COSIT 2009

Aber Wrac'h, Brittany, France

Tutorial, Monday 21st September 2009

PROCESSES AND EVENTS IN GEOGRAPHICAL SPACE

Antony Galton

University of Exeter, UK

Contents

SESSION 1: Scope

1 What kinds of processes and events are of interest to GI Science?

If anyone should doubt whether processes, events, and more generally *change*, are of interest to geography, and therefore to GI Science, the following list should dispel those doubts:

Kinds of Geographical Change

- Landform
  - Coastal erosion and deposition
  - River courses (meanders, etc.)
  - Sand-banks, spits, dune-formation
  - Glaciers retreating/advancing
  - Lakes drying up, splitting, merging
- Climate and weather systems
  - Flooding
  - El Niño, Gulf Stream, jet streams
  - Global warming
  - Cyclical vs progressive changes
- Vegetation
  - Natural changes—ecological succession
  - Forest clearance/plantation
  - Ploughing up grassland
  - Rotation of crops
- Land-use. Changes amongst arable, pasture, forestry, recreation, industrial, urban, conservation. And within these categories:
  - Arable: rotation of crops
  - Forestry: Age structure of plantation
  - Urban: new buildings, demolition, road layout, railway closure
- Stipulative (status), leading to physical changes via human behavioural response.
  - Land ownership, administration (boundary changes etc)
  - Status of roads (upgrade/downgrade)
  - Designation as common land, conservation area, military zone, etc
  - Imposition of congestion charges, parking restrictions, speed limits, etc
  - Name changes
- Human population
  - Immigration/emigration, birth-rate, mobility, wealth

  - Communication and transport
  - Political and economic factors

- Animal populations
  - Response to any of the above
  - Hunting, shooting and fishing
  - Domestication

Although everything in geography eventually changes, it is important to realise that things change at very different rates. Many things change sufficiently slowly that, for most purposes, they can be regarded as static. If this were not so, there could be no use for maps. Things which change over short time scales cannot usefully be mapped: e.g., what would be the point of a road map which claims to show the location of every vehicle at the time it was published? A road map shows a relatively static situation — a particular configuration of roads — within which motions occur that are too fast to be mapped. Similarly, a blob of green on a map may represent a woodland whose shape and position remain constant over many years: but within this woodland there is constant change as trees grow and die, as well as the yearly seasonal cycle.

Even though individual changes may be too fast to be mapped, *patterns* of change can be stable over a long period. Hence we can find static maps showing ocean currents, wind circulation, the paths of migratory birds, etc. But over a longer period, a pattern of change can itself change.

We should distinguish on-going *processes* from individual discrete *events*. We can assert that a process is in operation right now, in the present, but usually by the time we are in a position to assert the occurrence of an event, it is already in the past. So events are historical in nature; but this does not mean they are not of interest to geography, since the current state of the world has been brought about by the events of the past. Events are sometimes shown on a map: e.g., a battle site with a date attached; and historians make use of maps showing, e.g., the movements of armies in a war by means of annotated arrows. These are events — though while they were happening what was observable were the processes constituting the events. I will have more to say on processes and events and their relationships in the next section.

2 Terminological matters: What do we mean by 'process' and 'event'?

To begin with, an important WARNING: *The word 'process' is used in several different ways!* In particular, I want to distinguish between two common uses:

1. *Process as an ongoing open-ended activity*.
   the flowing of a river or ocean current
   the back and forth movement of the tides
   the growth of a tree
   raining
   photosynthesis
   coastal erosion
   human activities such as walking, running, eating, singing, sleeping

2. *Process as a specific closed sequence of actions, leading to a definite endpoint*.
   making a pot of tea
   baking a cake
   shutting down a computer
   constructing a by-pass

boarding a plane
performing an appendicectomy
giving birth
spinning a web

To avoid confusion, I shall reserve the word ‘process’ to refer to the first kind; I shall refer to the second kind as **routines**. The main differences between processes and routines are listed in Table 1.

Processes	Routines
At sufficiently coarse granularity, processes may be conceptualised as homogeneous	Each instantiation of a routine is an event, which at sufficiently coarse granularity may be conceptualised as point-like.
A process can in principle stop at any time without thereby being considered ‘incomplete’	There can be incomplete instantiations of a routine, which are interrupted before they finish
A process is like an ordinary object in that it can be meaningfully said to undergo change (e.g., becoming faster or slower)	It does not seem to make sense to ascribe change to routines

Table 1: The main differences between processes and routines

The word ‘event’ is also used in several different ways, but always, an event is regarded as *a discrete individual occurrence with a more or less definite beginning and ending*. Important general types of event are:

- **Transitions.** A transition from a situation in which some proposition holds to one in which it does not, or vice versa. Typically this is the onset or cessation of some process or state, e.g., the water starts to flow, the sun rises or sets, it starts or stops raining.
- **“Chunks”.** These are *bounded instantiations of processes*, e.g., someone walks, runs, sings, eats, or sleeps for a while [= a “chunk” of walking, running, singing, eating, sleeping], an object falls to the ground [= a “chunk” of falling], a bird flies from one tree to another [= a “chunk” of flying].
- **Instantiations of routines.** These are specific occurrences consisting of complete or incomplete instantiations of some routine, e.g., a particular occasion of someone making a cup of tea, or giving birth.

Although chunks are defined in terms of processes, they are distinct from the processes themselves. A process is

- *open-ended*, i.e., it does not include start and end points, and
- *dissective*, meaning that, e.g., any subinterval of an interval on which running occurs is an interval on which running occurs.

By contrast, a *chunk* of running is

- *closed*, being delimited by starting and stopping events which form an essential parts of the chunk,
- *non-dissective* (or *unitary*), meaning that no part of a chunk of running is itself a chunk of running.[1]

[1]Note the analogy here with matter: compare ‘gold’ with a ‘chunk of gold’ — the latter, but not the former, includes a surface as part of its definition, the surface being what delimits the extent of the chunk, analogous to the starting and stopping events of a process-chunk. The analogue of (non)-dissectivity is that whereas parts of gold are gold, the parts of a chunk of gold are not chunks of gold (because they are not completely delimited by a surface).

Note that events may be *punctual* (i.e., literally instantaneous, such as the onset of motion of a particle) or *durative* (i.e., having duration, possibly very short). But the distinctive thing is that with an event there is always some granularity level at which it can be conceptualised as point-like. E.g., consider *the extinction of the dinosaurs* — this event may have taken thousands of years, but from the perspective of many millions it marks a point in the Earth's history.

3 The relationship between events and processes

Events are generally dependent on processes. This dependency may take several forms:

- A durative event is typically "made of" processes. This is obvious when it is just a uniform "chunk" of some process — e.g., *He walked for an hour*, this reports an hour-long event made of walking (compare a metre-long plank made of wood).
- A durative event may be an instantiation of a complex routine, in which case it may be composed of a number of distinct process chunks representing different phases of the event as a whole. How it is analysed will often depend on the level of detail required.
- A punctual event is usually the onset or cessation of a process ("It started raining") [or state ("the aircraft touched down")].

Similarly, processes are generally dependent on events:

- A process may consist of events, in the sense of being an open-ended repetition of some event or sequence of events. E.g., the process of hammering consists of a repetition of individual hammer-blows; the process of leaf-fall consists of many individual falling events (each of which is individually a chunk of falling process enacted by a leaf — thus showing how there can be multiply-layered dependencies spanning a range of levels of detail).
- A "higher-level" process may exist by virtue of some complex event (e.g., a routine) being under way, e.g., a house is being built — this is a process which takes different forms at different stages, but we can (at least for some purposes) treat what is going on at these different stages as all one process by virtue of its relationship to the completed event. Again, this is a matter of granularity — at a finer granularity we may lose sight of the relationship to the larger event and only see a sequence of distinct process-chunks.[2]

4 Types and Tokens

With both events and processes we can distinguish between generic *types* and individual instances of those types, known as *tokens*.[3]

The type/token distinction is fairly straightforward to apply in the case of events. A general term such as "earthquake" denotes an event type. Its tokens are all the individual earthquakes that actually occur, e.g., the Lisbon earthquake of 1755, and the great San Francisco earthquake of 1906. (This is similar to the case of objects — e.g., Everest, Kilimanjaro and Aconcagua are all tokens of the generic type "mountain".)

With processes, the type/token distinction becomes more problematic. The difficulty is to determine what counts as an individual instance of a process. Consider a process such as raining:

[2] This relationship between a routine and its higher-level process is no doubt the reason why the term "process" is applied to routines as well as processes in the sense understood here.

[3] This is independent of any particular philosophical stance taken concerning the ontological status of types — e.g., some people take types to be abstract inhabitants of some platonic realm distinct from the physical world, whereas others would locate types in the physical world, existing in or through their instances.

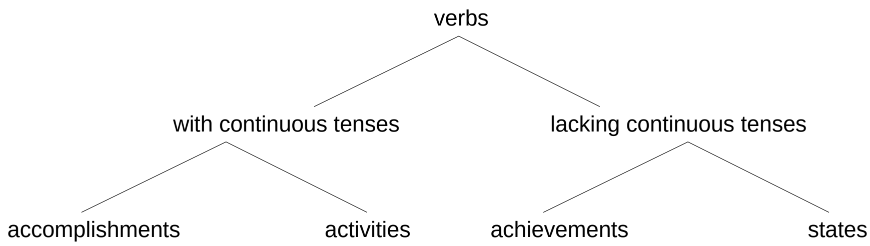

Figure 1: Vendler's classification [84]

- We say "The rain became heavier" — if we take the ontological implications of "the rain" seriously, this must mean that one and the same process was less heavy at time t_1 and heavier at later time t_2. (Otherwise, we have to say that the raining process at t_1 was *replaced* by a different, heavier raining process at t_2 — but if it is raining continuously between t_1 and t_2, becoming steadily heavier, there is no non-arbitrary point at which we can say the replacement occurs; one might say that there is a different process in operation at each instant, but this seems bizarre.)
- "The river dried up June and began flowing again in September". Is this the same individual flowing process, as it were "reincarnated", or a new process of the same type?

