are to be understood as applying to particulars, not to types. They feature primarily in causal facts; insofar as they feature in causal laws, it is because causal laws have causal facts as instances. The taxonomies of eventualities and their temporal dependency relation (causal and constitutive) are shown in Figure 9. The eventualities are shown on the left, the relations on the right. Solid-headed arrows are “is a” links; the open-headed arrows linking the two halves of the diagram show what types of eventuality are related by each of the relations, e.g., a process results in an event — shown by an arrow from “process” to “results in” and an arrow from “results in” to “event”.

In practice we find certain patterns of relations recurring with great frequency when we analyse different scenarios using this “tool-kit”. An example is the *perpetuation pattern* shown in Figure 10. Here the initiation of process X causes the initiation of process Y. The continuation of Y is then perpetuated by the continuation of X. The termination of X causes the termination of Y. (In the diagram, time is read from left to right.) An example is a box moving across the floor (Y) because someone pushes it (X).





Figure 10: Perpetuation pattern

First of all I start pushing; this causes the box to start moving, and initiates my pushing process. This process perpetuates the box's moving process that was initiated by starting of the motion. Eventually the process is terminated when I stop pushing. This causes the box to stop moving — the cessation of motion being the event which terminates the motion process.

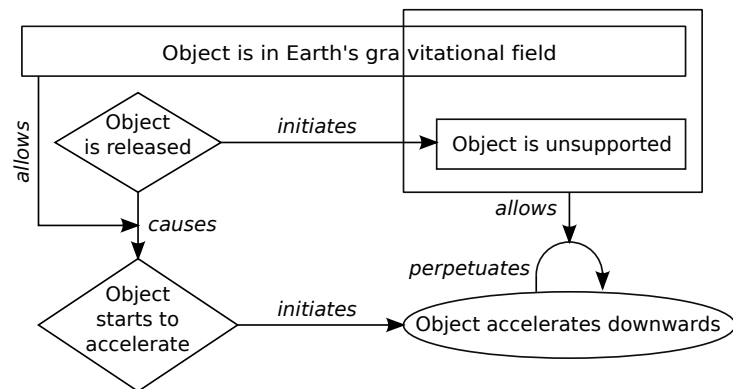
Case study 1: The Water Tank

Case Study 2: A Kettle of Water is Boiled


Case Study 3: Unsupported object in a gravitational field



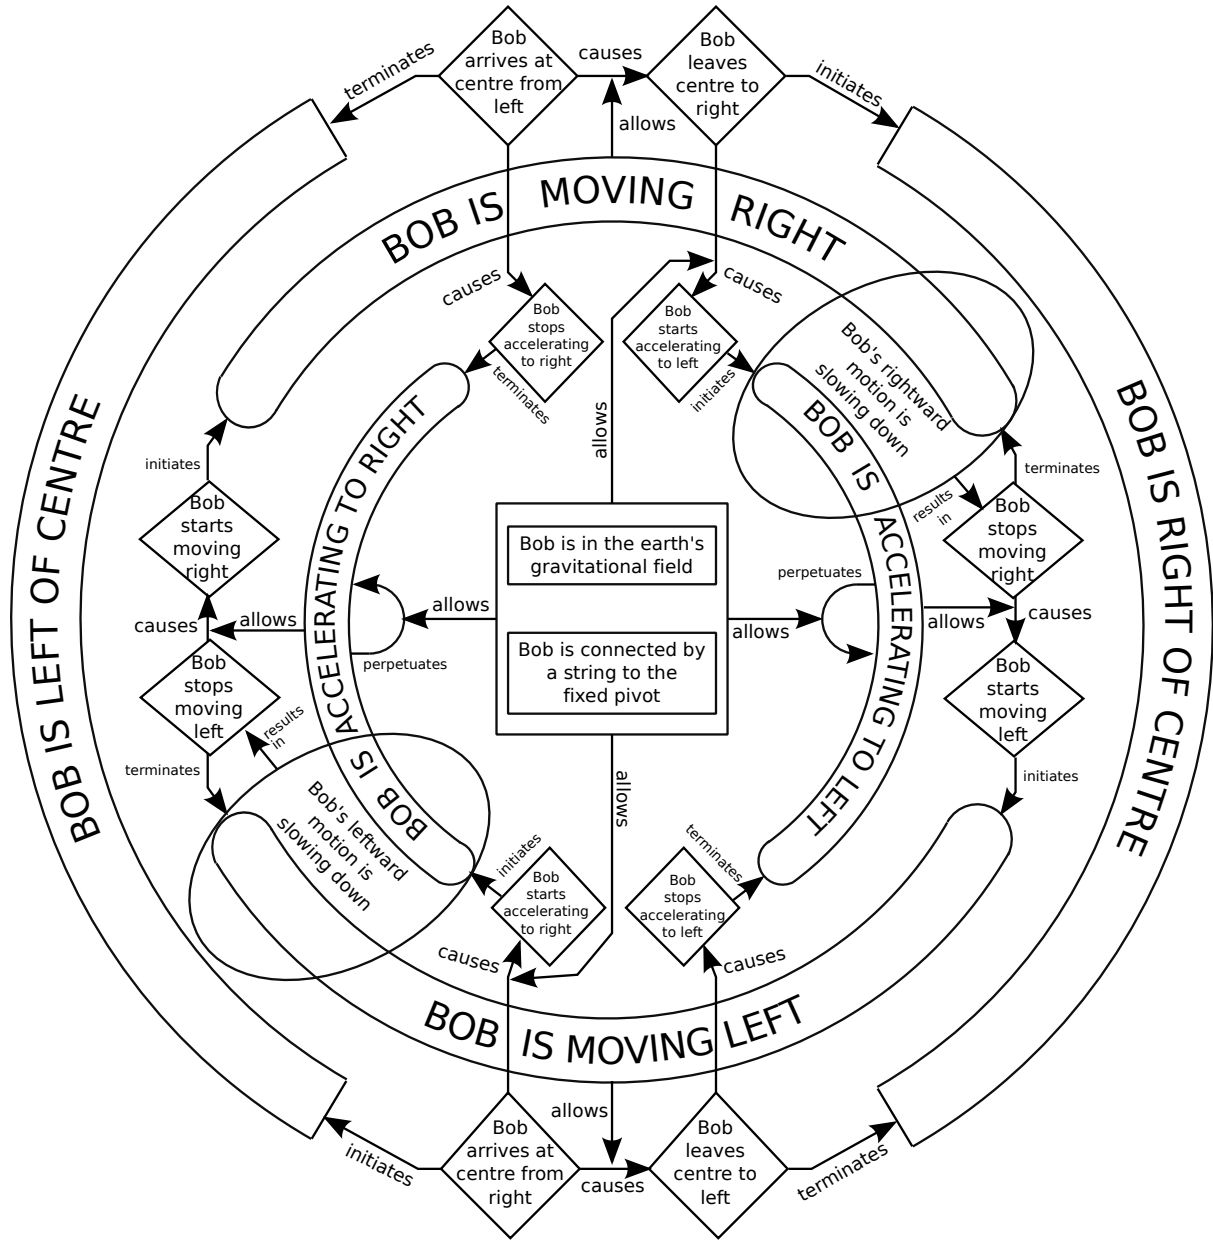
Case Study 4: An object is dropped from a height



Case Study 5: A ball is thrown up into the air



Case Study 6: A Simple Pendulum



Note that in this diagram each eventuality box here stands for multiple particular instances. The circular form of the diagram should be understood as “unwrapped” to form a line, in which each eventuality box denotes a unique particular occurrence.

References

- [1] James Allen. Towards a general theory of action and time. *Artificial Intelligence*, 23:123–54, 1984.
- [2] Robert Arp, Barry Smith, and Andrew D. Spear. *Building Ontologies with Basic Formal Ontology*. The MIT Press, Cambridge MA and London UK, 2015.

