Applied Ontology 0 (2009) 1–0
IOS Press

The Water Falls but the Waterfall does not Fall: New perspectives on Objects, Processes and Events

Antony Galton [a] and Riichiro Mizoguchi [b]

[a] *School of Engineering, Computing and Mathematics, University of Exeter, UK*
[b] *The Institute of Scientific and Industrial Research, Osaka University, Japan*

Abstract. We challenge the widespread presumption that matter and objects are ontologically prior to processes and events, and also the less widespread but increasingly popular view that processes and events are ontologically prior to matter and objects. Instead we advance a third view according to which each of these pairs of categories is ontologically dependent on the other. In particular, taking a cue from an ontology of devices, we identify the object as an interface between those processes which are internal to it and those which are external to it and which it may be said to enact, thereby linking objects intrinsically to the processes in which they are involved as well as providing a more powerful determinant of object identity than more traditional, non-dynamic criteria based on demarcation from the environment. The internal processes are themselves external processes in relation to the components of the object which enact them, leading to a potentially open-ended recursive decomposition of both objects and processes in a complex web of mutual interdependency. We also discuss how matter is related to objects, and processes to events, bringing the four categories together in a diagram which clarifies the relations between them — often considered problematic — and establishes a framework for a highly general top-level ontology.

Keywords: object, process, event, change, device, role

1. Introduction

We begin by contrasting two views on the relationship between objects and processes, on the one hand the traditional object-centred view that is reflected in the ontologies that have dominated western thought at least since Aristotle, and on the other hand the radical process-centred view which has repeatedly surfaced in the history of philosophy, even as far back as Heraclitus, and more recently with Bergson, James, and Whitehead [35].

- The 'object-priority' view: Matter and objects are prior to processes and events. Viewed as an ontological claim, the 'object-priority' view asserts that in some important sense matter and objects are all that there is in the world; events and processes exist solely by virtue of the distribution of matter and objects in space and time.
- The 'process-priority' view: Processes and events are prior to matter and objects.[1] This too is usually presented as an ontological claim, that processes and events are all that there is. Matter and objects are to be analysed as in some way constructed from, or emergent from, processes and events. It is sometimes said that what we commonly call an object is in fact an event or process, although it is hard to assert this using ordinary vocabulary without courting gross categorical confusion.[2]

[1] For a general overview of the process-priority view, see [35]. A recent collection of papers with a broadly 'processist' orientation is [39]. A 'process ontology' is one of four fundamental types of ontology identified by Zemach in [57]; he does not, however, accord any priority to the process ontology, pointing out that each of these ontologies is self-sufficient and need not be used in conjunction with any of the others. This contrasts with the position of [17], whose object-based ('SNAP') and process-based ('SPAN') ontologies are treated as complementary, neither being complete in itself — but again, with neither being given priority over the other.

[2] This has been pointed out on numerous occasions, e.g., by Wiggins [53, p.25, n.12] in response to a remark by C. D. Broad that there is no important categorical difference between a flash of lightning and the cliffs of Dover.

In both these cases we see a reductionist tendency at work, perhaps motivated by a sense of embarrassment at an over-rich ontology.

In this paper we advance the view that neither of matter/object nor process/event is ontologically prior to the other; but rather, each is dependent on the other. To put it as simply as possible (the complications and ramifications will come later), (a) *matter and objects by nature presuppose the participation in processes or events*, and (b) *processes and events by nature presuppose the existence of matter or objects.* This view is not in itself new: for example, Moravcsik [32] argued that "there is a mutual dependency between material bodies, events, actions and processes". Citing an assertion by Strawson that "we could not have the concept of a strike without having such concepts as those of tools and factories", he counters this with the observation that "neither could we have the concepts of tools and factories without having such concepts as those of production, manufacturing, &c." We are in agreement with this; but we believe that the arguments presented in this paper are more detailed and far-reaching than those advanced by Moravcsik. Likewise, Davidson [5], again arguing against Strawson, argues for a "symmetry of conceptual dependence" between objects and events. Again, we are in agreement with this, but would include processes as well as events in the dependence.

In relation to (a), we note that even maintaining a state of inactivity is a limiting case of process.[3] Indeed, where we see inactivity at one level of granularity there is typically substantial activity at another level. For example, a person sitting 'doing nothing' is all the while breathing, maintaining a posture through continuous muscular adjustments, and, inside the body there is a continuous activity of metabolic processes, digestion, circulation of the blood, and so on. This can be pursued below the cellular to the molecular or even the atomic or subatomic levels. In fact, if miraculously all activity at every level were to cease, Sleeping Beauty-like, then the world would surely cease to exist. This highlights the inadequacy of any attempt to model the history of the world as a succession of (possibly 'continuum many') 'snapshots' each of which is individually static — yet models of this kind are precisely what seems to come most naturally to anyone seeking to incorporate the time-dimension into a pre-existing atemporal information system.[4]

As regards (b), we note for now that any process or event must have some material 'host' which may be said to enact it — here 'material' should be understood in a broad sense to include, for example, energy and fields of force. On that understanding, there can be no processes or events in empty space, just as there can be no moves on an empty chessboard.

The purpose of this paper is not to provide a new formal ontology or to make specific corrections to existing ones. Rather, it is intended to lay down some general ideas which we hope will in due course be able to form the basis for a radical new ontology in which the mutual interdependence between key concepts such as 'object', 'process', and 'event' plays a key role. Taking as it does the form of a preparatory groundwork for such an enterprise, the presentation is largely discursive rather than formal-analytical, although there will be pointers towards a more formal treatment in some places. Many of the ideas discussed here are not presented with a claim to originality; some of them belong to a common stock of received wisdom and may indeed have the appearance of 'plain common sense', but where we are aware of specific antecedents we have cited them. Where we would wish to advance a claim to originality is in the overall synthesis of ideas to form a picture which, we believe, has not been explicitly articulated previously, and which, while not yet formalised, may provide a useful starting point for a new formalisation.

The particular view of the world presupposed by our discussions is that of an informed and reflective common sense. We are not concerned, or not immediately so, with aspects of the world that are revealed by sciences such as physics, chemistry, and biology, but which are inaccessible to common sense except through the intermediary of specialist training, equipment, or vocabulary. In this regard, our approach somewhat resembles Artificial Intelligence research in the field of 'commonsense reasoning' [21, 6, 23], although as we have stated, we have not yet ventured onto the level of formalisation to be found in many

[3] Salmon expresses this from a processist perspective, assimilating objects to processes: "As I shall use these terms, even a material object at rest will qualify as a process" [37, p.140].

[4] Compare Worboys [56]: "Until now, the most common approach to spatiotemporal models has been the view of the world as a succession of temporal snapshots of spatial configurations of objects".

such works. Thus we are concerned to model reality as it appears to humans engaged in ordinary human activities; for this reason the kinds of consideration admissible as 'evidence' for our proposals may be drawn for the most part from everyday experience and the language we use to describe it. We would be the first to admit, however, that this is not a precisely defined domain; it shades off into what may be roughly characterised as 'the scientific worldview' on one side, into analytical philosophy on another, and indeed eventually into all human enterprises, since they all take the commonsense world as their starting point. For this reason, some of what we say may be informed by ideas originating in areas somewhat removed from the everyday, but we do not believe that this compromises the essentially grounded and 'down-to-earth' nature of our theorising.

2. Objects, Events, Matter, and Process

An important goal of this paper is to establish the relationships amongst the key ontological categories of matter, object, process and event. We will take as our starting point a brief discussion of objects and events. Most ontologies draw a sharp distinction between objects and events, and we do not disagree with this.

Some key respects in which objects and events are often said to differ are as follows:

1. *Relation to time*. An object is present as a whole at each moment of its existence; an event only exists as a whole across the interval over which it occurs.
2. *Nature of parts*. An object can have spatial parts, but does not have temporal parts; an event has temporal parts and may or may not have spatial parts.
3. *Change*. An object can have different properties at different times, and is therefore able to undergo change; it does not make sense to speak of an event changing.

All of these points will be amplified further in what follows. Suffice it to say that these distinctions are widely (if not universally) accepted in the ontological literature, and we take them to be largely uncontentious. It should also be noted that the three distinctions are strongly interrelated: it is hard to accept any one of them without thereby also accepting the others. Together they characterise the distinction between *continuants* and *occurrents* that is widespread in both the philosophical and ontological literature.[5]

As well as these substantial and important differences, objects and events exhibit some important similarities. Like the differences, these are interrelated, and perhaps amount to merely different ways of referring to the same property. They are

4. *Discreteness*. Both objects and events are discrete individuals which may be referred to using count nouns.[6]
5. *Non-dissectivity*. The parts of an object or event are not themselves objects or events of the same type, e.g., half of a chair is not a chair, the first half of a walk to the station is not a walk to the station.
6. *Definite extension*. Objects and events have well-defined extensions: an object occupies a region in space, an event takes up an interval in time.[7]

[5]The terms *endurant* and *perdurant* are also used, sometimes but not always synonymously with *continuant* and *occurrent* respectively. There is an enormous literature concerned with the various kinds of entity that should be recognised in a metaphysically robust ontology, along with their properties and interrelationships; the central difficulties are concerned with how the various types of entity are related to time, and in particular what may be called their modes of persistence. The view here characterised as 'widely accepted' has also been widely attacked, and much of the literature consists of the cross-fire between these camps. See, for example, [41, 22, 20, 40, 19].

[6]In the case of events, these may be nominalisations of verbs, e.g., a walk or a run, but in many cases they are not, e.g., a battle, an accident.

[7]An exception here is provided by what may be called 'boundary objects' and 'boundary events'. A boundary object is not a material object but depends on a material object for its existence, e.g., the surface of a ball, the tip of a knife. Similarly a boundary event is an event which is not constituted by any process, but which depends for its existence on some process, e.g., a ball's starting to move, or reaching the highest point of its trajectory — these are the 'achievements' of Vendler's well-known

	SPACE	TIME
COUNT	object	event
MASS	matter	process

Fig. 1. The four key categories of physical phenomena

The negations of these three properties characterise a pair of categories that are related to, but distinct from, objects and events, namely *matter* and *process*. Whereas objects and events are discrete individuals, matter and process may be regarded as the 'stuff' from which those individuals are made. Matter and process both have the following properties:

7. *Non-discreteness*. Matter and process do not come in the form of discrete individuals, being referred to by mass nouns (e.g., 'milk', 'motion') rather than count nouns.[8]
8. *Dissectivity*. The parts of a portion of matter or process are portions of the same type of matter or process (e.g., half of an amount of water is an amount of water, half of a period of singing is a period of singing).
9. *Indefinite extension*. Matter and process can occupy arbitrarily small regions of space and time respectively: we can think of matter equally as filling a region and its subregions or as existing at each point in any region it fills; and of process as going on throughout an interval and its subintervals or at each instant during any interval throughout which it is going on.

Here we have used the presence or absence of dissectivity and indefinite extension to distinguish matter and processes on the one hand from objects and events on the other. It should be noted that in the case of matter and objects, these properties must be interpreted with respect to space (i.e., unlike objects, matter is spatially dissective and indefinitely extended in space), whereas in the case of process and events, they are interpreted with respect to time (so processes, but not events, are temporally dissective and indefinitely extended in time). In this way we can contrast matter and objects on the one hand (as 'spatial' entities) with process and events on the other ('temporal' entities). The specific distinctions between objects and events listed as (1)–(3) above do not, however, readily extend to matter vs process.

The relationships amongst the four terms can be illustrated in the diagram in Figure 1. This is a beautifully neat picture, which in some form or other has been widely recognised in the literature [1]. In the next section we shall introduce a disturbing influence which will force us to complicate the picture; this complication may be unwelcome, but in fact it will provide the key to a deeper understanding of what the picture represents.

Some authors have tried to characterise the distinction between process and event as one of *duration*. Thus Salmon [37]: "The main difference between events and processes is that events are relatively localized in space and time, while processes have much greater temporal duration ... In space-time diagrams, events are represented by points, while processes are represented by lines." It seems to be easy to mistake the boundedness of events for shortness of duration. Thus Sylvan [48] states "Events themselves are not normally regarded as processes, as they do not go on suitably, but happen and are finished. But no doubt

classification [51], where they sit alongside 'accomplishments', 'activities', and 'states' (see also [31]). Processes, as described in this paper, correspond to Vendler's activities, while events correspond to his achievements and accomplishments. It should be noted, however, that our processes include many apparently 'passive' items of the form *persisting in such-and-such a state*. We would maintain that this process of persistence is conceptually distinct from the state itself, even though in some cases it may amount to little more than 'doing nothing'.

[8]Note that we very often refer to processes using gerunds, e.g., 'walking', 'singing'. These generally behave as mass nouns rather than count nouns: 'How much singing have you done this week?'. It is only in the rather artificial discourse of philosophers that we frequently encounter talk of 'walkings' and 'singings'.

under a different stretching of the term *process*, events may be encompassed, as point or short duration processes." We maintain, on the contrary, that so far from being a mark of short duration, boundedness is a precondition for the assignment of *any* definite duration: processes endure, but only once we have assigned bounds to them can we speak of duration, and the act of assigning bounds means that we have switched our attention from the process to an event. However, underlying this confusion there is a point worth making: the boundedness of an event means that there will always be some temporal granularity at which it is appropriate to conceptualise it as a point. With processes, on the other hand, there is a granularity at which it is conceptualised as a homogeneous line. To switch between these conceptualisations is to view some happening alternately as an event and a process, according as the boundedness or the 'ongoingness' is highlighted.

The property here called dissectivity has often been referred to under the name *homogeneity*; Vendler, for example, [51] notes that processes such as running 'go on in time in a homogeneous way; any part of the process is of the same nature as the whole'. On the other hand, it has also been argued [7, 49] that processes are only homogeneous down to a certain granularity: at sufficiently fine temporal resolution, the running process is seen as a succession of alternating leg movements, none of which on its own would constitute running. However, although it is true that there cannot be an episode of running lasting only one millisecond, this does not mean that running does not happen during a millisecond within a ten-minute period of running. The movement occurring during that millisecond counts as running by virtue of its contribution to the running that occurs over the longer period in which it is embedded. We should like to say, in fact, that the running process comes into operation when the running starts, and operates continuously from then until the running stops. This is even the case with those processes — such as the dripping of a tap — in which there are clear gaps in which nothing seems to be happening; we can say that the tap was dripping all day, just as we might say that a bowl is filled with rice notwithstanding the many interstices between the individual rice grains.