These questions may seem rather remote from practical concerns, but if we are to provide a viable systematic ontology of processes for use in an information system, then we need to find consistent answers to them. In my experience, this is much harder than you might think.

5 Classic Taxonomies

The classification of states, processes, and events was studied by philosophers and linguists for many years before this became an interesting topic for geographers and information scientists.

An influential and much cited classification is due to Vendler [84], who used the terms *State*, *Activity*, *Achievement* and *Accomplishment* to designate four 'common time schemata implied by the use of English verbs', as follows (see also Figure 1):

- Verbs that possess continuous tenses
  - Verbs denoting actions with 'a "climax" which has to be reached if the action is to be what it is claimed to be' (*Accomplishments*)
  - Verbs denoting actions with 'no set terminal point' (*Activities*)
- Verbs that do not possess continuous tenses
  - Verbs that 'may be predicated only for single moments of time' (*achievements*)
  - Verbs that 'may be predicated for shorter or longer periods of time' (*States*)

Since Vendler is classifying *verbs*, his criteria are expressed in linguistic terms (e.g., continuous tenses, predication); but since verbs are used to describe what goes on, it is natural that much of the discussion should focus on the actions, activities, etc, themselves rather than on the specific linguistic expression of them. This is reflected in the fact that many of Vendler's examples (see below) are given as verb-phrases rather than verbs on their own.

Vendler gives several examples of each of these categories, as follows:

- States: *knowing something, believing something, loving someone*

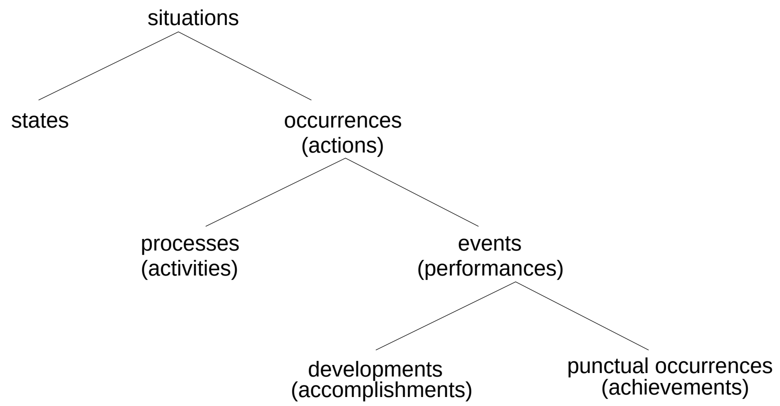

Figure 2: Vendler's classification as remodelled by Mourelatos [57]

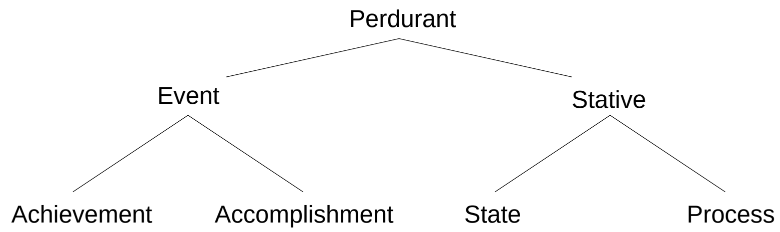

Figure 3: Perdurants in DOLCE

- Activities: *running, pushing a cart*
- Accomplishments: *running a mile, drawing a circle*
- Achievements: *reaching a hilltop, winning a race, recognising something*

It should be noted that all these examples involve a human (or at least animate) subject; Vendler does not explicitly consider other kinds of examples, although the definitions of the four categories does not restrict them to human actions and activities.

One of the problems with Vendler's classification is the ostensible dependence on the forms of expression in English. Indeed, it is not actually true that state and achievement verbs do not possess continuous tenses; what *is* true is that the interpretation of the continuous tenses is different in these cases (e.g., 'He is believing more and more nonsense every day', 'He's winning the race'). A related problem is that States and Achievements do not seem to make a natural pairing. For such reasons, other authors have modified Vendler's classification in various ways to do justice to what they perceive as a more natural interpretation of the facts.

Thus Mourelatos [57] retains Vendler's four categories under a different name, but presents their relationship in the form shown in Figure 2. Here the categories are given "topic-neutral" designations, which do not presuppose the involvement of an agent; the agentive forms of these terms are given in brackets, and correspond to Vendler's terminology.

A different remodelling of Vendler's classification has been incorporated into the DOLCE [51] top-level ontology as subcategories of the class of "perdurants" [for the meaning of which, see below]. This is illustrated in Figure 3.

Influenced by Vendler's system, but differing from it in some significant respects, is the five-fold classification of Moens and Steedman [56], shown in Figure 4. Note that here, so far from being contrasted with events, as in our account above and the Mourelatos and DOLCE classifications, processes

	EVENTS		STATES
	atomic	extended	
+conseq	**CULMINATION** recognise, spot win the race	**CULMINATED** **PROCESS** build a house eat a sandwich	understand love, know resemble
-conseq	**POINT** hiccup, tap, wink	**PROCESS** run, swim, walk, play the piano	

Figure 4: The classification of Moens and Steedman [56]

are subsumed as a subcategory of a rather broader category of events, which covers everything other than states. Note also that, as with Vendler, this is presented as a classification of verbs and verb phrases, although it could equally well be understood as a classification of what these linguistic items are used to refer to in the world. The criteria for the two-by-two cross-classification of events are, in the horizontal dimension, ‘atomic’ vs ‘extended’ (whose meanings are self-evident) and, in the vertical dimension, ‘$\pm$conseq’. This refers to whether the event has a ‘culmination’ after which some consequent state holds (e.g., the existence of a house or the location of the sandwich inside the body).

6 Perspective

By perspective I mean the *point of view* from which some phenomenon is described. This can make a lot of difference to how it is described. Of particular importance is the three-fold distinction between what I shall call ‘participant’, ‘bystander’ and ‘synoptic’ perspectives. I shall illustrate these with respect to two different phenomena.

1. A hurricane.

    - Bystander perspective. This is the point of view of an inhabitant of a village in the path of the hurricane, for whom it is an event which begins when the hurricane first arrives at the village and ends when it moves away from it. The effect of this event is the disruption and destruction it causes to the village.
    - Participant perspective. This is the point of view of the hurricane itself, for which the village is just one of perhaps many different places visited during the period of its existence. From this point of view the hurricane is an ongoing process which affects a sequence of different places.
    - Synoptic perspective. This is the point of view of a weather satellite observing it from several miles above, for which the hurricane is an object which comes into existence at a certain place, moves along a path, and eventually fizzles out.

2. Road traffic.

    - Bystander perspective. This is the point of view of someone standing by the roadside, perhaps looking for an opportunity to cross the road. From this point of view the traffic consists of a succession of vehicles going past. If there are traffic lights nearby, the motion of the vehicles will exhibit a periodicity, with stationary periods interspersed with periods of moving traffic.

- Participant perspective. This is the point of view of the driver of one of the vehicles. From this point of view, one sees a succession of different places along the route, and may experience a sequence of starts and stops as one encounters traffic lights, jams, etc.
- Synoptic perspective. This is the point of view of, say, a police surveillance helicopter ranging over the city and surveying the state of the traffic, locating where the hold-ups are, etc. From this point of view the traffic is like a fluid coursing through a network.

The Bystander perspective is the view of a dynamic spatial phenomenon as seen from a fixed spatial location. The Participant perspective is the view of the phenomenon as seen by a (possibly notional) agent taking part in the phenomenon itself. These perspectives are related to the Eulerian vs Lagrangian approaches to fluid flow, which we discuss in Session 4 (§16.1). The Synoptic perspective is a "bird's-eye view", providing a global description of the phenomenon as a whole. This may reveal features of the phenomenon that are not visible from either the bystander or participant perspectives. Sometimes what appears as an event from the bystander perspective may appear as a process from the participant perspective and as the motion of a discrete object from the synoptic perspective; I have discussed in more detail in [26].

7 Granularity

Granularity has already been mentioned several times. In general terms this is a measure of the level of detail present in a representation. It may be thought of as determined by the resolving power of the "lens" through which the representation views the world, but this should not be interpreted too literally as necessarily referring to something visual, which only captures one aspect of granularity. A crude way of expressing the granularity of a representation is by specifying, for space, the minimal distance which can be distinguished by the representation and, for time, the minimal duration, but again, this fails to capture the more subtle forms of "conceptual granularity".

As a matter of terminology, we speak of *finer* and *coarser* granularities: finer granularity has, as it were, as *smaller* "grain size" and shows *more* detail, whereas coarse granularity has *larger* grain size and shows *less* detail. For this reason it is best to avoid expressions such as "increasing the granularity", as the listener may not be sure whether what is increased is the grain size or the amount of detail.

The effects of varying spatial granularity are most obviously illustrated by considering different map scales. On a large-scale map of a city, for example, individual streets and buildings are shown. At an intermediate scale, the streets may be shown, but not the buildings, the presence of buildings being indicated by a distinctive colouration signifying 'built-up area'. At a small scale, e.g., that of a road atlas, the whole city is represented by a polygon showing the extent of the built-up area, with only the major roads passing through it indicated. At a smaller scale still, the city is just represented by a single symbol, say a dot with a label. At even smaller scales, the city simply doesn't appear at all. Here there is a clear correlation of granularity with scale, but it should be noted that granularity is a more general concept than scale, and does not just apply to maps; also it is perfectly possibly to have two maps at the same scale but with different levels of granularity. The process of reducing the level of detail in a map, and thereby coarsening the granularity, is known as *generalisation*.