- [3] Susanne Bleisch, Matt Duckham, Antony Galton, Patrick Laube, and Jarod Lyon. Mining candidate causal relationships in movement patterns. *International Journal of Geographical Information Science*, 28(2):363–382, 2014.
- [4] David Dowty. *Word Meaning and Montague Grammar*. D. Reidel, Dordrecht, 1979.
- [5] Antony Galton. *The Logic of Aspect: an Axiomatic Approach*. Clarendon Press, Oxford, 1984.
- [6] Antony Galton. Eventualities. In M. Fisher, D. Gabbay, and L. Vila, editors, *Handbook of Temporal Reasoning in Artificial Intelligence*, pages 25–58. Elsevier, Amsterdam, 2005.
- [7] Antony Galton. On what goes on: The ontology of processes and events. In B. Bennett and C. Fellbaum, editors, *Formal Ontology in Information Systems: Proceedings of the Fourth International Conference (FOIS 2006)*, pages 4–11. IOS Press, 2006.
- [8] Antony Galton. Experience and history: Processes and their relation to events. *Journal of Logic and Computation*, 18(3):323–340, 2008.
- [9] Antony Galton. States, processes and events, and the ontology of causal relations. In M. Donnelly and G. Guizzardi, editors, *Formal Ontology in Information Systems: Proceedings of the Seventh International Conference (FOIS 2012)*, pages 279–292. IOS Press, 2012.
- [10] Antony Galton. Outline of a formal theory of processes and events, and why GIScience needs one. In S. I. Fabrikant, M. Raubal, M. Bertolotto, C. Davies, S. Freundschuh, and S. Bell, editors, *Spatial Information Theory, Proceedings of the 12th International Conference, COSIT 2015*, pages 3–22. Springer, 2015.
- [11] Antony Galton, Matt Duckham, and Alan Both. Extracting causal rules from spatio-temporal data. In S. I. Fabrikant, M. Raubal, Michela Bertolotto, Clare Davies, Scott Freundschuh, and Scott Bell, editors, *Spatial Information Theory, Proceedings of the 12th International Conference, COSIT 2015*, pages 23–43. Springer, 2015.
- [12] Antony Galton and Riichiro Mizoguchi. The water falls but the waterfall does not fall: New perspectives on objects, processes and events. *Applied Ontology*, 4(2):71–107, 2009.
- [13] Antony Galton and Michael Worboys. Processes and events in dynamic geo-networks. In M. A. Rodríguez, I. F. Cruz, S. Levashkin, and M. J. Egenhofer, editors, *GeoSpatial Semantics: Proceedings of First International Conference, GeoS 2005, Mexico City, Mexico, November 29-30, 2005*, pages 45–59. Springer-Verlag, 2005.
- [14] Pierre Grenon and Barry Smith. SNAP and SPAN: Towards dynamic spatial ontology. *Spatial Cognition and Computation*, 4(1):69–104, 2004.
- [15] Mark Heller. *The Ontology of Physical Objects: Four-dimensional Hunks of Matter*. Cambridge University Press, Cambridge, 1990.
- [16] John Lyons. *Semantics*. Cambridge University Press, Cambridge, 1977. 2 volumes.
- [17] C. Masolo, S. Borgo, A. Gangemi, N. Guarino, and A. Oltramari. WonderWeb deliverable D18. Technical report, Laboratory for Applied Ontology, ISTC-CNR, Trento, Italy, 2003.
- [18] Marc Moens and Mark Steedman. Temporal ontology and temporal reference. *Computational Linguistics*, 14:15–28, 1988. Reprinted in Mani, I., Pustejovsky J, and Gaizauskas, R, *The Language of Time: A Reader*, Oxford, 2005.
- [19] A. P. D. Mourelatos. Events, processes, and states. In Philip Tedeschi and Annie Zaenen, editors, *Tense and Aspect*, pages 191–212. Academic Press, New York, 1981.
- [20] A. Özgövde and M. Grüninger. Foundational process relations in bio-ontologies. In A. Galton and R. Mizoguchi, editors, *Formal Ontology in Information Systems: Proceedings of FOIS 2010*, pages 243–256. IOS Press, 2010.

- [21] J. Pustejovsky. The syntax of event structure. *Cognition*, 41:47–81, 1991.
- [22] Johanna Seibt. Process philosophy. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Spring 2016 edition, 2016.
- [23] Theodore Sider. *Four-dimensionalism: An Ontology of Persistence and Time*. Oxford University Press, Oxford, 2001.
- [24] Barry Smith. *Basic Formal Ontology 2.0: Specification and User’s Guide*. Accessible from <https://github.com/BFO-ontology/BFO>, 2015.
- [25] John F. Sowa. *Knowledge Representation: Logical, Philosophical, and Computational Foundations*. Brooks/Cole, Pacific Grove, CA, 2000.
- [26] J. G. Stell and M. West. A four-dimensional mereotopology. In A. C. Varzi and L. Vieu, editors, *Formal Ontology in Information Systems: Proceedings of the Third International Conference (FOIS 2004)*, pages 261–272. IOS Press, 2004.
- [27] Rowland Stout. Processes. *Philosophy*, 72:19–27, 1997.
- [28] Rowland Stout. The life of a process. In Guy Debrock, editor, *Process Pragmatism: Essays on a Quiet Philosophical Revolution*. Rodopi, Amsterdam/New York, 2003.
- [29] Zeno Vendler. Verbs and times. In *Linguistics and Philosophy*, chapter 4. Cornell University Press, Ithaca, 1967. Reprinted in Mani, I., Pustejovsky J, and Gaizauskas, R, *The Language of Time: A Reader*, Oxford, 2005.
- [30] Michael F. Worboys. Event-oriented approaches to geographic phenomena. *International Journal of Geographical Information Science*, 19(1):1–28, 2005.