While existing systems of formal ontology generally agree on the necessity to distinguish continuants and occurrents at a fundamental level, the literature displays a strong tendency to gloss over the distinction between processes and events. The chief exception is DOLCE [27], whose treatment of occurrents (perdurants, in DOLCE terminology) mirrors Vendler's four-fold division into achievements, accomplishments, states, and processes [51]. In DOLCE, only the first two of these categories count as events, which are thereby clearly separated from processes. In BFO [10], the class of occurrents is divided into spatiotemporal regions, temporal regions, and "processual entities", but within the latter category, the most pertinent distinction from our present perspective is between processes and boundaries of processes. Process boundaries correspond to what we referred to as boundary events above — all other events would be incorporated into the BFO category of processes. It is interesting to note that BFO treats states as continuants, not occurrents, thus radically disagreeing with DOLCE's classification of states and processes as 'stative' perdurants. We will argue below that processes have a continuant-like aspect: this point of view can be seen as combining DOLCE's alignment of states with processes and BFO's treatment of states as continuants. Like BFO, other ontologies such as SUMO [34] and CYC downplay or ignore the distinction between processes and events.

The relationships amongst our four categories are illustrated in Figure 2. As already noted, matter and process are the 'stuff' of which objects and events are composed. The simplest kinds of object and events are simply 'chunks': a chunk of matter (e.g., a piece of stone, that is, a quantity of stone bounded by a surface), or a bout or episode of some process (e.g., a shower of rain — that is, a period of rain bounded by a beginning, when the rain starts, and an ending, when the rain stops). Specifically, if a process P comes into operation at time t_1 and continues until t_2, when it ceases, then this constitutes an event which takes place on the interval $[t_1, t_2]$. Note that because processes are dissective but events are not, we can say that the process P was in operation over any subinterval of $[t_1, t_2]$, or even at any instant in its interior; but the event does *not* occur on any proper subinterval of $[t_1, t_2]$ — this is easy to see, since although the process P is in operation over the subinterval, it is not, in that subinterval, bounded by both a beginning (when P starts) and an ending (when P stops).

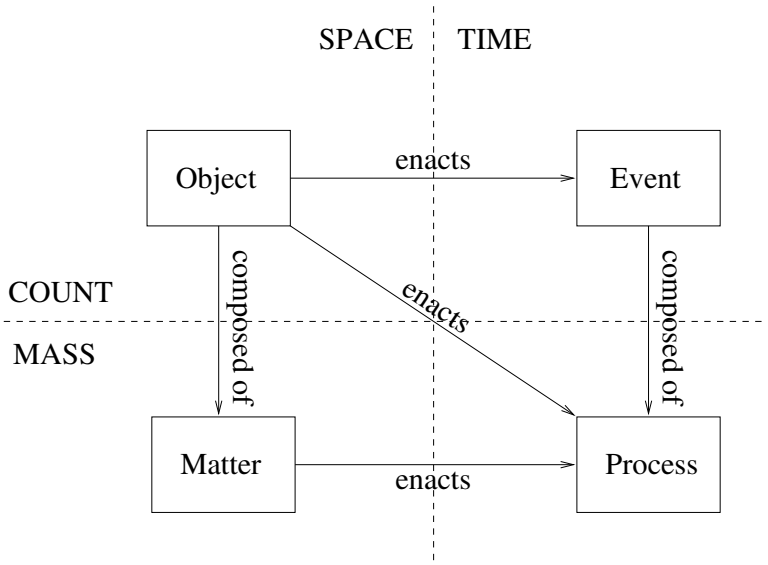

Fig. 2. The four key categories of physical phenomena

Most objects and events for which we have names are more complex than these most simple ones. A table is not a chunk of 'table-stuff' (for there is no such stuff); it consists of matter arranged in such a way as to fulfil the function of allowing objects of a certain range of sizes to be supported at a suitable height above the ground to be manipulated by seated or standing humans. Typically it will consist of several simple chunks of matter (though a plastic table moulded in one piece is itself a simple chunk). Similarly, an event such as a journey might consist of a sequence of qualitatively distinct process-chunks, e.g., an episode of walking, an episode of waiting at the station, an episode of boarding the train, etc, which have to be ordered in a certain way in order for the ensemble to qualify as a journey. This is all familiar enough; we mention it here in order to make it clear that even though the manner in which the chunks of matter or process are assembled may be indefinitely complicated, still it remains true that objects are made of matter and events are 'made of' process. These relationships are shown by the vertical arrows in the diagram.

The other arrows show the basic relationships between the spatial and temporal, i.e., between continuants and occurrents. We shall call this relation *enactment*. Any process, and hence any event, must be the activity or action[9] of some material subject. This is obvious in the case of human activities: walking involves a walker, singing a singer. For processes where intention is not involved, it is often equally clear (e.g., falling requires something that falls, rotating something that rotates), but nature also exhibits many processes where it is harder to pinpoint exactly what is enacting them: thus we might debate whether raining is the activity of a cloud or of the rain (i.e., falling water droplets) itself. Lightning involves the movement of charged particles, the shining of a light involves photons and also a material source that emits them. Hence we must allow 'matter' to range over more than just the tangible materials of everyday life, which for the most part consist of aggregations of atoms. Matter, thus understood, may enact processes (e.g., water enacts the flow process in a river or ocean current), and so may objects (e.g., a boat enacts its motion downstream). Objects can also be said to enact events by virtue of enacting the processes that are constitutive of the events: e.g., a man enacts a journey by enacting the walking, waiting, travelling, etc., that constitute the journey. But we shall not say that matter enacts an event: the reason for this is that since an event is anchored in space and time, the matter involved in the event takes the form of a definite and delimited quantity of matter, which therefore can be understood as an object or collection of objects. For example, the process of raining is enacted by water (i.e., matter), but a rain-storm (which is an event) is

[9]Here these terms are to be understood as neutral with respect to the notion of agency or intention: inanimate objects will be said to enact processes as well as animate ones. We have had considerable difficulty finding a suitable English term for this.

enacted by the particular quantity of water that falls during that storm. This particular quantity of water constitutes an object, which enacts the storm by falling from the clouds to the ground.

3. On change

We now introduce a disturbing factor into the beautifully neat picture presented in Figure 2. This comes about from a consideration of *change*. The key observation here is that change is always ascribed to *something*, which is the subject (or object) of change. One, superficially paradoxical, way of putting this is to say that in order for there to be change, there must be something that does *not* change; otherwise, why would we describe the situation as 'change' rather than just 'difference'? Suppose, for example, we are given two temperature measurements: 20°C at 2 p.m., and 25°C at 4 p.m.; has there been a change? Only if the measurements refer to the same thing: either the same place of measurement, or the same object being measured. We need the *same* place/object to have different temperatures at different times in order for there to have been a change of temperature. What makes the place/object the same at the two times is that it has an identity which remains constant despite some of its accidental properties (such as temperature) changing.[10]

A central notion here is that of *identity*. Despite all the changes he has undergone in the intervening period, John is the same person as he was on his fourth birthday. Little or none of the actual matter of which he is composed has remained the same, and doubtless many of his characteristic patterns of activity have changed too. But something persists, and that is what constitutes John's identity; this is actually something rather mysterious and difficult to account for without begging the question, but whatever it might be, it is this identity whose constancy supplies the justification for singling John out as a whole which can be the object of change. On the process-priority view it can only be a construct from lower-level internal processes, whereas on the object-priority view it is the bearer or host of higher-level external ones. In our view, these are complementary aspects which are both needed for a satisfactory account of what an object is. This will lead in due course to our view of an object as an interface between internal and external processes.

Can events and processes change? Persuasive arguments can be found in the literature (e.g., [8, 28, 18, 41]) that events cannot change, and our view accords with this. The argument is essentially that the event as a whole occupies an interval of time; if in its early stages the event has a certain property which it lacks in its later stages, then it is not the event as a whole which either has the property or lacks it, but rather one part of the event has the property and another part lacks it. Hence in accordance with our criterion above, the event does not change. This is closely connected with the fact that the unitary character of an event is temporal, i.e., events are temporally non-dissective. The properties of an event are properties of the event as a whole. It has temporal parts which may have different properties but this is not change since it is not the same thing that has the different properties.

With processes, though, it is different. We have said that a process is temporally dissective, but in some ways this is misleading. The intended contrast here is with events, whose temporal non-dissectivity can be expressed by saying that none of its temporal parts shares the character of the whole. This would suggest that the dissectivity of a process amounts to the parts of the process sharing the character of the whole. But a process, as we understand it, is not a whole and cannot be said to have parts. Here the analogy with matter is close: just as we do not say that water, or sand, forms a whole with parts, so too with walking or singing. None the less, we can take an episode of walking and say that the walking at t_1 is the *same* walking as the walking at t_2; if the speed or direction of the walking at t_2 is different from what it was at t_1, then we can say that the walking has changed between t_1 and t_2. In the case of objects, change is

[10]This point of view is essentially that expounded by Aristotle in (*Physics*, I.7). Cf. $190^{b}10$: "...that which comes to be is always composite, and there is one thing which comes to be, and another which comes to be this, and the latter is twofold: either the underlying thing, or the thing which is opposed" (translation from [4]). In our temperature example, the 'thing which comes to be' is the new temperature (25°C), the 'underlying thing' is the object or place whose temperature we are talking about, and 'the thing which is opposed' is the old temperature (20°C).

possible because the whole object is present at different times, and can have different properties at those times; this suggests that when we talk of a process changing, if we are to understand anything at all by the 'whole' process, this must be something that is present at each of the times at which the process is in operation. In saying this, we are treating the process as being in at least one respect more like an object than an event, calling into question the neat division into continuants and occurrents we have assumed up to now.

It is interesting to note that this point of view, though not widely adopted, has found expression in the philosophical literature, notably [46, 47],[11] and even where the view is not made explicit we may find pointers in this direction. For example, when Smart [42] states that in a sentence like 'the battle became fiercer' the term 'battle' "takes on some of the properties of a substance word", he is in fact adverting to a use of 'battle' in the sense of a battling *process*, as constrasted with the completed event referred to when we say, for example, that the battle of Hastings took place in 1066 — and in saying that in the former sense the term takes on some of the properties of a substance word, he is coming close to saying that processes are, at least in some respects, to be regarded as continuant-like. A similar view is implicit in [36]. In [18] we find the statement that 'Clearly, in the case of processes we talk of changes without hesitation', but this statement is not supported by any argument relating to the identity of a process.

Let us consider some examples in detail to support our claim. Consider the case of an episode of running by an individual, say Mary. The episode as a whole is, of course, an event, and therefore cannot be said to change. But we are thinking of the process of running as it is realised at different moments during the course of the episode. For example, it may start off slow and then speed up; at a certain moment the direction of the running might be to the north, later on to the west. What is changing here? We can try to ascribe these changes to Mary: we might say that initially Mary was slow and then became fast, that initially she was facing north, and then west; but this is not really a satisfactory solution. To say that Mary is slow at a particular time is meaningless until we specify in what respect she is slow; this could be any of a range of activities such as walking, speaking, thinking, etc, and since one such activity must be specified in ascribing slowness to Mary it is clear that it is the activity rather than Mary herself that is being described as slow. With direction the situation is just as bad, since we must distinguish the direction that Mary is facing from the direction in which she is moving: if she walks backwards or sideways, then these are different.

Our argument for change in processes can be cast in a way that is exactly analogous to change in objects, illustrated above in the case of a change of temperature. We will consider now a swinging pendulum, and specifically consider changes in the amplitude of the swing. Suppose, then, that we have two amplitude measurements: 25° at 2 p.m., and 20° at 4 p.m. We can only say that there has been a change if these measurements relate to the same thing: the same swinging process. And in that case it is natural to say that the swinging has changed: that is, its amplitude has decreased.

There are three counterarguments that might be raised here by someone who was determined to deny that processes can change. We shall consider each in turn, and show how they can be rebutted.

1. First, it might be argued that we are simply mistaken in identifying the swinging process that was measured at 2 p.m. with the swinging process that was measured at 4 p.m.: they are distinct processes, which happen to be connected by a spatio-temporally continuous chain of similar processes, but there is no ground for saying that they are the same.
   One problem with this counterargument is that it would make it difficult to talk about the higher-order rates of change such as acceleration. To define acceleration, we need to compare the velocity at two time points, say t and $t + \delta t$; but this presupposes that it is the same process existing at t and $t + \delta t$. In effect, the concept of acceleration presupposes the notion that processes can change.
   Another weakness of the first counterargument is that there is no obvious reason why it should not apply equally well to change in *objects*. Why can we not just as well deny that it is the same thing that has temperature 20°C at 2 p.m., and temperature 25°C at 4 p.m.? To prevent this we would

[11] See also the following papers by one of the present authors: [12, 13, 15].

need a robust principle of identity for objects that does not apply to processes; but it should be noted that in both cases the grounds for ascribing identity rest on some form of spatio-temporal continuity criterion.

2. The second counterargument takes the opposite line from the first: instead of denying that the process measured at 2 p.m. is the same process as the one measured at 4 p.m., the argument is that the process as a whole extends over an interval that covers both those times and everything in between, and that what is present at 2 p.m. and 4 p.m. are different temporal *parts* of the process. The change consists in the process having temporal parts with different properties, and therefore it does not make sense to say that the process itself changes: in short, because the process *is* a change, it cannot itself undergo change.

   This counterargument misses the mark. What it picks out as something unchanging — namely the change occupying the whole interval from 2 p.m. to 4 p.m. — is indeed unchanging; but that is not what we are referring to here. Rather, we are referring to the process that exists at each moment during the interval. This is indeed a process of change (where 'change' is here a mass noun rather than a count noun), that is, it involves change in some object(s) or matter — namely whatever it is whose properties are changing by virtue of the process — but in addition the process itself has properties which can be different at different times, and this constitutes change in the process itself.
3. Finally, it might be argued that given the two amplitude measurements, then there is indeed change, but what is changing is not any process but simply the pendulum itself.

   This argument depends on our being able to ascribe to the pendulum, rather than to the swinging process, the property with respect to which change occurs. But it is not the pendulum that has an amplitude, but the swinging. And in general, when we want to talk about a process changing, it is because the property that is changing value belongs to the process itself and not to that which enacts it.

   To illustrate this with another example, consider the flow of water through a pipe. Suppose we have two measurements: 2 litre/second at t_1 and 1 litres/second at t_2. So the flow has slowed down. But there is no sample of water that has undergone that deceleration. The water passing through the pipe at t_1 was going at a rate of 2 litres/second, and the *different* water passing through the pipe at t_2 was going at 1 litre/second, with intermediate samples at intermediate speeds. So no one sample of water has slowed down. Nor can we say that the pipe has slowed down, since the pipe does not move or flow at all. So the third counter-argument only works, if at all, for those processes which involve the same material participants throughout their lifetime. In that case one might sometimes plausibly ascribe the changes in the process to the participants instead — but not always, as the example of a person's motion changing direction without the person herself changing direction shows.