An important feature to notice here is that at different granularities the same spatial phenomenon may appear as an object (the city), a collection of objects (the buildings), or a texture (the built-up area) representing a spatial field (see below, §11). A similar thing happens in the case of temporal granularity, where in this case we may find that one and the same temporal phenomenon appears as an event, a sequence of events, or a process.

To illustrate this, consider a situation where a coastal cliff is being eroded, with the result that over a long period of time the coastline moves steadily in the landward direction. This may be understood as a process which is in continuous operation over many years. If we observe it at a finer level of granularity, however, we see that it actually consists of a sequence of discrete, localised cliff falls, each of which is an individual event. Look more closely still, and each of these cliff falls can be seen to consist of a variety of processes: perhaps some undercutting of the bottom of the cliff face, followed by a cleavage developing between the resultant overhang and the main body of the cliff, and then the tumbling of rock and earth down onto the beach — the exact details will vary with the nature of the local geology. Within any of these processes, many smaller-scale events occur, e.g., the fallings of individual chunks of rock.

An important point to make here is that in evaluating any model one must take account of the level of granularity presupposed by the model. Most statements about any phenomenon only apply to it over a range of granularities, and ideally one wants a more complex representational framework within which one can vary the granularity in order to represent different aspects of the phenomenon under study. One often hears remarks to the effect that such-and-such a phenomenon is "not an event but a process"; a more careful statement would be that what appears at one level of granularity as an event can be seen at a finer granularity to consist of various processes. Thus in some sense it is *both* an event and a process, but not in the same view.

8 Varieties of process

Within the category of processes we may distinguish a number of broad types which differ in their characteristic temporal "profile". Which type a given process is assigned to will generally be dependent on the level of granularity at which it is being described, as discussed below. The four main types (as presented in [24], pp.207–211) are *steady-state*, *cumulative*, *periodic*, and *irregular*.

- A process may be described as steady-state if it is being described in a way that abstracts away from any change that occurs as a result of the process. A candle flame, considered as a burning process, and viewed at a spatio-temporal granularity too coarse to register any movements in the flame resulting from local air currents, may be regarded as steady-state process. Another example is the flow of a river under a bridge.

- A process is described as cumulative if it results in the steady increase or decrease of some observable quantity or measurement. Any linear or approximately linear motion may be regarded as a cumulative process in that as time goes by, the moving object is displaced ever further from its starting point. Similarly the process of writing is cumulative as more written material accumulates while it is in progress, and so is the process of eating, as more and more food accumulates inside the body. Cumulative processes may (but need not) be self-limiting: that is, they may have built-in endpoints beyond which they cannot be continued further.[4] An example is the process of pouring water into a non-leaking container; this is cumulative because as it progresses the amount of water in the container steadily increases, but it cannot continue once the container is full (or rather, if one continues pouring water into the container, an overflowing process begins, and what was cumulative becomes steady-state).

- A periodic process is one in which a particular sequence of phases is repeated over and over again, the repetitions being identical or qualitatively similar. This category includes all rotations and oscillatory processes, e.g., both the rotation of the earth about its axis and the revolution of the earth about the sun, as well as many processes which keep step with these, e.g., the seasonal alternation of growth and dormancy in vegetation. It also covers the motion of the legs when someone is walking, the sequence of gyrations made by the hand of someone writing, and a process such as hammering which consists of a more or less regular sequence of hammer blows.[5] A periodic process appears steady-state when viewed at too coarse a granularity to distinguish

[4] This is known in the AI literature are *auto-termination*.

[5] Cf. the remarks about such processes being "made of" events, above.

the individual cycles that make it up: a good example is the shining of a fluorescent light, which appears continuous — i.e., steady-state — to humans but in reality flickers on and off in step with the alternating current of the electricity supply — the “flicker fusion rate” of the human eye imposing too coarse a granularity to detect this.

- Periodic processes shade off into irregular processes as either the periodicity or the qualitative nature of the repetition becomes less uniform. An example is radioactive decay, where although the rate of decay follows a precisely predictable trend, the pattern of decay events at the level of individual atoms is unpredictable and irregular. (In the long run, of course, radiocative decay is a cumulative process.) The flow of traffic along a road is often irregular.

SESSION 2: Handling processes and events in GIScience

9 How Time enters GI Science

This section is based on the "Brief History of Time" presented by Worboys [91, 89]. Here the incorporation of time and temporal phenomena into GIS is presented as a sequence of stages in the course of which the temporal dimension becomes increasingly fully integrated into what starts off as a purely spatial information system.

9.1 Stage Zero: Static GIS

This is the condition of traditional spatial and geographic information systems. The information contained in the system is not time-indexed at all; most usually it implicitly represents "the present", though whether this means the time the data was collected, the time it was incorporated in the system, or the time the user accesses it, may be left unstated. This approach is only really viable when the information in question relates to features of the world which do not change, or which change only slowly, so that it is equally valid at all three times mentioned above. At Stage Zero there can be no representation of change and hence no representation of processes and events either.

9.2 Stage One: Snapshots

If we treat a representation at Stage Zero as presenting a 'snapshot' of the state of the world at a single moment of time, then the simplest way of taking account of the passage of time is to present a *sequence* of snapshots, each indexed by a different time. An example is shown in Figure 5, which shows four maps of Exeter (a city in the south-west of England), in the years 1809, 1898, 1946, and 2006.

Inspection of the maps in Figure 5 reveals a number of changes that have taken place, for example:

- Between 1809 and 1898, Exeter expanded considerably, absorbing the outlying villages of St Thomas, St Sidwell, and Heavitree. Moreover, during this period three railway lines were constructed, the Great Western Railway following the course of the River Exe, the London and South Western Railway coming in to Exeter from the east, with its branch line to Exmouth.

- Between 1898 and 1946, an additional railway line was constructed, branching to west from the Great Western Railway past Alphington. The city continued to expand, absorbing Whipton.

- Between 1946 and 2006, the city expanded still more, incorporating Pinhoe, Alphington, and Countess Wear, the branch line of the Great Western Railway was taken out of use, and some new major roads were built, notably the M5 motorway to the east of the city and another major road to the west.

It is important to note that these changes are *inferred* from the snapshots. Because they are not explicitly represented, there is no way to tell exactly when they occurred during the periods to which we have assigned them — e.g., did the branch-line disappear before or after the construction of the motorway? Moreover, we cannot tell how long the changes took, or whether they occurred in a continuous manner or in one or more discrete steps.

To some extent these problems could be overcome by interpolating a larger number of snapshots. If we had a snapshot for every year between 1809 and 2006 we could build up a much more complete picture of the exact sequence of changes undergone. The storage implications of this are considerable, of course, and in many ways a more satisfactory solution is not to store the sequence of snapshots themselves but only an initial snapshot and then the *changes* that take place. This means that any feature that remains the same over a long period only needs to be recorded once: the fact that it remains the

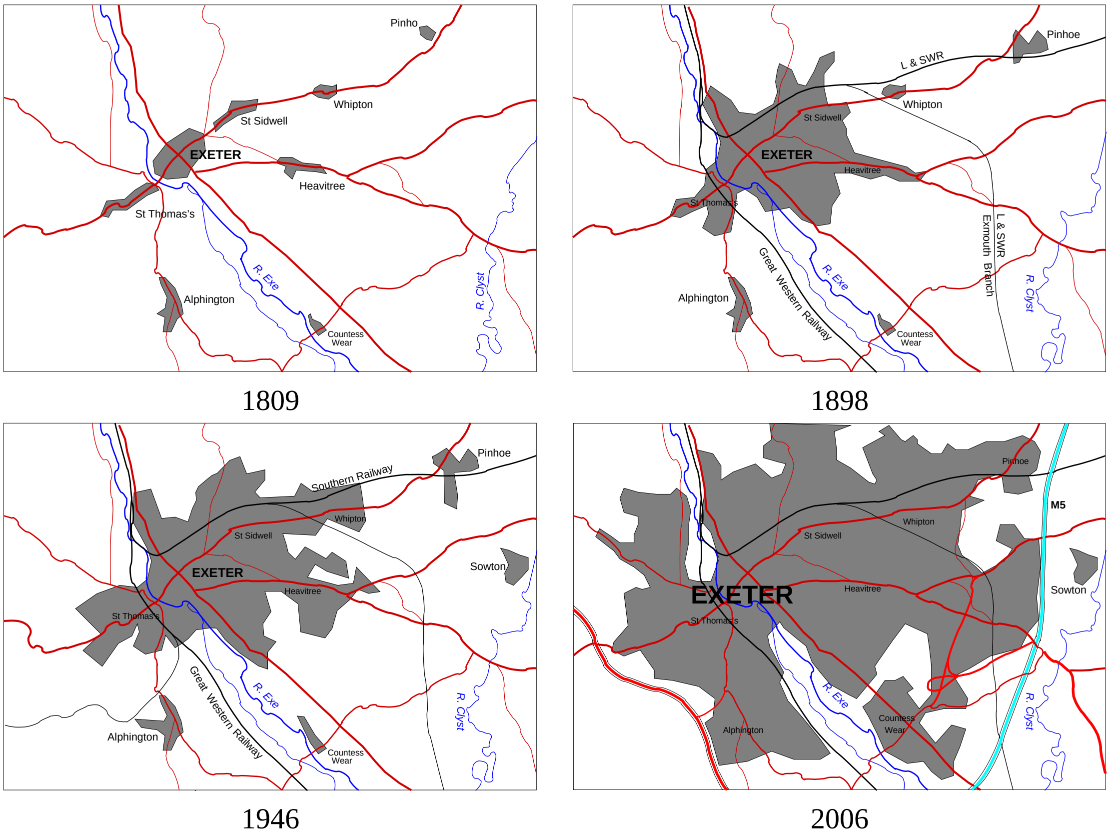

Figure 5: “Snapshots” of Exeter, UK (Loosely based on maps provided by the Ordnance Survey)

same at later times can be inferred from the fact that there is no record of any changes to it. This idea leads naturally on to:

9.3 Stage Two: Object change

At this stage, the focus is on the life-histories of individual objects and features in the area of interest, rather than on snapshots. In relation to our sequence of snapshots of Exeter, for example, the facts concerning the development of the railways shown there are as follows:

- The Bristol & Exeter Railway, entering Exeter from the north along the Exe valley, was opened in 1844. The South Devon Railway, continuing the line southwards along the Exe valley, was opened in 1849. These were both amalgamated within the Great Western Railway in 1976.
- The London & South Western Railway reached Exeter from the east in 1860, and opened it branch line to Exmouth in 1861. They became incorporated into the Southern Railway in 1923.
- The branch line to the west, past Alphington, was opened as an extension of the already existing Teign Valley Line in 1903. This was subsequently closed in 1958.
- All the surviving railways became part of the nationalised British Railways (later called British Rail) in 1948. Subsequently, during 1994–1997, this became privatised again, with track and rail services operated by different companies.

This kind of information is not easy — even if possible — to convey using snapshots alone. A snapshot cannot really portray an event: at best, an event can be implied by a pair of snapshots portraying the 'before' and 'after' states.

If we have in addition the following information:

- The M5 motorway was extended as far as Exeter in 1977

then we can readily answer the question posed above: the branch-line was closed 19 years before the motorway was opened.

The idea of object change presupposes a logically prior idea of *object identity*. We may distinguish two kinds of change: changes which affect the attributes of objects without altering the objects' identities, and changes which affect the identities of the objects involved.

- When Hawaii became the 50th state of the USA in 1959, the identity of the USA was not affected — that is, the USA continued to exist, but as a 50-member union rather than a 49-member one. Hawaii itself likewise continued to exist, this being the latest in a series of changes of status it had undergone since its formation as a single unit in 1810 (first as a Kingdom, then in 1894 it became a Republic, and four years later a US Territory). However momentous any of these changes may have been, we can see them as exemplifying the paradigm of a object maintaining its identity through a sequence of changes to its attributes.[6]

- The division of Czechoslovakia into the Czech and Slovak republics at the beginning of 1993 is an example of a change that affects identity. In fact Czechoslovakia as a nation-state ceased to exist after that date, and the Czech and Slovak republics came into existence, with neither having a claim to retaining the identity of the former Czechoslovakia. On the other hand, with the re-unification of Germany in 1990, the resulting unified state retained the title and, more importantly, many of the institutions of the Federal Republic of Germany, a term which had previously applied to 'West Germany'; thus this is probably best described as a case of the Federal Republic retaining its identity while absorbing within itself the former German Democratic Republic (i.e., 'East Germany'), which thereby ceased to exist as a geographical entity.

With these geo-political examples, much hangs on the various formal constitutions, treaties, etc, by which they are governed, and in many cases there will be explicit agreements that such-and-such an entity ceases to exist or comes into existence or becomes part of some other entity at a given time. With physical geographical features it may be much less clearcut. An example that is often discussed concerns a lake that divides into two water-bodies as the water table lowers. In principle there are several different ways of describing what is going on here:

- The original lake continues to exist as a single entity with two disjoint parts.

- The original lake continues to exist and is to be identified with the larger of the two descendent water-bodies, the other being a new lake which has been "spawned" by the old one.

- The original lake no longer exists; it has been replaced by two new lakes occupying part of the area of the old one.

Usually in cases like this there are no hard facts one can appeal to to decide which description is correct. One might use the names given to the various water-bodies in question as an indication, but the truth is that there is nothing in nature to decide the matter, and it becomes a matter of possibly arbitrary convention. This means that many geographical concepts such as "lake", "mountain" and "forest" cannot

[6]Here it should be stressed that we are using the word 'identity' in a rather neutral philosophical sense, not in the more psychogical sense associated with such phrases as "identity crisis" or "sense of identity".

be regarded as natural kinds — which in turn means that there may be no definite answer to a question such as “How many lakes are there in Finland?”.[7]

The kinds of changes that may affect the identities of their participant objects include such things as creation, destruction, reincarnation, aggregation, disaggregation, fusion, and fission, formal definitions of which are given by Medak [52]. Such changes have been studied systematically under the heading “identity-based change” by Hornsby and Egenhoger [39, 40, 41].

9.4 Stage Three: Events, Processes and Actions

The changes handled at Stage Two are all given as parts of the life-histories of individual objects. But not everything that happens is conveniently described in this way. Stage Three introduces what Worboys calls “a full-blooded treatment of change”, in which events, processes and actions are accorded first-class status as fully-fledged members of our ontology. This enables us to handle, for example, phenomena such as migrations, wildfires, hurricanes, epidemics, traffic congestion, battles, coastal erosion, and many others, which while they may involve many different individual objects may be best viewed as objects of study in their own right rather than simply an amalgamation of episodes from individual life-histories. An example of this approach is the geospatial event model (GEM) of [88], which is described further in §15.2.

At this stage we can start to describe the network of *causal* relations between events, processes and states, thus providing the basis for explanation and prediction, which are increasingly seen as desiderata for the next generation of geographical information systems. As so often in this field, the terminology for the various kinds of causal relation is confused, and in particular there is some confusion about what kind of thing can cause what: e.g., can an event cause a state, or vice versa, or is causation strictly a relation between events? Mike Worboys and I [31] have attempted to clear up the confusion by introducing a systematic set of terms each with a clearly defined role, as follows (here E is an event and S is a state — which could include the state of some process being active):

- ‘E_1 causes E_2’ means that E_2 occurs as a result of E_1’s occurring.
- ‘E initiates S’ means that if E occurs at t then (1) S must hold over an interval beginning at t, and (2) there is an interval ending at t throughout which S does not hold.
- ‘E terminates S’ means that if E occurs at t then (1) S holds over an interval ending at t, and (2) there is an interval beginning at t throughout which S does not hold.
- ‘S_1 enables S_2’ means that if S_1 holds at t then it is possible for S_2 to hold at t.
- ‘S_1 disables S_2’ means that if S_1 holds at t then it is *not* possible for S_2 to hold at t.
- ‘S_1 perpetuates S_2’ means that if S_1 and S_2 both hold at t, and S_1 holds throughout an interval i which starts at t, then S_2 will also hold throughout i.
- ‘S allows E’ means that if S holds at t then it is possible for E to occur at t.
- ‘S prevents E’ means that if S holds at t then it is not possible for E to occur at t.

These relations are portrayed diagrammatically in Figure 6.

The problem facing the developer of a Stage Three system is how best to describe, represent, reason about, and manipulate these events and processes, and how they should be related to the objects which participate in or are affected by them. In the remainder of this session we shall look at a number of different approaches which have been proposed for doing just this.

[7] According to `http://www.ymparisto.fi/default.asp?node=12335&lan=en` (accessed 15/7/09), the answer is 187,888 — this assumes that a lake is a body of standing water larger than 500 m^2, but that of course is just one definition amongst many that are possible.

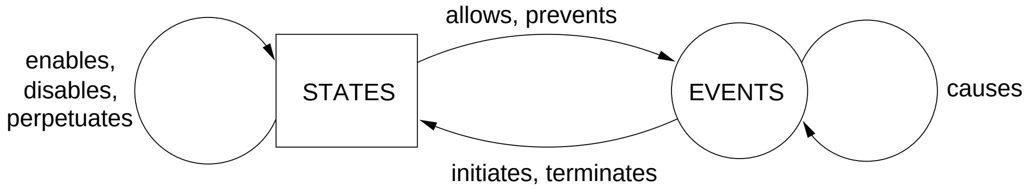

Figure 6: Kinds of causal relationships between states and events (from [31])

10 The Ontology of Continuants and Occurrents

10.1 Three-dimensionalism vs Four-dimensionalism

How should we conceive of the world? Is is a three-dimensional space whose three-dimensional contents come into existence, move about, change, and pass away as time passes? Or is it a four-dimensional space-time whose four-dimensional contents are fixed once and for all? Common sense seems to suggest the former picture, but a number of developments over the past 100 years have exerted pressure on us to move towards the latter.

The original impetus for these developments came from the Special Theory of Relativity, according to which the spatial and temporal separations between two events are not absolute but depend on the state of motion of the observer; in their place, relativity posits the *spatio-temporal* separation of the events as independent of the observer. The idea is that any observer will resolve this absolute spatio-temporal separation into a spatial component and a temporal component, but this will lead to different results when done by observers who are moving relative to each other. This is illustrated using the Minkowski diagram, introduced in 1908, in which positions in space-time are plotted on a two-dimensional graph in which one axis (with gradient less than 1) represents all spatial dimensions bundled together and the other (with gradient greater than 1) represents time. An example is shown in Figure 7, where the two Minkowski diagrams show the coordinate systems corresponding to observers in relative motion; the red and blue dots represent two events. Their spatio-temporal separation, represented by their actual distance apart on the page, is the same in both diagrams, but their spatial and temporal separations differ for the two observers.