How should we react to the conclusion that processes can change? On the face of it this is problematic, since processes are generally included in the category of occurrents, and it is generally regarded as a feature of occurrents that they do *not* change. If we use the terms 'mutable' and 'immutable' to refer to ontological categories that respectively do and do not admit change, then we seem to have the following situation:

	objects	processes	events
temporally extended	No	Yes	Yes
mutable	Yes	Yes	No

Our classification now hangs on whether we are prepared to admit the possibility of temporally extended things being mutable. We require processes to be both: temporally extended, in order to fill the role of being the material from which events are composed, and mutable for the reasons argued in detail above. Unfortunately, this conjunction of characteristics seems to contradict the received wisdom that temporally extended entities cannot be said to undergo change. The justification for this, as already noted, is that in

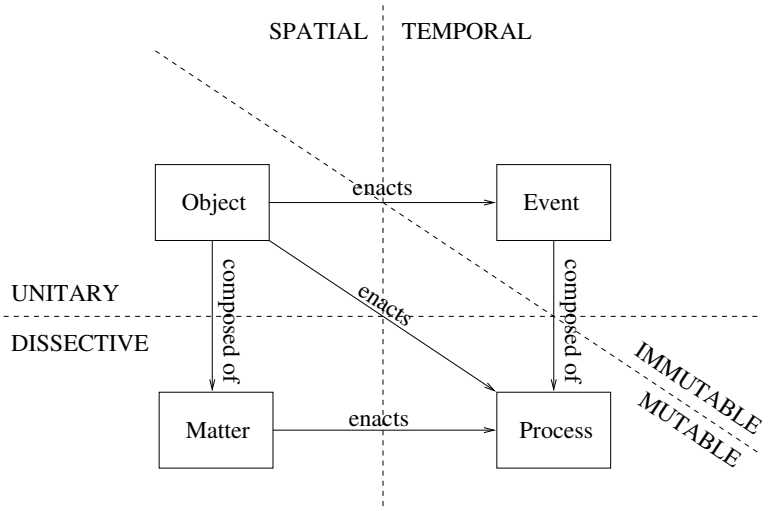

Fig. 3. Four categories of physical phenomena

order for something to change, the very same entity must possess different properties at different times, whereas with a temporally extended entity what exists at different times is distinct parts.[12]

We have here a puzzle not unlike that confronting physicists in the early years of the twentieth century, when they needed to reconcile the apparently particulate nature of light with the fact that in many circumstances it manifests interference phenomena characteristic of a wave. The solution we shall propose to our puzzle is very loosely analogous to the idea of a wave packet used in quantum mechanics to resolve the wave-particle dilemma. This solution is presented in the next section.

Meanwhile, we should note that the distinction between mutable and immutable is not aligned with either the vertical or horizontal divisions in Figure 2 — this is why we referred to change as a disturbing influence on the neat symmetry of our diagram. A revised diagram, incorporating the element of mutability, is shown in Figure 3.

4. Events and Processes

As already noted, an important contrast between events and processes is that the former possess a unitary (or non-dissective) character that the latter lack. This goes along with the idea that whereas an event is a distinct individual episode with a definite beginning and end which define its character as a whole, a process seems to occupy time in a more indefinite, open-ended way.

To illustrate this, consider once again our example of Mary's running. Suppose she begins running at 10 a.m. and stops running at 11 a.m., having been running all the time during the intervening hour. In this situation we have an event — let us call it E — which consists of Mary's starting to run at 10 a.m., her continuing to run for the next hour, and her stopping running at 11 a.m. The event E takes place on the interval [10.00,11.00]; it does not take place on any proper subinterval of that interval, since no such subinterval can include both the starting and the stopping which are essential parts of the event. Thus event E is a single whole, and as such we can meaningfully speak of it as having parts: for example, if Mary was in fact running round a circuit, of which she completed ten laps, then we can talk about the first lap, the second lap, and so on, as events in their own right, each forming a part of the complete running event E. Even when there is no 'natural' (or *bona fide*) way of dividing the event into parts, we can still do so by arbitrary fiat, e.g., the first ten minutes, the second ten minutes, and so on; whether these should really

[12] Compare Mellor: 'Change may ... be defined as a variation in a real property of something, provided the variation does not reduce to a difference between different parts of it' [28, p.9].

count as parts may depend on one's particular metaphysical predilections, in exactly the same way as with *bona fide* vs *fiat* parts of ordinary material objects [43, 45]

We can contrast this event E with the process which constitutes it: Mary's running. Let us call this process P. In referring to the process of running one is making no presuppositions about its starting and stopping, and in fact we regard the process as being *open-ended* in the sense that if someone is running then, at least in principle, they can continue running indefinitely. This means that considered as a process, the running does not constitute a unitary whole, and as such one cannot meaningfully speak of it as having parts either. Rather, we should say that Mary's running exists (or is in operation) throughout the interval (10.00,11.00), and this implies that we can look at any subinterval of the interval and find Mary running — the process P itself, and not any 'part' of it — going on over that subinterval. The event E is identified with the 'chunk' of process P that is delineated by the boundary events of starting and stopping. The one-hour duration of E reflects the fact that P was in operation over a one-hour period between the starting event and the stopping event, but whereas the identity of E is tied to that hour-long interval, giving the event its unitary character, the identity of the process is distributed over the hour in such a way that whatever time (be it a subinterval or an instant) we select from the hour, we will find the same process in operation at that time. In this way processes are more like continuants than events: the analogous case for ordinary objects is that whatever instant we choose from a person's lifetime, we will find the same person alive at that instant. The difference from ordinary objects is that we have to allow processes to have temporal extension as well in order for them to be the raw material for events.

In general, if some process P becomes active at t_1 and becomes inactive at t_2, and is continuously active during (t_1, t_2), then we can say that there is an event E which occurs on $[t_1, t_2]$ precisely by virtue of these facts. This event is what we mean by a 'chunk' of process P. The notion of a process being active is of course intimately tied to the identity of the process: being active for a process is the analogue of being present (at a time) for an object. In particular cases there may be uncertainty about whether a process active at t_1 is the same as a process active at t_2, and in that case there may be a corresponding uncertainty as to what process chunks there are; but this possibility does not invalidate the general characterisation of a process chunk.

Now suppose that at 10.20 Mary was running at 15 km/hr, whereas at 10.40 she was running at 20 km/hr. Then the running — process P — has become faster. The change in the running reflects the fact that at one instant it has one property and at a later instant it — i.e., the very same running — has another property. Thus process P exists or is in operation at instants as well as over intervals: if this were not the case we could not meaningfully ascribe a speed to it at an instant, which would mean that we could not, for instance, speak of the speed as continuously increasing over an interval.[13]

We are very used to regarding speed as a property of processes that naturally attaches to instants rather than intervals, but we must not forget that the notion of instantaneous speed is an idealisation derived by a limiting process from the more readily grounded notion of average speed over an interval. We can consider a moving object at different temporal granularities. If the limit of our temporal resolution is, say, 0.1s, then we can think of the motion process as occupying a little 'temporal window' of duration 0.1s, which as it were moves forward through time. On any such window we can assign a speed to the motion (namely, the distance travelled during the window, divided by 0.1s), and if the speed is different in different windows then we can say that the speed of the motion is changing — and hence that the motion process itself is changing with respect to the attribute of speed.

If the motion is smooth and homogeneous in character, then in principle there is no limit to how short a window of time we can consider the process as occupying. But many, perhaps most or even all, processes do not appear to be homogeneous. Our example of running is a case in point: here the process is cyclical, with the legs alternately exerting a force on the ground to propel the body forward. It has often been claimed (e.g.,by [7]) that such processes cannot readily be said to take place over intervals shorter than

[13] Of course, to *compute* the speed at time t we look at the distance/time ratio over a nest of intervals of the form $(t - \delta t, t)$, but arguably the speed itself is that property of the process exactly at t which explains why the intervals $(t - \delta t, t)$ exhibit the particular distance/time ratios that they do.

their intrinsic granularity. Our temporal window idea is very well adapted to the description of change in such processes: the running process, let us say, occupies a window of time during which the minimal amount of change happens that is required to characterise what is going on as running. This window moves forward continuously through time (we can think of it as something like the 'specious present' of the psychologists [26], although its duration is derived from the process itself rather than, for example, the conscious experience of an observer), and the contents of the window may be different at different times, so that the running process itself changes.

On this picture, the running process can still be regarded as homogeneous in the sense that, in any moment during an interval over which running takes place, there is running going on. What *makes* it running is, of course, the fact that any such moment falls within a temporal window containing a sequence of movements characteristic of the running process. It is no argument against homogeneity, thus understood, that a precisely similar moment (if such be possible), occurring in an overall context from which the sequence characteristic of running is absent, would not count as a moment of running.

The length of the window should not be considered fixed: since it is determined by the intrinsic granularity of the process, if that granularity changes, so will the length of the window. Moreover, we cannot usually assign sharp boundaries to the window. There may be different candidates for what counts as the 'intrinsic granularity' of the process, and this may vary depending on what attribute of the process we are considering. To take a very different example from that of running, consider the case of a lecture which becomes more interesting as it progresses. In line with our previous remarks, we know that what is becoming more interesting is not the lecture considered as a completed event, but rather the lecturing process that is constitutive of that event. Now it is arguable that it makes no sense to ascribe a level of interestingness to this process at an instant or over a very short interval: it has to do with the content that is being conveyed, and any piece of content takes some time to get across. Thus the interestingness is being ascribed to the lecturing within a window that is long enough for some content to be conveyed. This does not prevent us from registering a change in the level of interestingness, If it is 'rather boring' for the first five minutes and 'fascinating' for the last five minutes, there has clearly been a change, and this is because the lecturing process in operation over the first five minutes is the same lecturing process that was in operation over the last five; this same process has the property 'rather boring' over the former interval and 'fascinating' over the latter. The temporal window under consideration need not be as long as five minutes: it just needs to be long enough for it to be possible ascribe a level of interest to the lecture during that window. Thus the mutability of processes comes from the fact that a process can be said to have different properties during different time windows during the period that it is in operation.

Our solution to the puzzle about events may be summed up as follows. The puzzle is how to reconcile the fact that processes can change with the fact that they seem to be extended in time. The solution is, in effect, that a process's temporal extension itself moves forward through time. The extendedness of the process defines what we have called its temporal window, which must be large enough to accommodate the intrinsic inhomogeneities of the process (such as the movement of the legs when running); the mutability arises from the fact that while the process is in operation, its temporal window moves forward and the attributes of the process can assume different values at different positions of the window.

Having introduced the notion of a temporal window we must partially retract our earlier statement that a process, as we understand it, cannot be said to have temporal parts. In the limiting case of a completely homogeneous process for which the duration of the window shrinks to zero, this statement remains correct: the process is present at each instant that it is in operation, and it is the same process that is present at different instants during the course of its life. In this case the process can be regarded as a true continuant, and indeed should probably be classified rather as a *state*. For any other process, the temporal window means that the process can be regarded as extended in time, and different parts of the window may be said to contain different temporal parts of the process. For example, when someone is running, the window will contain both forward motion of the left leg and forward motion of the right leg, and these can be regarded as parts of the process: thus we say that moving the left leg forward and moving the right leg forward are both temporal parts of running. But if we consider two different temporal windows, what they contain is

the same process at different stages in its lifetime, not different temporal parts of some entity (such as the *life* of the process) whose temporal extension encompasses both windows.

So far we have focused on a case where it is not hard to see why people might be confused about the difference between a process and an event. The difference between the event which is Mary's hour-long run between 10 a.m. and 11 a.m., and the process which is the running Mary was doing between 10 a.m. and 11 a.m., may seem at first sight to be rather subtle — properly considered, however, the difference is considerable. Following the analogy with material objects and matter, it is similar to the difference between a gold ring and the gold that it is made of. In the case of the ring, the particular annular form assumed by the gold — in other words the configuration of its surface — forms part of the ring but not of the gold. Similarly, the starting and stopping form part of the event (the run) but not of the process (the running). This event has minimal structure; but many events are composites consisting of many distinct sub-events constituted by different processes (think, for example, of a battle or a picnic), and there it is easier to appreciate the difference.

Mary's starting to run, and her stopping running, are both events in their own right, but of a rather different character from the running as a whole. Like other events, they are dependent on processes, but in a different way: not through being constituted by processes, but by forming a boundary of a chunk of process. In idealisation at least, these events are instantaneous. Thus some events occur over intervals and others at instants; but in fact it is not essential to an event that there should be a precisely determined time of occurrence, whether an interval or an instant. It can even seem that some events have the open-endedness that we regard as a characteristic feature of processes. An example is the event of *departure*. This has a definite beginning, when, say, the train starts moving, but the ending is somewhat indeterminate: should we say the train has departed when it is out of the station, when it has got up to speed, or when it has gone out of sight? This is underdetermined by our concept of departure. In that sense, departure seems to be open-ended. Likewise, arrival is open at the beginning. We can even have an event which is open, in this sense, at both ends: passing by something (e.g., one car overtaking another). Despite this apparent open-endedness these are all clearly events rather than processes. The key requirement is that an event should have one or more *anchoring points*; normally there will be two anchoring points, a beginning and an end, but a departure event is only anchored at its beginning (it has no well-defined end) and an arrival event at its end (no beginning); an overtaking event is anchored in the middle.[14]

The intimate dependence of events on processes has often led people to think that events and processes must be ontologically rather similar. The terminology used is notoriously slippery, and the term 'process' has often been used in a broader sense than we have been using here, even to the extent of including all events within its scope.[15] Here we would like to emphasise the very *different* nature of processes and events, as illustrated by our discussions above. In everyday usage the contrast between events and process is often drawn quite sharply, particularly in order to draw attention to the difference between some 'one-off' happening and a protracted on-going process. An internet search for the phrase 'not an event but a process' shows that this contrast is at home across a wide range of different subject areas: thus, to cite only a few, we find that 'marketing is not an event, but a process'[16], that a security audit is 'not an event but a process'[17], that 'risk management ... is not an event but a process'[18], that 'fertilization ... is not an event,

[14]The spatial analogue of this notion of anchoring can be used similarly to underpin the idea of vague or indeterminate spatial location; see [16, 24, 14]. A mountain, for example, is anchored at its summit, the outer boundaries being underdetermined by the concept 'mountain'. This corresponds, in Varzi's supervaluationist account [50], to the fact that any precisification of a given mountain name (e.g., 'Everest') must include the summit. It should be emphasised that in all these cases, the indeterminacy belongs to the concepts and words we use, not to the reality we are using them to describe.

[15]See [15] for an extended discussion of the variant classifications that have been proposed.

[16]http://www.gmarketing.com/articles/read/13/ (accessed 14/8/08).

[17]http://www.securityfocus.com/infocus/1697 (accessed 14/8/08)

[18]http://ahds.ac.uk/creating/information-papers/risk-management/index.htm (accessed 14/8/08)

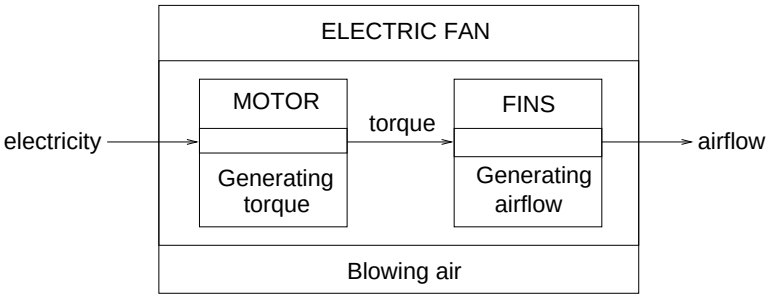

Fig. 4. An electric fan as a device.

but a process' [19], that 'for Augustine justification is not an event but a process'[20], and that 'revelation is not an event but a process' [21].