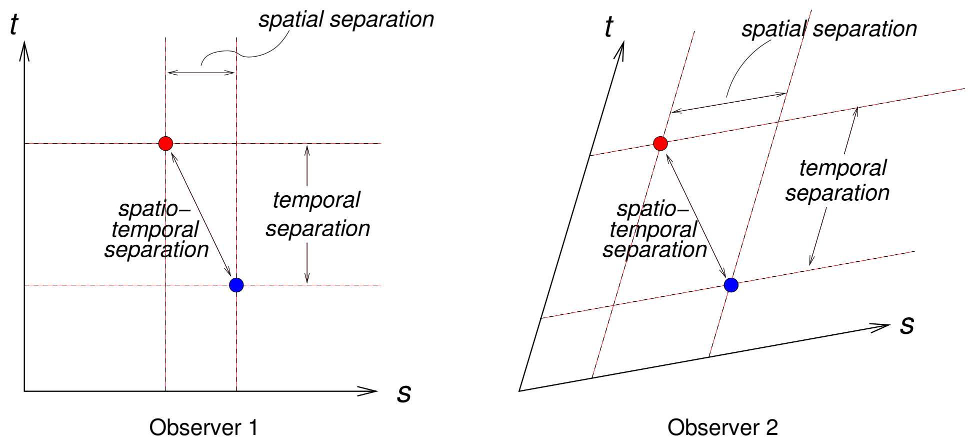

Figure 7: Minkowski diagrams, showing relative spatial and temporal separations, but absolute spatio-temporal separations, between a pair of events measured by observers in different states of motion.

The Theory of Relativity does indeed call into question many of our everyday ("commonsense") notions of space, time, motion, and other physical concepts, and because of this has exerted a considerable influence on researchers even in disciplines remote from physics who are concerned in any way with space and time — Geography being a notable case in point. Arguably, however, the relativistic effects shown in Figure 7 are irrelevant to geography, since these effects do not become apparent until the relative motions in question are appreciable fractions of the speed of light — in other words much faster than any speeds encountered by human beings when going about their business on or near the surface of the Earth.

Despite this, the general idea of portraying space and time together in a single diagram can be very useful and informative even in situations where we do not need to worry about the relativistic effects (i.e., in situations where the orientation of the axes is treated as unvarying). This is illustrated by the *time geography* diagrams of Hägerstrand [35], illustrated in Figure 8. Here, at the left of the figure, a space-time diagram is used to visualise the motions of three objects; on the right are shown three "snapshots" corresponding to the times marked by the three planes orthogonal to the time-axis in the left-hand diagram.

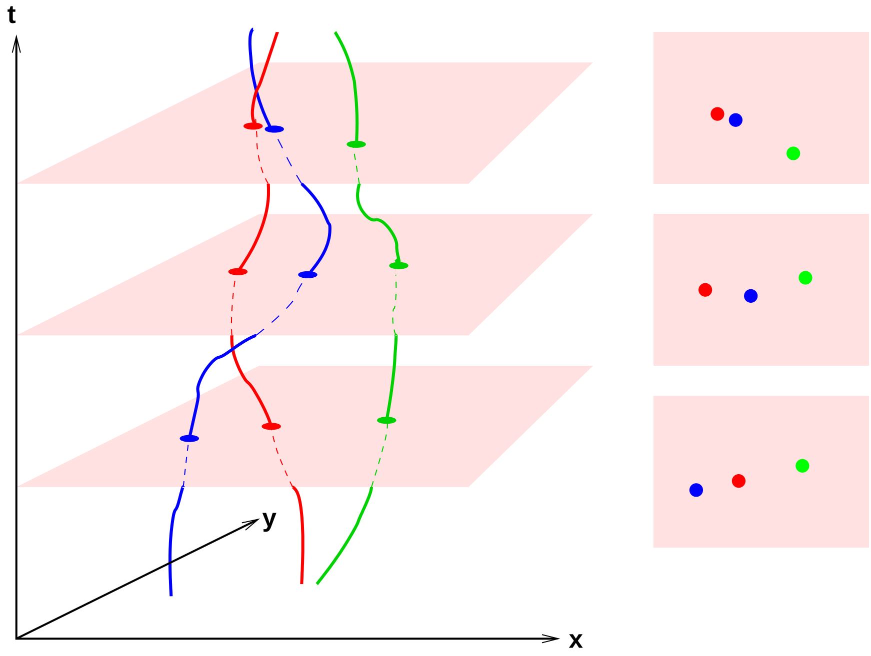

Figure 8: Lifelines of three moving objects, with snaphots

Such diagrams suggest an alternative way of viewing change and motion. Ordinarily, we think of objects as moving through space, that is, occupying different positions at different times, as represented by the sequence of snapshots on the right-hand side of Figure 8. On the alternative view, an object's motion is represented by its "life-line" in space-time. Although it is tempting to think of the object as "moving along" or "tracing out" its lifeline, this is highly misleading. The only real motion an object does is through space. It exhibits this motion by occupying different positions at different times. The life-line collects together all the different (time,position) pairs thereby determined, in effect treating time as another spatial dimension. The term *geospatial lifeline* has been used [42] for the lifeline of a moving object in geographical space, i.e., when considered at a sufficiently coarse granularity for purely local, non-geographical movements to be invisible.

From here it is a short step to the fully-fledged four-dimensional view, according to which an object is identified with its life-line. On this view, what you see standing in front of you is not a complete person: rather, it is merely a cross-section, orthogonal to the time-axis, of a complete person. On this view, a person not only has *spatial parts* (arms, legs, etc) but also *temporal parts* (childhood, adulthood, etc). In fact, one can no longer clearly distinguish objects from events. Both are simply the contents of selected regions of space-time.

It is not always appreciated how radical a departure from our everyday ways of thinking and speaking this four-dimensional view is. Most of our everyday forms of expression are adapted to the three-dimensional view. If we want to espouse four-dimensionalism, we must either give up our ordinary language or radically reinterpret it. For example, "John became taller", instead of saying that the object named John changed from being less tall to being taller, must be understood as saying that one cross-section of John is taller than another one with smaller temporal coordinate. Talk of one three-dimensional thing changing is translated into talk of different three-dimensional parts of some four-dimensional thing having different properties. Of course, we can still *call* this change: and the challenge for four-dimensionalism is to show that all our everyday temporal language (at least to the extent that it is, in fact, coherent) can be reconceptualised in the four-dimensional fashion without loss of meaning.

Authors sympathetic to, or actively advocating the four-dimensionalist point of view include philosophers [38, 75], computer scientists [58, 36], and GI scientists [68].

10.2 Continuants and Occurrents

On the three-dimensional view[8] it is common to draw a radical distinction between the categories of *continuant* and *occurrent*:[9]

- A continuant is an entity which endures through time. At any given time during its life it is said to exist. The parts of a continuant are coexisting continuants occupying distinct spatial regions at any given time. At different times the same continuant may have different properties, and therefore the continuant may undergo change.

- An occurrent is an entity which extends through time. It is said to occur (not exist). At different times, different parts of the occurrent are present. An occurrent is not said to undergo change.

Continuants include most of the things we would normally categorise as "objects": people, tables, chairs, buildings, roads, mountains, clouds, planets, etc. The clearest examples of occurrents are *events*, e.g., the assassination of Julius Caesar, the Second World War, a person's life, an earthquake, a volcanic eruption. There is broad agreement about these amongst those who espouse a three-dimensionalist view. With states and processes, on the other hand, the issue is less clearcut. For example, to consider just two popular modern ontologies, states are regarded as continuants in BFO [33] but as occurrents in DOLCE [51]. Regarding processes, although these are generally treated as occurrents, I have advanced the view that they are in many respects more like continuants [27, 29].

The distinction between continuants and occurrents originated in philosophy, but has been championed as an important distinction in the context of information system ontologies, where it has been enshrined in the distinction between SNAP and SPAN ontologies, the former comprising continuants, the latter occurrents.[33, 34]. These form part of the Basic Formal Ontology (BFO) mentioned above.

11 Fields and Objects

Another distinction that is important when working with geographical information is that between *field-based* and *object-based* approaches. This distinction arose as a result of the coexistence of two different approaches to recording and displaying geographical data in the early days of GIS, using *raster* or *vector* formats.

[8] Which should really be called "three-plus-one"-dimensionalist, since it distinguishes the three dimensions of space from the single dimension of time.

[9] The terms *endurant* and *perdurant* are also used.

In the raster-based approach, the space of interest is divided into *cells* (or — in the context of an image used to display the data, pixels), usually, though not necessarily, forming a rectangular grid, and the data takes the form of *values* assigned to each of the cells. A common example is an *elevation field*, where to each cell of the grid is assigned a numerical value representing the height above sea-level of the patch of ground represented by that cell.

By contrast, in the vector-based approach, the data comprises records of individual features of interest — e.g., trees, rivers, buildings, towns, coastlines — each of which is represented by a geometrical element such as a point, polyline, or polygon, specified by listing the coordinates of a sufficient set of its constituent points (e.g., the vertices of a polygon).

Although it has often been presented as a debate as to the relative superiority of the two approaches, it is generally recognised that for some types of data one approach may be more appropriate, and for other types the other one. More significantly, it also became apparent that to present the distinction in the low-level, implementation-oriented terms of raster *vs* vector obscured a much more fundamental dichotomy which relates to higher-level concerns about the conceptualisations we use for organising our understanding of the world: this is what has become the distinction between the field-based and object-based approaches [63, 15].

In the field-based approach, the world is conceived of as a collection of *fields*, i.e., functions associating values with locations. Examples are the elevation field, which associates a height above sea-level with each point on the earth's surface, a land-cover field, associating with each point a value recording the nature of the surface at that point (e.g., type of vegetation, soil type, whether built on, etc), a land-use field, and fields giving seasonal average rainfall and temperature values. On maps, such fields are typically represented using a range of colours to represent different value ranges.

In the object-based approach, on the other hand, the world is conceived of as a collection of discrete individual entities, each of which has a particular size, shape, and location — and of course other, non-spatial attributes can be associated with such entities as well. An example would be a collection of data on cities, where each city is represented (using a small scale) as a point or (using a larger scale) polygon, to which are attached such non-spatial data as the name of the city, its population, its administrative status, etc. Note that it would be hard to represent this as a field: one could have an "urbanisation" field, which assigns to each point a value to indicate whether or not that point falls within the boundary of a city, but there is no straightforward way of marking off the cities as distinct individual entities.[10]

Although this discussion so far has been entirely spatial, things become more complicated when time is brought into the picture [25]. Time is like space in that it provides a tightly organised collection of elements which serve as *indices* for locating thematic elements — thereby forming part of the *Place – Time – Theme* triad regarded as fundamental to geospatial representation [64]. As with space, we can approach temporal representation either in terms of fields or in terms of entities:

- A temporal field is a function from times to values. This is what is known in the Artificial Intelligence literature as a *fluent*. A typical example is a record of how the temperature (rainfall, wind speed, wind direction, air pressure, ...) measured at a particular weather station varies over the course of a day (week/month/year/...). In these cases the values are numerical, but a temporal field can also take qualitative values (e.g., states of alertness in an emergency response system), and in the limiting case these are just two, i.e., the Boolean values 'true' and 'false'. In this case we are simply tracking the truth-value of some proposition as it varies over time (e.g., whether a mountain pass is open or closed). Temporal fluents record *states*, either the state of an object or spatial field, or the state of some ongoing process.

[10] In particular, the obvious solution, to identify the cities as maximal connected sets of points marked as urban, would not work for a case such as Leeds and Bradford in the north of England, which are distinct cities but are not separated by any tract of non-urban land.

- A temporal entity, on the other hand is a discrete individual chunk of history to which a label or identity has been attached: in other words, it is an event. The entity-based approach to recording temporal information will maintain a list of events, each annotated with salient properties such as its location in space, its time of occurrence, perhaps the objects (continuants) participating in it, and its causal relations with other events.

As in the spatial case, temporal fields are in some way more ‘fundamental’ or ‘primitive’ than temporal entities: they will appear earlier in the sequence of stages identified by Worboys, and indeed are already at least implicitly present in the ‘sequences of snapshots’ model.

Field-based approaches are essentially ‘local’ in the sense that each item of information recorded in such an approach pertains to one place at one time. The power of such an approach comes from applying analytical tools to large collections of such individual items of information, e.g., to discern large-scale patterns in the spatio-temporal distribution of the field values. Entity-based approaches can have a more global reach, in the sense that there is no limit, in principle, to the spatial or temporal extent of the entities represented — hence the Gobi desert, the Amazon river, or the Second World War can be treated as individual entities, whereas in a field-based representation individual locations may be marked as desert, or water, or at war, without there being any way of bringing together a specific collection of such locations to identify as a single entity.

SESSION 3: Formal treatments

A *formal system* is a system for representing facts about some chosen domain in such a way as to allow those facts to be manipulated using rules that refer only to the *form* of the representations, and not their meaning.[11] As an example, *arithmetic* is a formal system: once all the facts about how much the employees of a company earn, how much tax, insurance, etc, they have to pay, and so on, have been entered into the system, all the subsequent manipulation of these facts can be handled by purely formal 'number-crunching', i.e., using the rules of arithmetic, which apply to numerical representations generally, irrespective of whether they refer to amounts of money, meteorological variables, populations, or anything else. Of course, once the calculations have been performed, the end result has to be interpretated again in terms of money, so that the employees can actually be paid.

The key to a successful formalisation is to design the formal rules so that they deliver results in accordance with what is expected in the domains to which the formal system is applied. It is always relevant to ask, in relation to any particular domain, whether or not, or to what extent, it is possible to capture the key relations and regularities of that domain within a formal system. In general terms, the aspects of domains to which formalisation can be most successfully applied are the highly general 'topic-neutral' aspects such as number, time, and space, which exhibit their characteristic regularities over a wide range of more specific domains within which they play a part.

It should be recognised, though, that representing and manipulating information in electronic form on a computer is always, strictly considered, an act of formalisation, and therefore what we can do with computers must be limited to what can be formalised. The drive to imbue computers with 'semantic' capabilities is, properly considered, a quest for increasingly sophisticated formalisations, in which more and more elements of some domain of intended application are brought within the sphere of purely syntactic manipulations.[12]

12 Temporal Logic: modal vs first-order approaches

There are two main approaches to representing temporality in a logical calculus. These are called the *modal* approach and the *first-order* approach; they differ according to the logical form in which temporal information is encoded.

12.1 The modal approach

On the *modal* approach, temporal information is expressed by means of *operators*, that is, expressions which may be applied to existing propositions to form new propositions. In *Tense Logic* [66, 67], the temporal operators have a function analogous to the tenses in natural language. Thus if the formula r is used to express the proposition "It is cold", tense operators can be applied to modify the meaning as follows:

Formula	Interpretation
r	It is cold
Pr	It was cold, it has been cold
Fr	It will be cold
Hr	It has always been cold
Gr	It will always be cold

[11] In other words, by *syntactic* manipulation that does not refer to the intended *semantics*.

[12] I am aware that these marks may be contentious; they are in line with the well-known arguments by John Searle concerning the impossibility of generating true meaning and understanding by means of digital computation alone [73]. The premise of 'Strong AI' is essentially that *all* semantics is ultimately nothing but fancy syntax, and hence can be achieved using formal symbol manipulation.

Operators can be combined, e.g., HFr "It has always been going to be cold", FPr "It will have been cold", and such combinations occur in formulae that may be posited as axioms to serve as a basis for reasoning. An example is

$$p \rightarrow GPp$$

which states that if a proposition p is true, then it will always have been true (expressing the idea that what's done cannot be undone).[13]

For practical purposes, this simple form of temporal logic is extremely limited, since much of the temporal information we want to record cannot be expressed purely in terms of tenses. We do not just want to say that something has been or will be the case; we want to say something about *when* it is the case. For this purpose we must find some way of referring to individual times within a logical system. One possibility is to stay within the modal approach, but introduce a special set of propositions which are used to assert what the time is. Thus if t is used as a proposition meaning "It is 12th July 2009", then the proposition $\Diamond(t \wedge r)$ (where $\Diamond p$ is an abbreviation for $Pp \vee p \vee Fp$ — i.e., "it has been, is, or will be the case that p") asserts that it rained on that day. Written in the form $@_t r$, this is a fundamental notation in *hybrid logic*, which seeks to exploit some of the advantages of both the modal and the first-order approaches (see `http://hylo.loria.fr/`).

12.2 The first-order approach

When reference to individual times is required, it is more usual to adopt the second kind of approach to temporal logic, the *first-order* approach. This uses the first-order predicate calculus (FOPC) as its base notation. An *atomic formula* in FOPC has the form $P(a_1, a_2, \ldots, a_n)$, where $a_1, \ldots, a_n$ are *terms* denoting individual entities, and P is a *predicate* expressing some relation that holds amongst individual entities taken n at a time. The formula asserts that this relation does indeed hold for the n entities picked out by the a_i. The terms a_i are called *arguments* of the predicate.

In first-order approaches to time, it is assumed that times are individual entities which can be referred to using terms, and these terms can then be used as arguments to predicates. For example, to express the assertion that it rained on 12th July 2009, we might write something like $Rain(day_{12-07-2009})$, and to say that Napoleon invaded Russia in 1812, we could write $Invade(napoleon, russia, year_{1812})$.[14] This is known as the *Method of Temporal Arguments*. Some other first-order approaches will be described below.

The method of temporal arguments does not discriminate between processes and events. In both the formulae $Rain(day_{12-07-2009})$ and $Invade(napoleon, russia, year_{1812})$, what is referred to might be either an event or a process. If $Rain$ is understood to refer to a shower of rain, i.e., a bounded episode with a definite beginning and ending, then $Rain(day_{12-07-2009})$ would naturally be understood as reporting that such an event occurred on 12th July 2009; whereas if $Rain$ is understood to refer to the process of raining as it might be realised at any given moment, then $Rain(day_{12-07-2009})$ would report that this process was in operation on that date. Similar remarks apply to the Napoleon example, where the contrast is between an invasion event, and the process of invading. An important difference between the event and process interpretations is that whereas, for an event, there must be a minimal interval over which the event occurs, such that the event cannot be said to occur over any proper subinterval of that interval, a process is typically understood to be in operation at each moment, and hence throughout each subinterval, of any interval over which it is in operation.

[13]Familiarity with basic logical notation is assumed. Here is a quick reminder. Where a and b are formulae expressing propositions, $a \wedge b$ is true just in case a and b are both true, $a \vee b$ is true when at least one of a and b is true, $a \rightarrow b$ is true unless a is true and b is false, $a \leftrightarrow b$ is true so long as a and b are either both true or both false, and $\neg a$ is true if and only if a is false. Thus these formulae express something like "a and b", "a or b", "If a then b", "a if and only if b, and "It is not the case that a", respectively.

[14]We adopt the convention that predicates are written with an initial capital, whereas terms are written with an initial lower-case letter.

In order to allow the expression of such distinctions, many researchers have adopted a different way of using first-order logic to express temporal facts. This uses a device known as *reification*.[15] Whereas in the method of temporal arguments the predicate expresses both the event and the fact of its occurrence, in a reified system these are separated, with the event expressed by means of a term, and the fact of its occurrence by a predicate. Thus instead of

$$Invade(napoleon, russia, year_{1812}),$$

one might write

$$Occurs(invade(napoleon, russia), year_{1812}).$$

Here $invade(napoleon, russia)$ is a term denoting an event type (see §4), and the general form $Occurs(E, t)$ states that an event of type E occurs at time t. Because the term refers to an event type, this is known as *type-reification*. This method was used by Allen in [2], a paper which has been extremely influential in AI.

Contrasted with type-reification, there is also *token-reification*, in which terms are used to refer to event tokens rather than event types. Under this system, the statement that Napoleon invaded Russia in 1812 would appear as[16]

$$\exists e(Invade(napoleon, russia, e) \wedge Occurs(e, year_{1812})).$$

This says that there is an event token e which is an invasion of Russia by Napoleon and which occurs in the year 1812. This method was introduced in Philosophy by Davidson [17], and in Artificial Intelligence by Kowalski and Sergot [43], where it forms the basis of a system known as the Event Calculus.