Perhaps the clearest way of highlighting the distinction between processes and events is to revert to the close analogy that this distinction bears to that between matter and objects [1, 33]. We mentioned above the idea of an event being ('as it were') composed of processes (or better, 'process-stuff'), and this is much the same as the way physical objects are composed of matter. In other cases, the relation is not so much composition as determination. Objects such as edges, corners, and surfaces, for example, are all determined by the spatial distribution of matter even though they are not themselves composed of matter.[22] Likewise, 'bounding' events such as arrivals and departures, or more generally startings and stoppings, are primarily defined in terms of a temporal distribution of process.

5. Internal and External Processes

In Figure 3 the enactment relation is prominent. In this section we examine this relation between continuants and occurrents in more detail. It will lead us to formulate a distinction between internal and external processes which will play an important part in our characterisation of an object in terms of processes it is involved in. A useful point of entry into this topic is to consider an ontology of devices, in which a device is seen as a 'black box' with input and output ports [29, 38]. The function of the device is essentially the generation of output from input. What happens in between, mediating between input and output, is purely internal to the device, and is not visible from outside so long as the device is regarded as a black box. But open up the box, and you will see other black boxes (the immediate components of the device), each with its own inputs and outputs in terms of which its function can be specified — and within those, other black boxes still, until we reach devices that we wish to regard as 'primitive'.

If an electric fan, for example, is considered as a black box, then its function does not include the rotation of the fins, even though, because they are conspicuously visible from outside, it is easy to make the mistake of thinking that it does. From an external point of view, the function of an electric fan is the conversion of electrical energy (input) into airflow (output). The fan's function should be described as blowing air rather than rotating its fins. The rotation of the fins is part of the internal behaviour by which the airflow is brought about; another part is the electric motor, which converts electrical energy to torque (see Figure 4).[23]

[19] http://www3.interscience.wiley.com/cgi-bin/fulltext/120015036/PDFSTART (accessed 14/8/08)

[20] http://pontifications.wordpress.com/2008/01/07/augustine-on-justification/ (accessed 14/8/08)

[21] http://www.js.emory.edu/BLUMENTHAL/Green.html (accessed 14/8/08)

[22] But see, for the case of surfaces, the discussion in [14].

[23] In this figure and other figures appearing later, each object will be given a name, a process that it enacts, an input, and an output. The enacted process should be understood as the means by which input is converted to output. Where one object is embedded within another, it is to be understood that the former is part of the latter — in certain cases, the two are actually equal (at least as regards their physical constitution), but are being considered under different descriptions (e.g., the human being and its body in Figure 6).

In this figure we can identify, for each object (i.e., the electric fan, the motor, and the fins), four elements: the object itself, the process we are considering it to enact, and the input and output to this process. The key notion here is the complementarity between object and process: to say that O enacts P is precisely to say that P is an external process of O. The relations 'enacts' and 'is an external process of' are converses. The object enacts its external processes, and the processes are what the object does, conceptualised in general terms as the generation of output from input. Input and output here do not necessarily have to be material stuff or things: we shall see examples where they are states of affairs.

What we have said about internal and external processes applies to events too; this is because events are composed of or defined in terms of processes. If the fins become disconnected from the motor, they will stop rotating; this is an internal event. A consequence of this is that the airflow stops; this is an external event. The function of some devices is intrinsically event-like (discrete) rather than process-like (dissective). An example is an electrical switch. The input here is the state of the current (flowing or not flowing); the output is the opposite state. It is by nature impossible to maintain a continuous transfer between input and output: it is effectively instantaneous.

In the fan, the rotating fins are visible from the outside — so there is something 'external' about the rotation process. In fact, what is internal or external depends on a point of view: what is regarded as input and output. Here a pertinent distinction is between intentional or designed function and actual (whether intentional or unintentional) function. The operation of any device can have unintended side-effects in addition to its intended functionality. It is no part of the purpose of a fan that it should display visible rotation. (Just as it is no part of the purpose of a clock that it should produce audible ticks — which is why many perfectly functional clocks do not do this!)

One might compare here a child's toy windmill and a real windmill. The function of the real windmill is to convert wind energy into useful work (e.g., turning a grindstone to grind corn, or turning a dynamo to generate electricity). The rotation of the vanes is an internal process by which it achieves this function; it is not what the windmill is *for*, even though it is the most visible aspect of its behaviour. A child's windmill also takes wind energy as input; it converts it to visible rotation of the vanes, and then ... nothing else! What was an internal process in the real windmill has become the desired output in the toy windmill, its function. In fact it is typical of toys and some kinds of models that their function is to reproduce the visible appearance of what they model, regardless of the role this appearance plays in its functioning.

In Figure 5 we show diagrams for our windmill examples, in similar style to Figure 4.

Invoking the device ontology in this way might be objected to on the grounds that very many of the objects we ascribe processes and events to are *not* devices in any normal sense of the word. Nonetheless, we would claim that the 'internal/external' view is of much wider application than just the device ontology. Although talk of functions seems to invoke the notion of intention — of the designer or user of a device — we can usefully extend the distinction between internal and external processes to non-designed objects such as biological organisms, and even to inanimate things such as rivers, waterfalls, and ocean currents, as we shall discuss below. We must stress here that we are not seeking to extend the notion of 'function', about which there is in any case already considerable dispute amongst philosophers and others, in particular with regard to the relationship between designed function (arising from the intentions of a designer) and 'natural' function (e.g., biological functions), where the appearance of design has arisen through non-intentional processes (cf. [2]). Rather, we are here viewing function as an instance of a wider notion of 'what something typically or by nature does', the most general definition of which is precisely in terms of the external processes enacted. The point of invoking the device ontology here is that it throws into sharp relief the distinction between internal and external processes and the relationship between them. When we speak of devices, we can unproblematically speak of their associated functions; but where in what follows we speak of 'devices' (in quotation marks), then we will not use the word 'function', instead referring to what the 'device' does or how it behaves.

The device ontology sees the world in terms of nested structures, as we have already described in relation to the fan. A more complicated example is a motor car. On a high-level view the function of the car is simply to *move* (typically along roads); the unit that moves is actually the car together with its driver, passengers, and luggage. At a rather lower level, the driver considers himself to be external to the car.

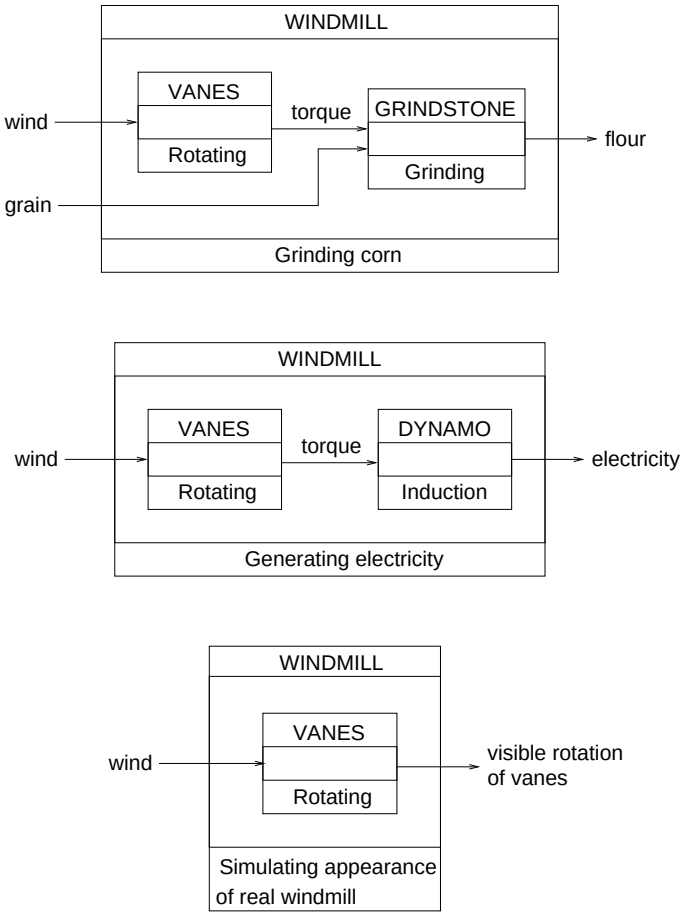

Fig. 5. Three windmills as devices. From top to bottom: windmill for grinding corn, windmill for generating electricity, toy windmill

On this view the car is a device whose input is the manipulation of the controls by the driver and whose output is the movement of the car along with its contents. Note that this input/output relation is common to almost any mode of tranport. That the car achieves its motion by burning petrol in an internal combustion engine and transmitting the power thereby generated to the wheels is *internal*. The components (engine, transmission, wheels) are individual devices each with its own function.

A similar description applies to the human body. Consider walking: at the level of the whole human, this is a behaviour whose overt effect is locomotion along a path. It is achieved by the coordinated behaviours of various subsystems that are internal to the human, most visibly the alternating backward and forward motion of the legs. Note that this backward and forward motion is not itself walking: thus we can say that the human walks by using his legs, but the legs themselves do not walk. This is obvious in the case of other parts of the body: the heart, lungs and brain do not walk either, although their activities are an essential part of the means by which the human is able to walk.[24] The brain plays a particularly significant role here, since if we try to identify the *input* to the walking, it would seem that in general terms the best one can say is that it is a perceived 'need to move', and the brain is instrumental in converting this into the nerve impulses which set the legs in motion; the legs thereby exert a forward force which imparts an

[24]Similarly, eyes do not see, ears do not hear, and brains do not think; rather, humans see with their eyes, hear with their ears, and think with their brains.

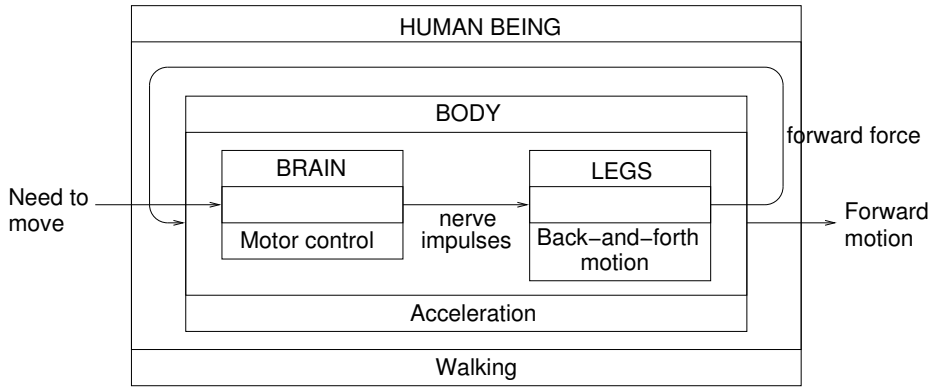

Fig. 6. A walking human as a device

acceleration to the whole body which results in its moving forward.[25] Using our diagrammatical notation for devices, the picture is somewhat as in Figure 6 (where the role of heart, lungs, etc has been ignored for simplicity).

The matter involved in a process may already exhibit a sufficient degree of coherence — *sticking together* — to allow us to identify it as a material object independently of the process. An example is a stone rolling downhill: rolling downhill is a process, but what enacts this process can be recognised as an enduring entity, the stone itself. This is not always the case. The tumbling of water over a precipice is a process; at different times, this process is being enacted by different samples of water. So there is no one sample of matter that enacts it. There is indeed an enduring pattern here which we recognise as an object — a waterfall — but this does not itself enact the tumbling process: the water falls but the waterfall does not fall.[26]

What does the waterfall do? How we answer this depends on exactly what we take the waterfall to be. In fact the most immediately obvious choices for what the waterfall is and what it does are mutually inconsistent: we mean here on the one hand the view of the waterfall as a particular persistent configuration of falling water; and on the other hand, the view that what a waterfall does is to transfer water from a higher elevation to a lower. The problem with this is that in modelling a waterfall as a device for transferring water, as suggested by the latter view, we are treating a waterfall as a kind of conduit: but a conduit does not include what it conducts as a constituent (e.g., water is not part of a water-pipe, or gas of a gas-pipe). The conduit view of a waterfall is indeed possible, but what it refers to is not the falling water itself (that is, what we normally mean by the waterfall) but rather the interruption to the continuity of the river bed resulting from the presence of the rocky precipice. This is what effects the transfer of water from higher to lower elevation by maintaining a spatial configuration and solidity which allows the water to fall.

So what does the waterfall, considered as a persistent pattern of falling water, do? In the short term, nothing very much: one can sit watching the waterfall for hours, day after day, and very little changes. But over a period of many years the waterfall may migrate upstream as it carves its way into the rock, and it is this long-term migration that supplies the true answer to our question. The process which the waterfall enacts is this larger-scale movement upstream. The waterfall may be regarded in this way as a 'device' for moving the rocky precipice upstream along the river channel, through the agency of the falling of water, which constitutes its primary internal process.

These two views of the waterfall — as a rocky precipice which facilitates the transfer of water from higher to lower level, and as a persistent configuration of tumbling water which over a long period pushes the rocky precipice upstream — are illustrated in Figure 7. In each of these cases we only show one

[25] Of course we are aware that the detailed mechanics of walking are extremely complicated; here we present merely a very schematic outline.

[26] Of course, one might doubt whether it is correct to regard a waterfall as an object in the first place; this question will be discussed at greater length in §7.

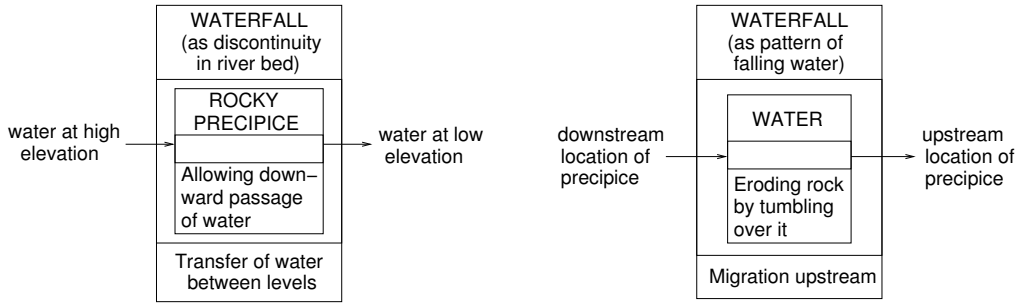

Fig. 7. Two ways of viewing a waterfall as a device

embedded subsystem, having the same input and output as the system as a whole. The subsystem might thus look similar to the system as a whole, but none the less they are conceptually distinct and enact different processes: we can say, for example, that the transfer of water between levels that is enacted by the discontinuity in the river bed (and is characterised in terms of the difference between the earlier and later positions of the water) is effected *by* the downward passage of the water, the enablement of which is enacted by the precipice (and which is characterised in terms of actual downward motion of the water).