On any of these methods there is potential ambiguity as to how the occurrence of an event is related to the time at which it is said to occur. When we say that Napoleon invaded Russia in 1812, we mean that the temporal extent of the invasion event is an interval which falls wholly within the named interval 1812. However, the $Occurs$ predicate is usually interpreted differently, so that $Occurs(e, i)$ means that i is exactly the interval over which E occurs, and not any larger interval. This interpretation is secured by an axiom such as

$$\forall e \forall i \forall i'(Occurs(e, i) \wedge i' \sqsubset i \rightarrow \neg Occurs(e, i')),$$

where $i' \sqsubset i$ means that i' is a proper subinterval of i. Thus this says that if e occurs on interval i, it cannot occur on any proper subinterval i' of i. If all we want to say is that e occurs *during* i, without insisting that the occurrence takes up the whole of the interval, then we could write

$$\exists i'(i' \sqsubseteq i \wedge Occurs(e, i')),$$

which could, if desired, be abbreviated $OccursIn(e, i)$.

One cannot go very far with this kind of analysis without developing a precise formalisation of time itself: that is, of instants and intervals and the various relations between them. This has been a thoroughly researched area, which I shall briefly describe in the next section.

[15] For a more detailed discussion of the various forms of reification, see [28].

[16] In first-order logic, a formula of the form $\exists xP$ states that there is *at least one* element x such that P holds. The symbol $\exists$ is called the *existential quantifier*. There is also a *universal quantifier* $\forall$, where a formula of the form $\forall xP$ states that P holds for *every* element x.

13 Time itself: Formal properties of instants and intervals

Our normal time-measurement practices include reference to both *instants* and *intervals*. Whereas an interval has a certain duration, instants are regarded as durationless: they may be thought of as the meeting-points of contiguous intervals. When we say that something happened at, say 2.45 p.m. exactly, then, at least in idealisation, we are pinpointing an instant. More often, we locate something in time by naming an interval during which it occurs, as in the example in §12.2: Napolean invaded Russia in 1812. We have names for many "standard" intervals such as years (1812), months (June 1812), days (24th June 1812), but we can also designate an arbitrary interval by naming its beginning and end points (e.g., from 4 p.m. to 5.30 p.m on 24th June 1812), or in terms of some event or state of affairs which marks the interval (e.g., during the reign of Louis XIV).

In the following subsections I shall discuss the logical formalisation of instants and intervals.[17]

13.1 The logic of instants

The standard mathematical model for temporal instants is provided by the real number system. If some instant is designated as the origin, to be assigned the real number 0, then any other instant can be assigned a real number which measures the number of seconds before or after the instant 0.[18] If we abstract all quantitative considerations away from the real number model, we are left with the following set of assertions about the *qualitative* nature of the temporal ordering. These are expressed in terms of a 'precedence' relation taken as the primitive basis for the ordering; in the logical notation used below, $t \prec t'$ states that instant t precedes (i.e., is earlier than) instant t'. We shall use the terms *predecessor*/*successor* to refer to an instant which precedes/follows a given instant.

1. *Irreflexivity*. No instant precedes itself:
$$\forall t(\neg t \prec t)$$
2. *Transitivity*. A successor of a successor of t is a successor of t:
$$\forall t \forall t' \forall t''(t \prec t' \wedge t' \prec t'' \rightarrow t \prec t'')$$
3. *Linearity*. Of any two distinct instants, one precedes the other:
$$\forall t \forall t'(t \prec t' \vee t = t' \vee t' \prec t)$$
4. *Past unboundedness*. Every instant has a predecessor:
$$\forall t \exists t'(t' \prec t)$$
5. *Future unboundedness*. Every instant has a successor:
$$\forall t \exists t'(t \prec t')$$
6. *Density*. Between any two distinct instants there is a third instant which precedes one and is preceded by the other:
$$\forall t \forall t'(t \prec t' \rightarrow \exists t''(t \prec t'' \wedge t'' \prec t'))$$

[17] These subsections are based on material included in [30]

[18] In the time protocol used for internet synchronisation, the zero point is midnight on 1st January 1900.

7. *Continuity*. If a set S of instants is non-empty, has non-empty complement, and contains the predecessors of all its members, then there is an instant all of whose predecessors, and none of whose successors, is in S:[19]

$$\forall S(\exists t S(t) \wedge \exists t \neg S(t) \wedge \forall t \forall t'(S(t) \wedge t' \prec t \rightarrow S(t')) \rightarrow \exists t \forall t'(t' \prec t \rightarrow S(t) \wedge t \prec t' \rightarrow \neg S(t)))$$

The continuity condition expresses the Dedekind property of the real numbers, distinguishing their order type from that of the rational numbers, which also possess the first six properties above. In logical terms, it is notable that the first six properties can be expressed in first-order logic, whereas the seventh property cannot.[20] Moreover, these six first-order properties are complete in the sense that *any* first-order-expressible property of the ordering on the real numbers necessarily follows, by pure logic, from the six given properties; as such, the six properties can be taken as providing an *axiomatisation* of the first-order logic of the qualitative ordering of real-valued time [7], and hence a suitable basis for qualitative reasoning.

For the purpose of qualitative temporal reasoning over particular domains, one may wish to omit or replace one or more axioms. For example, the density axiom might be replaced by

6′. *Discreteness*. If an instant has a predecessor (successor), then it has an immediate predecessor (successor):[21]

$$\forall t \forall t'(t \prec t' \rightarrow \exists t''(t'' \prec t' \wedge \forall t'''(t''' \prec t' \rightarrow t''' \preceq t'')) \wedge \exists t''(t \prec t'' \wedge \forall t'''(t \prec t''' \rightarrow t'' \preceq t''')))$$

The axioms 1–5 and 6′ characterise the first-order properties of the ordering of the integers, but again not exhaustively — to capture the integers uniquely one again needs to supplement these axioms with a property not expressible in first-order logic (e.g., that between any two instants there are at most finitely many others).

Discrete (integer-like) time can be a useful model for situations where the system under study evolves in a sequence of discrete steps, as for example the execution of a computer program, or the moves in a board game. If one is studying a phenomenon that only displays significant structure on a time-scale of more than a day, say, then it is reasonable to use a discrete sequence of days as an idealised representation of time for the purposes of modelling that phenomenon, and if one does this then the implicit logic is that of axioms 1–5, 6′ rather 1–6.

Another interesting possibility is to drop the linearity axiom (3). By doing so, one can build into the temporal model an asymmetry between the past and the future by which, at any point of time, there is a unique past but many possible futures. The possibility here may represent genuine indeterminacy, or simply uncertainty, so one can adopt this model without having to commit oneself to any particular philosophical stance on determinism and free will. This ‘branching time’ model can be obtained by replacing axiom 3 by

3′. *Past-linearity*: Of any two distinct instants which both precede some third instant, one precedes the other.

In Figure 9 is shown a simple branching-time model. Here any two instants which both precede the time labelled t must lie along the line of times stretching into the past from t; these form a linear series, and hence one of the instants must precede the other, as required by axiom 3′. But of the distinct times t and t^*, which do not both precede some common third time, neither precedes the other.

[19] If it is in S, this instant will be the *last* instant in S; if not, the first instant not in S.

[20] It requires *second-order logic* to express it, as shown by the bound predicate-variable S in our formula.

[21] An immediate predecessor of a given instant is a latest predecessor, i.e., a predecessor which is preceded bu all the other predecessors of that instant; and similarly for successor, with ‘precedes’ replaced by ‘follows’ (the inverse of ‘precedes’).

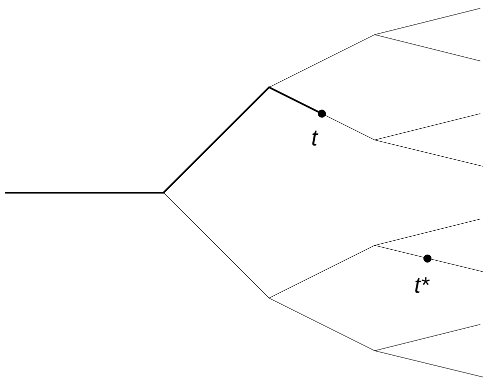

Figure 9: A simple branching-time model

13.2 The logic of intervals

Within AI knowledge representation, it was soon noticed that for many purposes we are more interested in intervals than instants. If we are talking about an event or a state of affairs in the world, and ask when the event happened or the state of affairs obtained, the answer is usually an interval rather than an instant: e.g., John travelled to London on August 26th; Mary was in the house from 2 p.m. to 3.45 p.m.