Similar remarks apply to rivers, where again we can distinguish the river-channel, as a configuration in the land which acts as a conduit for water from source to mouth but does not contain water as a part, and the river itself (which one can swim in or float on) which consists of water flowing through that channel. The former conception is presumably what we are referring to when we point to a dried up river bed and say 'That is the river so-and-so', but the latter conception corresponds better to our normal everyday idea of a river. Under this conception, the flow of water is an internal process of the river; its external processes are the gradual changes of form and position that the river may undergo over many years. In the case of an ocean current there is no analogue to the the conduit view, since there is no channel that the water runs through. Like the river, the current has the flow of water as its internal process, and longer-term changes of form and position as its external process.

What enacts the internal process of the river or ocean current is, of course, water: not a specific quantity of water, but different specific quantities of water at different times. Water is not a kind of object but a kind of matter. (Put differently, 'water' is a mass noun, not a count noun.) Of course, at a sufficiently fine granularity, water resolves itself into a vast number of discrete components, the individual H_2O molecules. At this level, the flow of water can be seen as a collective phenomenon; but it is a collective of sub-microscopic components. Other collectives, for example crowds of people or flocks of animals or birds, have components that are discernible as individuals at a scale that we humans can immediately apprehend.

Consider, for example, a moving crowd. As a whole it moves in a certain direction: this is an external process, enacted by the crowd. Individual members of the crowd can have a wide variety of different movements, but from the point of view of the crowd these are all internal processes. Indeed, membership of the crowd can be fluid: individuals can come and go without compromising the identity of the crowd itself; this kind of systematic replacement is characteristic of some forms of collective phenomenon, e.g., queues — and note the similarity of a queue to a waterfall in this respect.

The external processes enacted by a collective 'emerge from' the composition of the internal processes. Not all the processes that are spatially interior to an object contribute in any way to the external processes of the object — a clear example of this would be a train journey; in the train there are many passengers engaged in all manner of activities, but none of these contribute to the motion of the train; only the running of the engine does that. An interesting question is: When does a collection of individuals constitute a higher-order entity? An arbitrary collection of people (e.g., all those listed on a certain page of a telephone directory) will not usually exhibit the kind of coherence or coordination in the behaviours of the individual components that would lead us to regard it as a collective entity in its own right. We shall address this issue in the section on objects below. For now, just note that to analyse a collective in terms of its constituents

does not mean to analyse it *away*: the collective is none the less real for being analysable into smaller-scale constituents.[27]

6. Granularity and substructures

Let us pursue the thought with which we ended the previous section. The functioning of a device can be explicated in terms of the functionings of all its subdevices and their interaction. More generally we can say that a whole is dependent on its parts. One might be tempted to conclude from this that such a whole is a kind of 'illusion' or 'virtual object'. But then parts of the whole are wholes in their own right, so by the same line of reasoning must be regarded as similarly 'illusory'. Unless the recursion bottoms out at a level of irreducibly atomic objects, all of reality seems to dissolve into endless layers of illusion. This cannot be a worthwhile picture of the world. If we are to nip this potential regress in the bud, we must accept that at least some objects are real; but it seems arbitrary to pinpoint the onset of reality at any particular granularity level. This is a powerful argument for accepting objects at all levels of granularity as equally real.[28]

Having agreed to take parts and wholes seriously, we then have the interesting question of what makes one object a part of another object. It is well known that 'part' is not a single concept, and there exists a considerable literature exploring its many ramifications [41, 54, 3, 52, 25, 30]. One of the most important kinds of part is *component*. We believe that our analysis of processes can help to throw some light on the nature of this concept. In the case of devices and organisms, this may be identified with the notion of *functional part*, for which the important thing is how the function of the part is related to the function of the whole; this defines the *role* of the part within the whole. A key idea is that a functional part is something that enacts an internal process of the whole. The internal processes of a human being include, as we have said, such things as breathing, digestion, blood circulation. They are internal processes because they support the external processes of the human — speaking, eating, walking, and more generally *living*. The internal processes are enacted by the lungs, stomach, heart, blood vessels, and the blood itself, and this is why all these things count as functional parts of the body. The human itself can be seen as an interface between the internal and external processes: the human is *sustained* by the former and *enacts* the latter, and the stability of the human as a persistent object is maintained by the constant ongoing relationship between the internal processes and the external processes they support. In the more general setting in which we regard arbitrary objects as 'devices', we can no longer speak of functional parts, but we can still refer to components to denote those parts whose activities contribute in an essential way to the activity of the whole.

Precisely what objects we see in the world is a function of which processes we choose to view as important, and the roles we assign to these processes as internal or external. Consider the case of a train. Properly considered, there are many different views we can take of what a train is and what it does. Perhaps the most salient point of view is to regard the train as a *vehicle* whose function is the conveyance of passengers or goods.[29] Considered as a device, the train as vehicle takes as input fuel and the control exerted by the driver (or a computer in the case of driverless trains) and delivers as output the forward motion of passengers along the route defined by the track. This view can be expressed by the diagram in Figure 8(a).

The function by which the train carries its passengers can be divided into two distinct components: containment and motion. The train *contains* its passengers: what this means is that it constrains their

[27] See [11] for a discussion of some problems relating to collective phenomena, and [55] for an attempt at classifying such phenomena.

[28] Of course this does not mean that there cannot be 'illusory' objects as well, e.g., artefacts of a particular point of view, like a constellation which consists of stars which are in reality at very different distances from us but happen to lie along close together lines of sight as seen from Earth. Even a constellation has a certain minimal reality by virtue of giving rise to a characteristic gestalt perception in human observers; it is mostly 'illusory', however, in that the expectation that the gestalt perception is caused by something with additional causal powers proves to be unfounded.

[29] For simplicity we shall assume passengers in what follows.

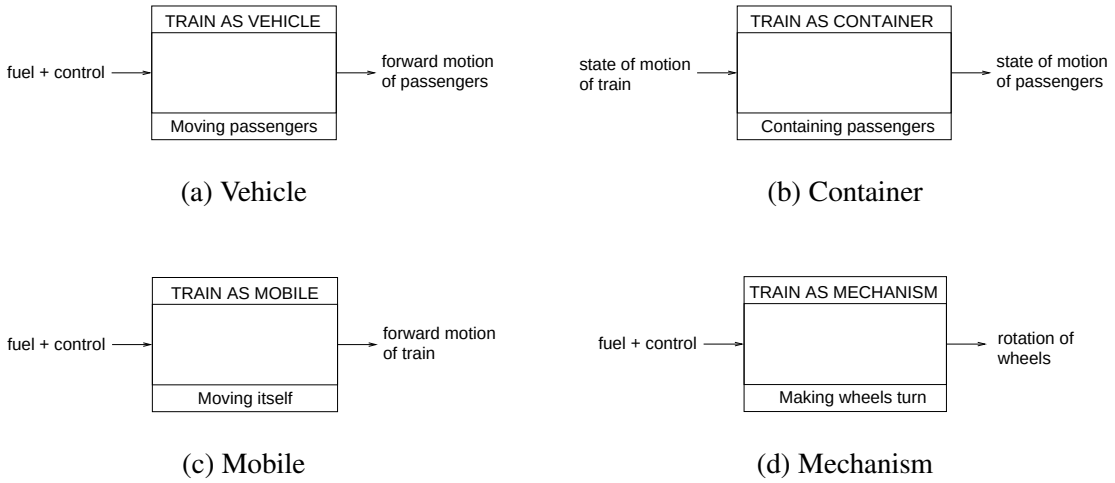

Fig. 8. Four views of a train

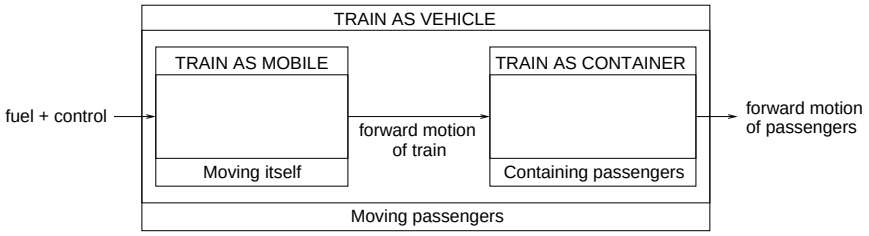

Fig. 9. Expanded view of the train as a vehicle.

positions in such a way that the passengers must share the state of motion of the train — if the train is stationary, then so are the passengers,[30] and if the train moves, the passengers move with it. But the train also *moves*, specifically in a linear fashion along the track. The combination of motion of the train and containment of the passengers by the train leads to the desired result, motion of the passengers.

This discussion has brought to light two additional views of the train. First, there is the train as a *container*. Its function is to convert the state of motion of the train to a state of motion of its contents, as illustrated in Figure 8(b), where the contents are represented as passengers (although more generally they might be passengers, goods, or a mixture of these). Moreover, the state of motion of the train and its contents may be replaced by the particular kind of motion (e.g., forwards, backwards, oscillatory) as required in a particular case. Second, there is the train as a moving, or more exactly *mobile*, object. Its function is to convert fuel and control into forward linear motion, as shown in Figure 8(c). Our top-level view of the train as a vehicle is obtained by combining the train-as-mobile and train-as-container, as shown in Figure 9.

Now of course the forward motion of the train could not happen without the rails on which the train runs (if the train is suspended in the air, the wheels may rotate, but this rotation will not be converted into forward motion). We can divide the process of converting fuel into forward linear motion of the train into two parts: first, converting fuel into rotation of the wheels, and second, converting rotation of the wheels into forward linear motion of the train. The first of these parts is the function of our fourth view of the train, the train as *mechanism*, as portrayed in Figure 8(d). The second part depends on the interaction of the wheels with the rails. The rails are not part of the train, but this interaction is an essential part of the functioning of the train as mobile (though not of the train as mechanism). In term of the theory of roles

[30] Here we are thinking at a level of granularity which regards motion between stations as significant but the motions made by passengers fidgeting in their seats or walking up the corridor as not.

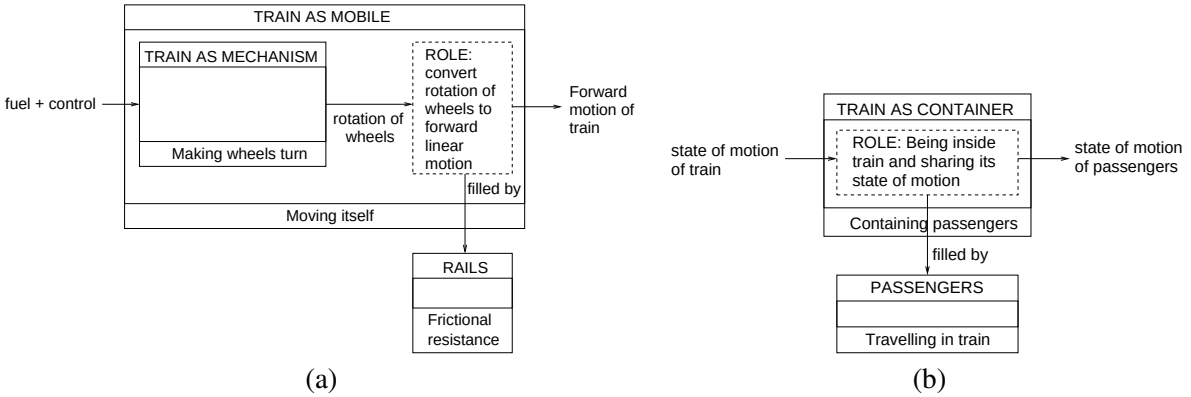

Fig. 10. Expanded views of train-as-mobile and train-as-container, showing roles and role-fillers.

[30], we may say that the train as mobile includes a *role* for something which converts rotation of wheels to forward linear motion. This role may be regarded as internal to the train as mobile, but what fills the role — the rails — is external. The rails fulfil the function defined by the role by exerting frictional resistence against the wheels as they turn.[31] This can be shown by expanding our diagram for the train as mobile as shown in Figure 10(a). Here the role is shown as an internal place-holder within the train object, to be filled by the object that plays the role; in this case the input and output are not shown explicitly for the role-player, since they can be understood to be the same as the input and output shown for the role itself. The same convention applies to other diagrams involving roles.

Another missing ingredient from our account is the passengers themselves, which are referred to, but not included, in the train-as-container view. Here again we can say that the notion of a container must intrinsically include a role for what is contained. In the case of a passenger train viewed as a container for passengers, this role is, of course, filled by the passengers, by virtue of their exercising the function of travelling in the train. This is shown in the expanded view of the train as container illustrated in Figure 10(b).

Putting all these pictures together, we obtain a diagrammatic analysis of the various views of the train and their interrelationships, as shown in Figure 11. Of course, even this is far from complete. Every component can be analysed further (for example, in the train as mechanism, we have not identified the engine and wheels as separate components: the engine generates torque from fuel, the wheels convert the torque into rotation), but we have shown enough to illustrate some important general principles.

The key notion is that an object, considered from a particular point of view, is characterised in terms of the processes it enacts. These are what we call the *external processes* or *behaviour* of the object. This behaviour arises as a result of various internal processes which causally contribute to it. Here 'internal' should be understood as referring to how these contributory processes are related to the overall external behaviour. They are not necessarily physically internal in the sense of taking place inside the object we are considering: for example, the interaction between the wheels and the rails do not take place inside the train, but they are internal processes of the train because of the role they play in generating the train's external behaviour.

Each internal process of the object is specified by a *role* which forms part of the definition of the object. We regard this as internal to the object, and it may therefore be called an *internal role* to distinguish it from any external roles which the object itself may play in relation to some larger system of which it forms a part. An internal role is filled by a *role-holder* which enacts the process specified by the role. The role-holder may be internal or external to the object itself. If it is internal, then it may be regarded as a *component* of the object, whereas if it is external we shall call it an *auxiliary object*, but in either case

[31] Here we are assuming that the rails are fixed to the ground. A more complete model would build in this assumption by including the ground and the interaction between the rails and the ground; for simplicity we are ignoring this.

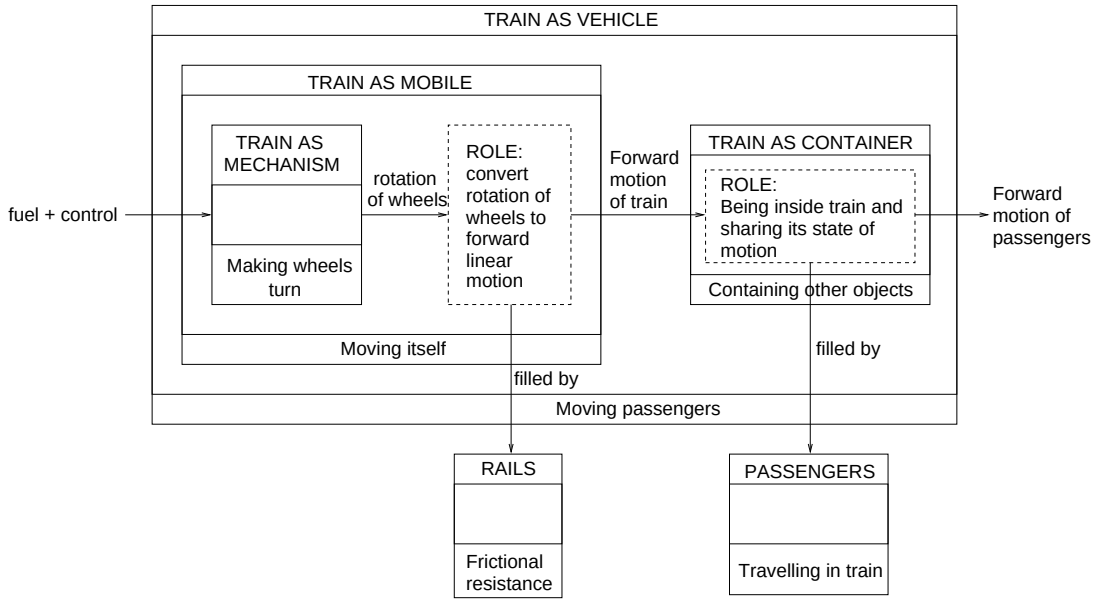

Fig. 11. How the four views of the train are interrelated

we may say that the role-holder *participates* in the external process of the main object.[32] An example of the difference between internal and external role-holders is provided by the power sources of a car and a bicycle, considered as mobiles: the role in question is to generate power, and in the case of the car this is filled by the engine, which being internal to the car is therefore a component of the car. In the case of the bicycle the role is filled by the rider, who is external to the bicycle, and therefore not a component of the bicycle. The rider is, however, a participant in the bicycle's locomotion process. This is illustrated in Figure 12.

The relationships amongst the key concepts introduced in the preceding paragraphs are shown in Figure 13.

We are now in a position to begin to systematise the key relationships between objects and processes. We shall take 'object', 'process', 'part', and 'enactment' as primitive concepts, and shall express the relationships amongst them by means of the following axioms and definitions:[33]

- Axiom 1: A process is enacted by exactly one object at any one time.
- Definition 1: A process p is an *external process* of object o if o enacts p.
- Definition 2: Assume some external process p of o is causally dependent on an external process p' of some object o', then
  - (a) o' is a *component* of o with respect to p if o' is physically a part of o,
  - (b) o' is an *auxiliary object* for o with respect to p if o' is physically disjoint from o,
  - (c) p' is an *internal process* of o,
  - (d) p' *contributes* to p,

[32] Note that, strictly speaking, in Figure 11, the relationship between train-as-mobile and train-as-mechanism should involve a role: the position occupied in the diagram by train-as-mechanism should be given to the role of making the wheels turn, with the train-as-mechanism the internal role-holder that fills this role. This makes the train-as-mechanism a component of the train as mobile, which perhaps sounds strange, but is conceptually harmless. Similar remarks apply to the relationships between train-as-vehicle and both train-as-mobile and train-as-container, and indeed to many of the part-whole relationships portrayed in the diagrams in §5. Where components are concerned (as opposed to auxiliary objects), we shall regard the explicit reference to roles as optional.

[33] Of course this is not intended to be understood as a full axiomatic theory! It is one of the 'pointers' to a more formal treatment mentioned in §1.

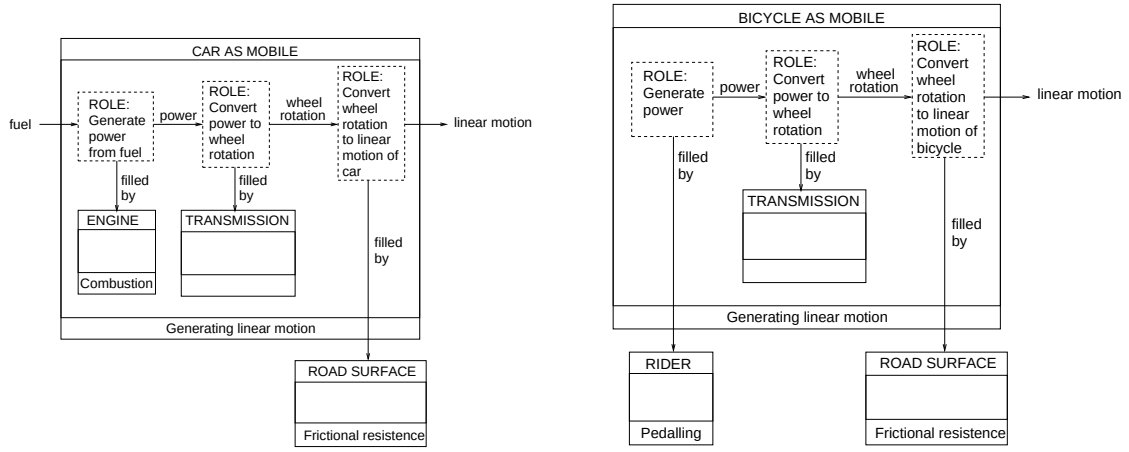

Fig. 12. Car with internal power source contrasted with bicycle with external power source.

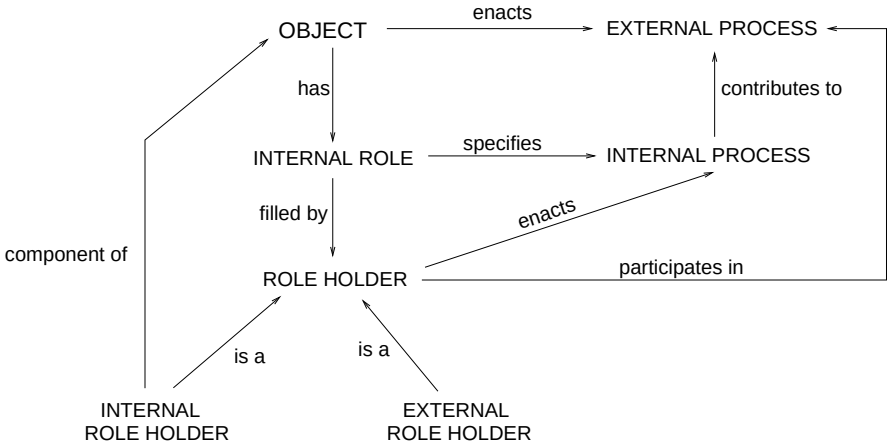

Fig. 13. The key concepts relating objects and processes

(e) o' *participates* in p, and
(e) process p' determines an *internal role* r of o, such that o' *fills* r.

It is important to note that components are here being understood as relative to some external process.[34] This is not always evident because for artefacts and organisms, some external processes are so salient that the components they define are generally regarded as functional parts in an absolute sense. The salient external processes of a human being, for example, are those which collectively form the life of the human being; this covers all our intentional bodily activities, and since these are supported by the internal processes of all the organs in the body, these organs are components with respect to those life processes. For this reason we would normally regard them as functional parts without further qualification. However,

[34]The relativity of functional parthood (which, as already noted, applies to components with respect to designed or evolved entities) is accorded due importance by [52], but this is with respect to a description under which the whole is considered rather than with respect to a process enacted by the whole. These are related, however, since in describing as object in a certain way one is implicitly referring to the characteristic behaviours of objects coming under that description, and thereby to the contributions of their parts to those behaviours. Here again we should emphasise that processes do not have to be dynamic: for example, the head of a bed (one of the examples considered by [52]) is a functional part by virtue of its role in supporting one end of the bed; here we would understand 'supporting the bed' to be a non-dynamic process which contributes to the bed's external processes of retaining a certain spatial configuration and supporting a sleeping person — but if a child is jumping up and down on the bed then the supporting function of the head of the bed becomes dynamic as it reacts to the varying forces imposed upon it .

consider now a process of *falling* (through the air, not just falling over), which humans occasionally enact. We would not normally regard this process as part of living. Now when a person falls, any arbitrary part of her falls, and it seems reasonable to say that the person as a whole falls if and only if all these parts do. That being the case, any arbitrary part of a person is a component with respect to falling. In this sense we could say that falling is *mereologically dissective*, meaning that any part of a falling object must itself be falling.[35] Some processes are *partially* dissective in this way: e.g., if an object is burning, then at least some parts of the object must be burning, but it is not necessary for all of them to be. But most of the processes which we ascribe to animate agents as agents are *mereologically unitary*: for example, when a person walks, no part of the person walks — in fact, the predicate 'walks' can only be meaningfully applied to whole organisms.

At first sight it may seem that the possibility of mereologically dissective processes contradicts Axiom 1. However, this is not the case. Consider again the case of the person falling. When the person falls, so does her head. Here we have one description, 'falling', for the process type that both the person and her head are enacting. But the particular individual processes are different: the person falling and the head falling. Each of these processes is enacted by a unique subject. In reality we would be unlikely to refer to the head falling unless the head were falling separately from the rest of the body (e.g., at an execution). If the whole person is falling then the process of interest is precisely that: the falling of the *whole* person. To attend to the falling of any of her parts is to open up the 'black box' and examine the internal processes which contribute to the external process.

The role played by the whole here is crucial; it is an immediate consequence of Axiom 1, together with the usual definition of proper part, which implies that a proper part is distinct from what it is a proper part of:

– Theorem 1. If an object is enacting a process, no proper part of the object is enacting that process.

If a process appears to be being enacted by several different individuals, in reality it is being enacted by a composite which has those individuals as parts. Each individual will be enacting its own process that contributes to the overall process, and these will of course be internal processes of the composite. If Tom and Kate shake hands, then Tom enacts the process of shaking Kate's hand, and Kate enacts the process of shaking Tom's hand, but if we want to regard the handshaking as a single process in its own right then we have to recognise the composite entity 'Tom and Kate' as what enacts it — and then Tom's shaking of Kate's hand and Kate's shaking of Tom's hand are internal processes of this composite entity. Again, if Tom and Kate together move a heavy table, then each of them enacts a process consisting of supporting one end of the table and pushing, pulling, etc; but the process of moving the table has to be regarded as enacted by a composite entity 'Tom and Kate' — and then Tom's manipulation of one end of the table and Kate's manipulation of the other are internal processes of the composite, contributing to its external process of moving the table. Of course the composite is not a mysterious third entity in addition to Tom and Kate: it just *is* Tom and Kate, considered together as a single subject.

It should be emphasised that these are by no means exceptional cases. Most of the objects we have considered — e.g., cars, bicycles, trains, people — are composites, and the processes they enact arise from the coordinated interactions between the processes enacted by their parts. Indeed, it is precisely such coordination that provides the ground for treating some collection of objects together as forming a composite: whereas unrestricted mereological fusion would allow the creation of innumerable 'fiat' objects whose only claim to existence is that someone has deemed them to exist, we feel that 'bona fide' composites must be grounded in something more substantial, and the existence of a process in need of a single enactor supplies precisely such a ground. We are not, of course, claiming that there is no difference between the tightly bound, relatively permanent composites exemplified by the common objects mentioned above, and the looser, temporary associations such as that formed by Tom and Kate while they

[35] This should be distinguished from the *temporal* dissectivity of processes discussed in §2, by which, e.g., any part of an interval of falling is an interval of falling. Nor should it be identified with *spatial* dissectivity, by which any part of a region in which it is raining is a region in which it is raining. This is related to the notion of a 'homogeneous predicator' introduced in [9].

are engaged in shaking hands or moving a table together. However, it is not easy to draw a hard and fast distinction between these two types, and for the purposes of the present discussion, there is no need to do so.

The temporal restriction ('at any one time') in Axiom 1 is needed in order to handle cases where numerically the same process is being enacted by different objects at different times. The falling of water in the waterfall is an example of this. As already discussed, this falling process is being enacted by different portions of water at different times.[36] So the temporal restriction in Axiom 1 is necessary.

7. What is an object?

As noted previously, any change must be a change *of* something. This is already an argument against a 'pure process' view of reality, since we cannot conceive of processes without their material support.

Nonetheless, phenomena like rivers (and even more, ocean currents) do present a problem. What, after all, is a river over and above a flow of water, i.e., a process? One might as well ask: what is a person over and above the sum of its internal processes? But what makes this sum worthy of consideration at all is that they constitute some kind of unity; the unity comes from the fact that there are other processes, its external processes, which it enacts. Thus these questions make the mistake of focussing only on the internal processes of the river or person, whereas the external processes play an essential a role in determining the identity of the object. Hence, rather than trying to characterise an object in terms of its internal processes (e.g., by identifying the object as the sum of those processes), we would rather say that *an object is a unity which is what enacts its external processes*. We could indeed say that the object is the *interface* between its internal and external processes: it is a point of stability in the world in virtue of which certain processes are characterised as internal and others as external.

There appears to be an element of circularity in this characterisation: not until we have determined the object can we decide which processes are to count as internal and which as external, and yet the object is supposed to be determined as something that mediates between the internal and external processes. This is a delicate issue which merits extended discussion.

First, a process cannot be classified as intrinsically either internal or external. Typically, a process can be seen as playing either role depending on the context in which it is considered. To repeat an earlier example, the swinging motion of the legs when a person is walking is an internal process from the point of view of the person but an external process from the point of view of the legs. Correlated with this swinging motion is a sequence of muscular contractions and relaxations: these processes are internal from the point of view of the legs, but external from the point of view of the muscles. The relativity of the internal/external distinction is important because it underpins the possibility of recognising functionally relevant structure in objects, to see objects as hierarchically decomposable into component objects. However, it also seems to suggest that there is an element of arbitrariness about how we divide the world up into individual objects, and to some extent this is indeed so; but this is unimportant — what is important is that the identification of an object is intrinsically bound up with the identification of the processes which it enacts and the processes that are internal to it.

To illustrate this, let us consider an example of an object which most people would readily consent to being in some sense 'virtual' or 'illusory': the mouse cursor on a computer screen. This enacts two kinds of external processes: it moves across the screen in response to the motion of the mouse, and it can change form depending on its on-screen location relative to other elements of the display. At any moment, the cursor consists of a configuration of distinctively coloured pixels on the screen. The cursor moves when the pixels undergo colour changes so that the distinctive cursor-configuration appears in a different position;

[36] One might be tempted, by an ontological sleight of hand, to avoid the temporal restriction by saying that at all times the falling process is being enacted by a single enduring object, 'the water in the waterfall', which, like the human body on a longer time scale, undergoes continuous replacement of material. But the trouble is that the position of this object, the water in the waterfall, is fixed: it does not fall, any more than the waterfall itself does! What falls are all the individual samples of water passing through it, and these are different at different times.

it changes form when the pixels undergo colour changes resulting in a different distinctive configuration. These colour changes are the cursor's internal processes. No pixel ever moves, and yet the cursor, which is a configuration of pixels, does move. One might well say that the 'motion' of the cursor is an illusion; and that therefore the cursor itself is an illusion; and yet if it is an illusion it is a very compelling one, and it seems highly unnatural to talk about the cursor in any terms other than as a fully-fledged object. The motion of the cursor is not so very different in this respect from the motion of a forest northwards, as trees at the southern margin die off and new trees spring up at the northern margin. No tree moves, yet the forest, which is made of trees, does move.

Nonetheless, there is an important difference between the motion of the forest and the motion of the cursor. In the terminology of Salmon [37], the former is a *causal process* whereas the latter is a *pseudo-process*. In a causal process, each successive phase is caused by the earlier phases (as, for example, new trees appearing at the northern margin of the forest provide the seeds from which later growth still further north will arise); whereas in a pseudo-process such causal links are absent, the successive phases being independently caused by some external process — in the case of the cursor this is the movements of the mouse (in Salmon's terms, a pseudo-process cannot 'transmit a mark'). Because the processes it enacts are pseudo-processes, the mouse cursor could perhaps be regarded as a 'pseudo-object', and this is the sense in which it is akin to an illusion. Salmon's processist orientation would presumably lead him to regard the cursor itself, and not just its motion, as a pseudo-process, thereby rendering the notion of 'pseudo-object' superfluous. Another example is a shadow and its associated movements — but we would normally describe this as insubstantial rather than illusory. The insubstantial or illusory nature of such processes might suggest that they can play no part in the causal history of the world: the processes themselves do not involve the transmission of energy, for example. However, this does not mean that they cannot be causes: the motion of the cursor gives rise to a perception by the computer user, which in turn affects his subsequent interaction with the computer; and the motion of a shadow can cause temperature changes in what it passes over.

Let us consider again the waterfall. How exactly is the waterfall determined as an object in terms of various processes conceived as being either internal or external? Consider in detail what is going on here. We assume to begin with that we are given a river as an object (thus the question of where the river gets its identity from is assumed to be irrelevant to the question regarding the waterfall). The river's internal process is the flow of water along a channel. At a certain point the normally smooth profile of the channel is interrupted by a more or less vertical rocky declivity. When it reaches this point, the flow of the water changes character so that instead of flowing horizontally along the channel, it tumbles over the edge of the declivity and falls freely through the air until it meets the continuation of the channel below, whereupon it eventually resumes its regular horizontal course along the channel. Over a long period of time, the flow of water over the edge of the declivity eats into the rock and wears it away, the resulting particles of eroded rock being carried downstream by the river. As a result of this, the declivity and the associated disruption of the flow of water in the river migrate slowly upstream.

We can identify a number of processes going on here:

1. Water falling.
2. The transfer of the river's flow from the high-level channel to the low-level one.
3. Rock being eroded from the edge of the declivity.
4. Particles of rock being carried off downstream.
5. Upstream movement of the declivity and its associated processes.

The first of these, 'water falling', can be analysed in more detail as the successive falling of an open-ended sequence of individual portions of water, each of which falls from the top of the declivity to the bottom. Water being effectively continuous, it is arbitrary how we decide to divide the water up into portions, but whatever portion we attend to, we can say that that water enacts its own falling process. These processes, in cumulation, result in the transfer of level of the river's flow, but the relation here is constitutive rather than causal. The actual cause of the water falling (and hence of the transfer of level) is the combination of the water's own gravity and the absence of any impediment to provide a counteracting upward force.

The latter may be regarded as the means by which the configuration of the rocky precipice between the upper and lower channels allows the falling to happen. In the device ontology, a device is regarded as exhibiting an external function which consists of the transformation of some input into some output; inside the 'black box' are the internal processes which constitute the means by which the external function is achieved. In this spirit, since absence of support resulting from presence of the rocky precipice is the means by which the transformation of the river's flow from the higher level to the lower is achieved, we can postulate the existence of a 'device' whose 'function' is precisely the latter transformation and whose internal 'process' is the maintenance of that absence of support which enables it to happen. This 'device' is the waterfall viewed as a conduit (illustrated in the left-hand diagram of Figure 7). Thus, on this view, we can say that the external process (transfer of level) is accomplished by the largely passive internal process of maintaining lack of support for the water. This internal process is enacted by the rock by virtue of its particular configuration.

As already noted, however, the conduit model does not correspond to our everyday idea of a waterfall, which essentially consists of falling water. On this view, the water, while it is falling, *is* the waterfall, it is not merely conveyed by it. Its falling is an internal process, and constitutes the means by which items 2–5 in the above list are brought about. This is the view illustrated on the right in Figure 7.

Note that not every question about the waterfall is automatically answered here, but this need not concern us: there should be no requirement that an ontology should be fully determinate. We have presented two views of the waterfall: one consisting of the rocky precipice without the water, the other consisting of the water without the rocky precipice. Could these be combined into a third view in which the waterfall consists of both the water and the precipice? In such a view, both the falling of the water and the erosion of the rock are internal processes; the external process is again the gradual migration of the waterfall upstream — only now it is the migration of the whole water/rock complex rather than just the rock face.

It might be objected that, in tying the identity — indeed existence — of objects to their associated processes, we are restricting our account of objects to those which exhibit complex behaviour, enabling them to be understood, even if only metaphorically, as devices of some kind. But at first sight not all objects are like this. As already noted, perhaps the most basic kind of object is a simple chunk of matter: a piece of rock, for example, or a gold nugget. In what sense can these possibly be regarded as devices? Although a simple piece of rock may seem inert and lifeless, in fact it exhibits some remarkable behaviour. If you leave it alone, it will just sit there — but significantly, it *stays in one piece*. The maintenance of the rock's integrity is a process: from the point of view of the rock, it is an internal process enacted by the constituent grains of the rock (crystal fragments, molecules, or whatever may be the appropriate level of description for the immediate constituents of the kind of rock under consideration), whose mutual cohesion holds the rock together. Without this constant ongoing process, the external processes enacted by the rock could not occur. These external processes are primarily various forms of motion: push the rock, and it moves; hold it above the ground and let go, and it falls; apply a twisting force to it, and it rotates. None of these things could happen if the rock did not maintain its inner integrity. Thus the rock can be considered to be a 'device' which converts external forces into various forms of motion; and the means by which it does this is simply by holding together so that its various parts move in a coordinated fashion. In this way we can see that the rock is not so very different from a complicated mechanism built from many moving parts. The latter may seem more significant to us, and this may be because, for the most part, chunks of matter matter to us not for themselves but as components of more complex objects which matter to us more.

What, then, is an object? Consider a situation from which we can isolate two collections of processes, called $\mathcal{I}$ and $\mathcal{E}$, with the following properties:

1. The collections $\mathcal{I}$ and $\mathcal{E}$ are disjoint.
2. There is a level of description at which the situation can be coherently described as containing the processes in $\mathcal{I}$ but not those in $\mathcal{E}$.
3. There is another, higher level of description at which the situation can be coherently described as containing the processes in $\mathcal{E}$ but not those in $\mathcal{I}$.
4. The processes in $\mathcal{E}$ are causally dependent on the processes in $\mathcal{I}$.

In this case, we say that there is an object, o, such that

5. $\mathcal{I}$ is a collection of internal processes of o.
6. $\mathcal{E}$ is a collection of external processes of o.
7. o enacts each of the processes in $\mathcal{E}$.
8. o is sustained by the processes in $\mathcal{I}$.
9. For each of the processes in $\mathcal{I}$ we can define a role in the internal description of o, and for each such role there is either a (functional) part of o or an auxiliary object which enacts it.

The level of description referred to in (3) corresponds to the viewpoint according to which o is regarded as a 'device' in the sense of the device ontology.

Note that we do not insist that the collections $\mathcal{I}$ and $\mathcal{E}$ are maximal: it is not necessary, in order to establish the existence of an object, that we enumerate *all* its internal and external processes. This would be an impossible demand to satisfy; but on the other hand, the more processes we can collect together in this way, the more robustly founded is the identity of the object thereby characterised.[37] Note also that the objects we recognise in any given situation depend critically on which sets of processes $\mathcal{E}$ and $\mathcal{I}$ we select. These have to be selected in a way which satisfies (1-4) above, but there may be many selections that are consistent with these requirements. Thus we can say that while there may be no absolute truth as to which objects a given situation contains, it is far from being the case that the identification of objects is a matter of arbitrary choice.

To test these ideas, let us see how they fare when applied to collectives. In §5 we asked 'When does a collection of individuals constitute a higher-order entity?'. We can now say that this happens just when we can identify certain processes as external and others as internal. Consider the passengers on a train. During the journey, the passengers engage in individual behaviours which are largely uncoordinated: some are reading, some are chatting to each other, some are sleeping, etc. These processes are all external to the passengers as individuals. However, there does not seem to be any higher-level process to which they causally contribute, and therefore we will not regard such processes as internal processes of any collective. On the other hand, each one of the passengers is also enacting a process of travelling, by the mere fact of being on the train and thereby participating in its motion. When the train reaches the terminus all the passengers alight from the train (for simplicity we assume there are no intermediate stops), walk to the exit, and go their separate ways: in other words, they *disperse*. These processes — the coordinated movement followed by dispersal — are causally dependent on the travelling processes enacted by all the individuals. Thus we can regard the former as our collection $\mathcal{E}$, the latter as our collection $\mathcal{I}$, and hence recognise the existence of a higher-order entity formed as a collective from the individual passengers. We are not stating that the collective has to exist in any absolute sense, but only that, to the extent that it is reasonable to draw a clear distinction between appropriately related sets of processes we can call 'internal' and 'external', it is also reasonable to designate a single entity to be the bearer of those processes. And it is clear that some collectives have a much stronger claim to being recognised as entities in their own right than others.

As a principle for object-identification, our process-based criterion should be regarded as supplementing, rather than entirely replacing, more traditional criteria based on demarcation from the environment. According to Smith [43], a *bona fide* object must have a *bona fide* boundary, which is to say, it must be demarcated from its surroundings by a genuine material discontinuity forming its surface or boundary. This is to distinguish *bona fide* entities from *fiat* entities (such as, for example, administrative divisions), which only exist by virtue of boundaries imposed by human cognitive acts [44]. Both *bona fide* and *fiat* objects, however, are characterised by means of boundary delimitation, the difference being the kind of boundary in question. But our process-based criterion makes no reference to boundaries, which is why it is able to pick out entities such as collectives through their behaviour even though it may be impossible to circumscribe them by any meaningful boundaries. These entities are not necessarily *fiat* — the processes

[37] Of course, by linking the identity of the object to the identity of certain processes, we open up the possibility of cases where the identity of an object is in doubt because the identity of certain processes is; but this fact in no way invalidates the proposal to tie objects to processes. Ontology is not, and never will be, a wholly cut-and-dried enterprise.

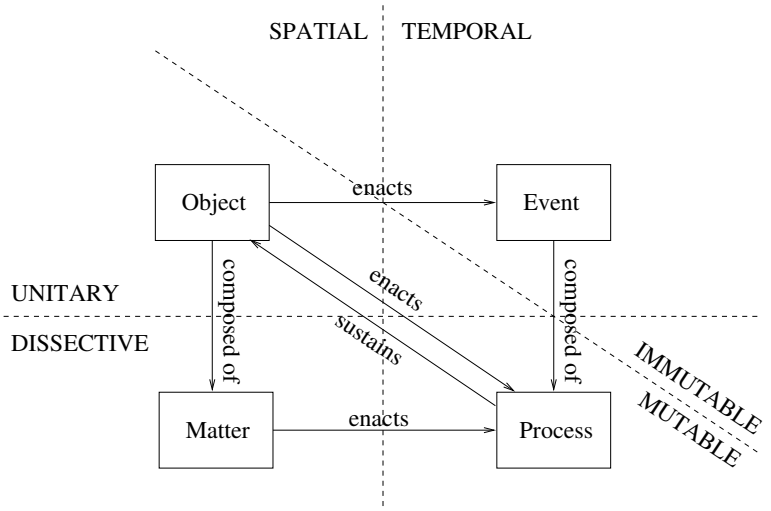

Fig. 14. The four key categories of physical phenomena

they enact do not exist solely by virtue of some human convention, although there may be an element of convention in the selection of process sets $\mathcal{I}$ and $\mathcal{E}$ used to characterise the entities — but they would seem not to be *bona fide* in Smith's sense either. Thus the process criterion has given us access to aspects of reality that the demarcation criterion is blind to.

8. Concluding remarks

The central theme of this paper is the mutual ontological dependency between processes and objects. No object can exist without enacting various processes; minimally, these will include a process of *persistence*, an external process that is typically supported by the internal processes by which the object's internal coherence is maintained, as discussed previously in relation to the piece of rock. More strongly, we can say that an object is what it is because of the processes it enacts, either actually or potentially. Equally, a process cannot exist without an object to enact it. This tight mutual dependence suggests that if we really want to find a minimal independent unit, it must be an object-process pair. However, since an object will typically be involved in many processes, and processes are causally dependent on other processes which in turn bring in more objects, a more realistic independent unit might be an object-process complex. The more complete we try to make our description of the original object the wider this complex will extend, and it seems likely that ultimately it must encompass the whole universe. For this reason we never deal with complete descriptions of anything, and that is why the idea of considering objects from different *points of view* is so important, as illustrated by the various views of the train in Figures 8–11.

The processes which determine an object's nature and identity (the object's external processes) are causally dependent on other processes, which we have described as internal processes of the object. Although the internal processes are not 'visible' to us when contemplating the object as a black box, the object could not exist without them. The operation of the internal processes sustains the object by giving rise to the external processes which constitute the object's life and being. We thus find that objects stand in two distinct relations to processes: on the one hand the object enacts its external processes and on the other it is sustained by its internal processes. Our basic ontological diagram (Figure 3) thus needs to be completed by inclusion of an arrow representing the 'sustains' relation. This is shown in Figure 14.

The existence of arrows going both ways between objects and processes may give the impression of a vicious circularity, but this appearance is misleading. There is no real circularity since the processes which sustain an object are not the same as the processes which the object enacts. The apparent circularity is dispelled in Figure 15, in which it is unfolded into a potentially infinite hierarchy of objects enacting their

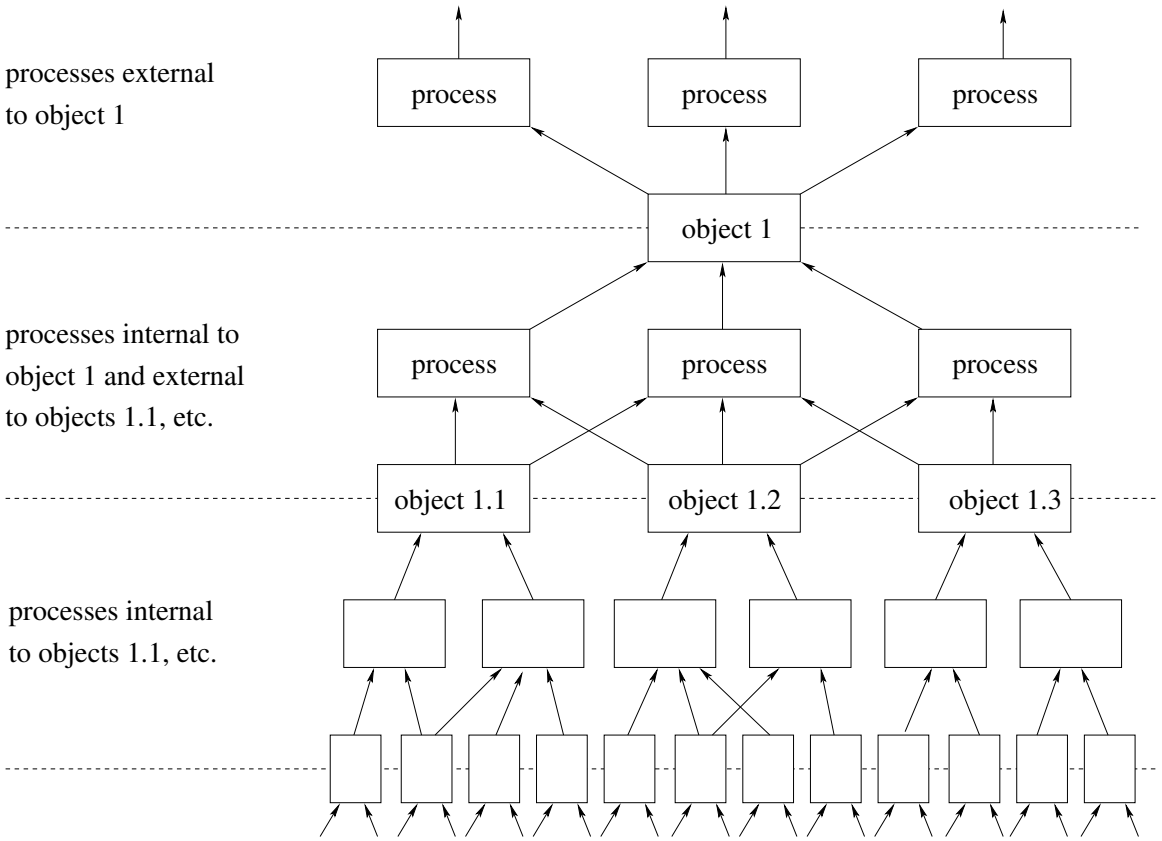

Fig. 15. Unfolding the apparent circularity

external processes and sustained by their inner processes which in turn function as external processes for a lower-level layer of objects. While this dispels any possibility of circularity between the definitions of particular objects and processes, or indeed between particular *kinds* of objects and processes, one may still be troubled by the impression of circularity arising from the mutual interdependence of the two ontological categories of process and object. But this is not, in fact, so very different from the relationship between, say, the notions of 'object' and 'quality': objects cannot exist without qualities, and qualities presuppose objects to have them. One might say the same concerning 'matter' and 'form', which provide the basis for Aristotle's hylomorphic dualism. We believe that despite the best efforts of formal ontologists this kind of thing is inescapable, and should not be regarded as a fatal weakness in the theory.

Does the hierachy 'bottom out' eventually in a 'bedrock' layer of irreducibly atomic elements? Such elements would indeed have to be object/process pairs: objects sustained by processes which they themselves enact, processes which are both internal and external to the same objects. Perhaps such things as photons or quarks can be described in this way, but that is no part of our business in this paper. In truth nobody knows whether the downward hierarchy of granularities is finite or infinite, and we have no wish to prejudge this.[38] Crucially, though, if there is such a bottom level, then we would be forced to retract property 1 (the disjointness of $\mathcal{I}$ and $\mathcal{E}$) in certain circumstances.

Another key theme of this paper is the *dual nature* of processes. In their tight coexistence with objects, they partake of many of the properties which characterise objects as continuants, in particular their *presence* as an actual ingredient of the world as it exists at a time. As such they, like objects, have histories and undergo change. But processes are intrinsically temporal in a way that objects are not: they extend

[38]Even so, it is tempting to see here an echo of the wave/particle duality intrinsic to quantum-level phenomena. A quantum particle can be understood as both an object and a process; if these two aspects are mutually dependent, then we would appear to have a self-sustaining bottom level for our hierarchy.

through time. Viewed in this way, processes can be seen as the 'stuff' of which durative events are made. These two aspects of processes, at first sight contradictory, are reconciled by means of the observation that the temporal extension of the a process may be regarded as occupying a temporal "window" which moves forward through time, thereby allowing the process itself to have a history and undergo change.

Events are dependent on processes in a way that is analogous, but not identical, to the way in which objects are dependent on matter. But in fact the dependence of events on processes is simpler than that of objects on matter because events cannot change whereas objects can. An event is a fixed configuration of 'chunks' of process, whereas an object may (and typically does) consist of different configurations of chunks of matter at different times. The explanation for the difference is that whereas events extend along the same dimension (i.e., time) as that in which change occurs, objects extend in space, orthogonally to the time dimension, leaving room for change to occur with respect to the way in which they are extended.

Finally, we must emphasise that all this is a work in progress. Much remains to be done to tie together more firmly the various strands which have informed this paper: considerations about time and change, matter and objects, processes and events, the device ontology as an inspiration for more general ways of looking at objects, and the important role of roles in elucidating the internal structure of the black box. But even in their current incomplete state, we offer our ideas in the hope of stimulating a much-needed extension of current practice in ontological engineering to encompass the dynamic aspects of the world in a more thoroughgoing way.

Acknowledgements

We would like to thank the three reviewers of this paper (Joost Breuker, Luc Schneider, and one who chose to remain anonymous) for their detailed commentaries on both the original submission and our first revision. We would also like to thank Laure Vieu for her careful reading of the paper and insightful comments.

References

[1] E. Bach. The algebra of events. *Linguistics and Philosophy*, 9:5–16, 1986.
[2] M. H. Bickhard. Process and emergence: Normative function and representation. In Johanna Seibt, editor, *Process Theories: Crossdisciplinary Studies in Dynamic Categories*, pages 121–155. Kluwer Academic Publishers, Dordrecht, 2003.
[3] R. Casati and A. Varzi. *Parts and Places: The Structures of Spatial Representation*. MIT Press, Cambridge, Mass. and London, 1999.
[4] W. Charlton. *Aristotle's Physics, Books I and II, translated with Introduction and Notes*. Clarendon Aristotle Series. Clarendon Press, Oxford, UK, 1970.
[5] D. Davidson. The individuation of events. In N. Rescher, editor, *Essays in Honor of Carl G. Hempel*. D. Reidel, Dordrecht, 1969. Reprinted in D. Davidson, *Essays on Actions and Events*, Oxford, Clarendon Press, 1980, pages 163-80.
[6] E. Davis. *Representations of Commonsense Knowledge*. Morgan Kaufmann, San Mateo, Calif., 1990.
[7] D. Dowty. *Word Meaning and Montague Grammar*. D. Reidel, Dordrecht, 1979.
[8] F. Dretske. Can events move? *Mind*, 76:479–92, 1967.
[9] C. Eschenbach. A predication calculus for qualitative spatial representations. In C. Freksa and D. M. Mark, editors, *Spatial Information Theory: Cognitive and Computational Foundations of Geographic Information Science (Proceedings of COSIT'99)*, pages 157–172, Berlin, 1999. Springer.
[10] Institute for Formal Ontology and Medical Information Science (IFOMIS). Basic formal ontology (BFO). http://www.ifomis.org/bfo/. Accessed 11th May 2009.
[11] A. P. Galton. Dynamic collectives and their collective dynamics. In D. M. Mark and A. G. Cohn, editors, *Spatial Information Theory*, pages 300–315. Springer, 2005.
[12] A. P. 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.
[13] A. P. Galton. Processes as continuants. In J. Pustejovsky and P. Revesz, editors, *13th International Symposium on Temporal Representation and Reasoning (TIME 2006)*, page 187. IEEE Computer Society, 2006.
[14] A. P. Galton. On the paradoxical nature of surfaces: Ontology at the physics/geometry interface. *The Monist*, 90(3):379–390, 2007.
[15] A. P. Galton. Experience and history: Processes and their relation to events. *Journal of Logic and Computation*, 18(3):323–340, 2008.

[16] A. P. Galton and J. M. Hood. Anchoring: A new approach to handling indeterminate location in GIS. In D. M. Mark and A. G. Cohn, editors, *Proceedings of COSIT'05*, pages 1–13. Springer, 2005.
[17] P. Grenon and B. Smith. SNAP and SPAN: Towards dynamic spatial ontology. *Spatial Cognition and Computation*, 4(1):69–104, 2004.
[18] P. M. S. Hacker. Events and objects in space and time. *Mind*, 91:1–19, 1982.
[19] S. Haslanger and R. M. Kurtz, editors. *Persistence: Contemporary Readings*. MIT Press, Cambridge, MA, 2006.
[20] K. Hawley. *How Things Persist*. Clarendon Press, Oxford, 2001.
[21] P. J. Hayes. The naïve physics manifesto. In D. Michie, editor, *Expert Systems in the Micro-Electronic Age*, pages 242–70. Edinburgh University Press, 1979.
[22] M. Heller. *The Ontology of Physical Objects: Four-dimensional Hunks of Matter*. Cambridge University Press, Cambridge, 1990.
[23] J. R. Hobbs and R. C. Moore, editors. *Formal Theories of the Commonsense World*. Ablex Publishing Corporation, Norwood, NJ, 1985.
[24] J. M. Hood and A. P. Galton. Implementing anchoring. In M. Raubal, H. J. Miller, A. U. Frank, and M. F. Goodchild, editors, *Geographic Information Science: Proceedings of the 4th INternational Conference, GIScience 2006*, pages 168–185. Springer, 2006.
[25] Y. Kitamura, Y. Koji, and R. Mizoguchi. An ontological model of device function: Industrial deployment and lessons learned. *Applied Ontology*, 1(3-4):237–262, 2006.
[26] R. Le Poidevin. The experience and perception of time. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. 2004. http://plato.stanford.edu/entries/time-experience/.
[27] C. Masolo, S. Borgo, A. Gangemi, N. Guarino, and A. Oltramari. WonderWeb Deliverable D18: Ontology Library (final). Technical report, LOA-ISTC, CNR, 2003.
[28] D. H. Mellor. *Real Time*. Cambridge University Press, Cambridge, 1981.
[29] R. Mizoguchi and Y. Kitamura. Foundation of knowledge systematization: Role of ontological engineering. In R. Roy, editor, *Industrial Knowledge Management — A Micro Level Approach*, pages 17–36. Springer-Verlag, London, UK, 2000.
[30] R. Mizoguchi, E. Sunagawa, K. Kozaki, and Y. Kitamura. The model of roles within an ontology development tool: Hozo. *Applied Ontology*, 2(2):159–179, 2007.
[31] M. Moens and M. Steedman. Temporal ontology and temporal reference. *Computational Linguistics*, 14:15–28, 1988.
[32] J. M. E. Moravcsik. Strawson and ontological priority. In R. J. Butler, editor, *Analytical Philosophy: Second Series*, pages 106–119. Basil Blackwell, Oxford, 1965.
[33] A. P. D. Mourelatos. Events, processes, and states. In P. Tedeschi and A. Zaenen, editors, *Tense and Aspect*, pages 191–212. Academic Press, New York, 1981.
[34] I. Niles and A. Pease. Towards a standard upper ontology. In C. Welty and B. Smith, editors, *Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS-2001)*, pages 2–9. ACM Press, 2001.
[35] N. Rescher. Process philosophy. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. 2008. http://plato.stanford.edu/entries/process-philosophy/.
[36] J. H. Roberts. Activities and performances considered as objects and events. *Philosophical Studies*, 35:171–185, 1979.
[37] W. C. Salmon. *Scientific Explanation and the Causal Structure of the World*. Princeton University Press, Princeton, New Jersey, 1984.
[38] M. Sasajima, Y. Kitamura, M. Ikeda, and R. Mizoguchi. FBRL: A function and behavior representation language. In C. S. Mellish, editor, *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95)*, pages 1830–1836, San Mateo, California, 1995. Morgan Kaufmann.
[39] J. Seibt, editor. *Process Theories: Crossdisciplinary Studies in Dynamic Categories*. Kluwer Academic Publishers, Dordrecht, 2003.
[40] T. Sider. *Four-dimensionalism: An Ontology of Persistence and Time*. Clarendon Press, Oxford, 2001.
[41] P. Simons. *Parts: a Study in Ontology*. Clarendon Press, Oxford, 1987.
[42] J. J. C. Smart. The river of time. *Mind*, 58:483–94, 1949.
[43] B. Smith. On drawing lines on a map. In A. U. Frank and W. Kuhn, editors, *Spatial Information Theory: A Theoretical Basis for GIS*, volume 988 of *Lecture Notes in Computer Science*, pages 475–84, New York, 1995. Springer-Verlag. Proceedings of International Conference COSIT'95.
[44] B. Smith. Fiat objects. *Topoi*, 20:131–148, 2001.
[45] B. Smith and A. Varzi. Fiat and bona fide boundaries: Towards an ontology of spatially extended objects. In S. C. Hirtle and A. U. Frank, editors, *Spatial Information Theory: A Theoretical Basis for GIS*, volume 1329 of *Lecture Notes in Computer Science*, pages 103–119, New York, 1997. Springer-Verlag. Proceedings of International Conference COSIT'97.
[46] R. Stout. Processes. *Philosophy*, 72:19–27, 1997.
[47] R. Stout. The life of a process. In G. Debrock, editor, *Process Pragmatism: Essays on a Quiet Philosophical Revolution*, pages 145–158. Rodopi, Amsterdam/New York, 2003.
[48] R. Sylvan. Process and action: Relevant theory and logics. *Studia Logica*, 51:379–437, 1992.
[49] B. Taylor. *Modes of Occurrence*, volume 2 of *Aristotelian Society Series*. Basil Blackwell, Oxford, 1985.
[50] A. Varzi. Vagueness in geography. *Philosophy and Geography*, 4(1):49–65, 2001.
[51] Z. Vendler. *Linguistics and Philosophy*. Cornell University Press, Ithaca, 1967.
[52] L. Vieu and M. Aurnague. Part-of relations, functionality and dependence. In Michel Aurnague, Maya Hickmann, and Laure Vieu, editors, *The Categorization of Spatial Entities in Language and Cognition*, volume 20 of *Human Cognitive Processing*, pages 307–336. John Benjamins Publishing Company, http://www.benjamins.nl, 2007.
[53] D. Wiggins. *Sameness and Substance*. Basil Blackwell, Oxford, 1980.
[54] M. E. Winston, R. Chaffin, and D. Herrmann. A taxonomy of part-whole relations. *Cognitive Science*, 11:417–444, 1987.

[55] Z. M. Wood and A. P. Galton. A new classification of collectives. In C. Eschenbach and M. Gruninger, editors, *Formal Ontology in Information Systems: Proceedings of the Fifth International Conference (FOIS 2008)*, pages 109–120. IOS Press, 2008.
[56] M. F. Worboys. Event-oriented approaches to geographic phenomena. *International Journal of Geographical Information Science*, 19(1):1–28, 2005.
[57] E. M. Zemach. Four ontologies. *Journal of Philosophy*, 67:213–247, 1970.