An early, and highly influential, interval-based theory of time was that of James Allen [2, 3]. Allen showed that the qualitative relationships amongst time intervals can all be expressed in terms of a single primitive relation m, read 'meets', where '$a \,\mathsf{m}\, b$' means that the interval a ends exactly when b begins. For brevity, we will write $i \,\mathsf{m}\, j \,\mathsf{m}\, k$ to mean $i \,\mathsf{m}\, j \wedge j \,\mathsf{m}\, k$. The systematic treatment of intervals, and the 13 distinct qualitative relations between them, has become known as the Interval Calculus. The relations are defined as follows:

Name	Symbol	Definition
is before	$<$	$a < b \equiv \exists j(a \,\mathsf{m}\, j \,\mathsf{m}\, b)$
meets	m	Primitive
overlaps	o	$a \,\mathsf{o}\, b \equiv \exists i \exists j \exists k \exists l \exists m(i \,\mathsf{m}\, j \,\mathsf{m}\, k \,\mathsf{m}\, l \,\mathsf{m}\, m \wedge i \,\mathsf{m}\, a \,\mathsf{m}\, l \wedge j \,\mathsf{m}\, b \,\mathsf{m}\, m)$
starts	s	$a \,\mathsf{s}\, b \equiv \exists i \exists j \exists k(i \,\mathsf{m}\, a \,\mathsf{m}\, j \,\mathsf{m}\, k \wedge i \,\mathsf{m}\, b \,\mathsf{m}\, k)$
finishes	f	$a \,\mathsf{f}\, b \equiv \exists i \exists j \exists k(i \,\mathsf{m}\, j \,\mathsf{m}\, a \,\mathsf{m}\, k \wedge i \,\mathsf{m}\, b \,\mathsf{m}\, k)$
is during	d	$a \,\mathsf{d}\, b \equiv \exists i \exists j \exists k \exists l(i \,\mathsf{m}\, j \,\mathsf{m}\, a \,\mathsf{m}\, k \,\mathsf{m}\, l \wedge i \,\mathsf{m}\, b \,\mathsf{m}\, l)$
equals	$=$	$a = b \equiv \exists i \exists j(i \,\mathsf{m}\, a \,\mathsf{m}\, j \wedge i \,\mathsf{m}\, b \,\mathsf{m}\, j)$
is after	$>$	$a > b \equiv b < a$
is met by	mi	$a \,\mathsf{mi}\, b \equiv b \,\mathsf{m}\, a$
is overlapped by	oi	$a \,\mathsf{oi}\, b \equiv b \,\mathsf{o}\, a$
is started by	si	$a \,\mathsf{si}\, b \equiv b \,\mathsf{s}\, a$
is finished by	fi	$a \,\mathsf{fi}\, b \equiv b \,\mathsf{f}\, a$
contains	di	$a \,\mathsf{di}\, b \equiv b \,\mathsf{d}\, a$

To see how the more complicated definitions work, look at Figure 10.

It will be seen that the relations come in pairs which are mutually inverse (e.g., if a meets b then b is met by a). The 'equals' relation is unpaired since it is its own inverse. This system of 13 relations is said to be JEPD, meaning 'Jointly Exhaustive and Pairwise Disjoint' — meaning that, for any two intervals, exactly one of the relations must hold.

Freksa [22] neatly characterised the 13 relations of the Interval Calculus in terms of relations between the endpoints of the intervals concerned. Figure 11, adapted from Freksa's paper, illustrates this; here α and A are the beginnings of the two intervals, and ω and Ω are the corresponding endings. Clearly

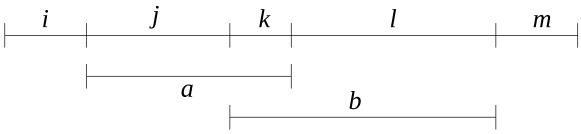

Figure 10: Illustration of $a \mathsf{o} b \equiv \exists i \exists j \exists k \exists l \exists m (i \mathsf{m} j \mathsf{m} k \mathsf{m} l \mathsf{m} m \wedge i \mathsf{m} a \mathsf{m} l \wedge j \mathsf{m} b \mathsf{m} m)$

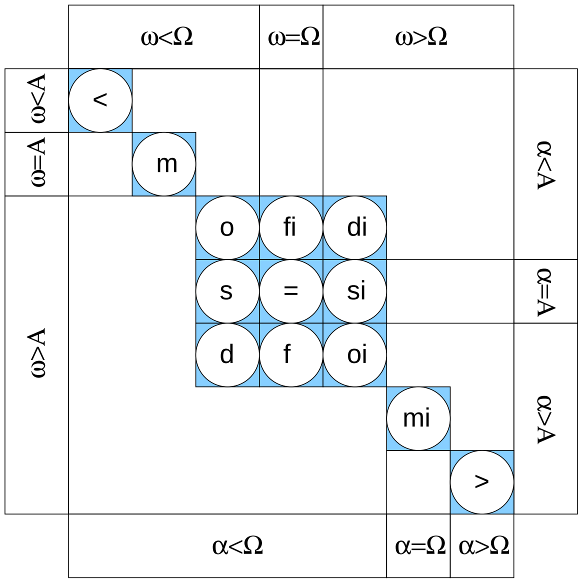

Figure 11: The Interval Calculus relations derived from the relative ordering of the endpoints (after [22]).

we must have $\alpha < \omega$ and $\mathrm{A} < \Omega$; the 13 interval relations correspond to the possible ways of consistently assigning a relative ordering to each of remaining pairs (α, A), (α, Ω), (A, ω), (ω, Ω), as shown in Freksa's diagram.

Freksa introduced the concept of *conceptual neighbourhood*, which he defined as follows:[22]

> **Definition 1** (Freksa 1992): Two relations between pairs of events are (*conceptual*) *neighbours*, if they can be directly transformed into one another by continuously deforming (i.e., shortening, lengthening, moving) the events (in a topological sense).

Thus, for example, the relation o is a conceptual neighbour of s since if a overlaps b, then by delaying the beginning of a until it coincides with the beginning of b, it then holds that a starts b. This can be read off from Figure 11 from the fact that the relations o and s occupy neighbouring cells in the diagram. Freksa distinguished three different neighbourhood relations corresponding to different allowed deformations, and conjectured that 'if a cognitive system is uncertain as to which relation between two events holds, uncertainty can be expected particularly between neighboring concepts'.

[22]Note that Freksa here refers to events rather than intervals; this fully accords with Allen's declaration that intervals are of interest only insofar as they are the times at which events occur.

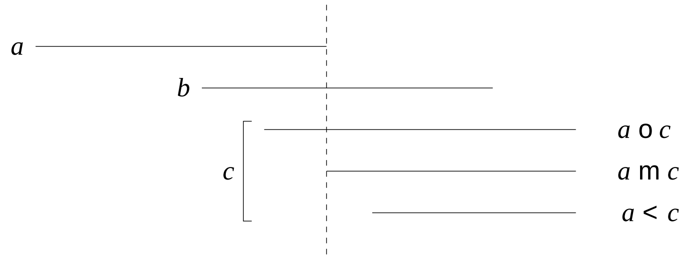

Figure 12: Composition of qualitative relations on temporal intervals

13.3 Compositional reasoning with intervals

The Interval Calculus is governed by a comprehensive set of *composition rules*, of which a typical example is:

$$\forall a \forall b \forall c (a \,\mathsf{o}\, b \wedge b \,\mathsf{o}\, c \rightarrow a < c \vee a \,\mathsf{m}\, c \vee a \,\mathsf{o}\, c),$$

i.e., if a overlaps b and b overlaps c then a either is before, meets, or overlaps c. This is illustrated in Figure 12, where each of the three possibilities for the relationship between a and c is shown. It should be noted that before, meets, and overlaps form a connected set ($\{<, \mathsf{m}, \mathsf{o}\}$) under conceptual neighbourhood — and indeed all the entries in the composition table have this property, corroborating Freksa's conjecture.

The complete set of composition rules is presented in a table (variously known as a 'transitivity table' or a 'composition table'). The idea is that a reasoner equipped with this table and an appropriate algorithm for deductive reasoning (e.g., using constraint propagation) is able to make inferences such as the following: Given that the time of the earthquake overlaps the time of the landslide, and the time of the landslide overlaps the collapse of the dam, it follows that the time of the earthquake overlaps, meets, or precedes the collapse of the dam.

Unfortunately, this kind of reasoning is rather limited, because many of the entries of the composition table contain too many alternative possibilites. The example in Figure 12 has three; but if all we know is that a is before b and b is after c, then the relationship between a and c could be any of the 13 interval calculus relations. In such a case we require information which goes beyond the purely qualitative, e.g., about the duration of the intervals, or the extent of their overlap.

More generally, for the purposes of reasoning, we need to consider not just the JEPD set of 'base' relations listed above, but also disjunctions of these. For example, if all we know about two intervals is that a starts before b, then the relation between them could be any of $<$, m, o, fi, and di (as can be read off from Figure 11, expressing the constraint as $\alpha < \mathrm{A}$), and hence may be represented as a *disjunctive* relation which we may denote $\{<, \mathsf{m}, \mathsf{o}, \mathsf{fi}, \mathsf{di}\}$. The total number of such relations is $2^{13} = 8192$, and these constitute the full *interval algebra*, often denoted $\mathcal{A}$. The composition table for $\mathcal{A}$ thus has $8192^2 = 67,108,864$ entries — these do not need to be computed individually, however, as any individual entry is readily derived from the composition table for the base relations.

Much work has been done on the mathematical and computational properties of the Interval Calculus and related systems. Of particular interest, as regards the applicability of any of these formalisms, is whether or not it is *decidable*, and if so, what is the *computational complexity* of its decision problem. A formalism is decidable if there is an algorithm, specifiable in advance, which will determine, for any collection of statements in the theory, whether or not they have a model, that is an actual example of a situation which is correctly described by those statements.

The most basic problem for reasoning over $\mathcal{A}$ is the *constraint satisfiability problem*. An instance of this problem consists of a set of constraints of the form 'i stands in relation R to j', where 'i' and 'j' are variables standing for intervals and 'R' is any one of the 8192 relations in $\mathcal{A}$. Given such a set, the problem is to determine whether it is possible to associate actual intervals (represented using real-number pairs, e.g., (1.5,2.7)) with the variables appearing in the set in such a way as to satisfy all the constraints. As a simple illustration, consider the following set of constraints:

$$a\{\,\mathsf{m}\,,\,\mathsf{o}\,\}b, \qquad b\{\,\mathsf{f}\,,=,\,\mathsf{fi}\,\}c, \qquad c\,\mathsf{mi}\,d, \qquad d < a$$

It is easy enough, in this case, to find an allocation of intervals to the variables so that these constraints are satisfied, e.g., $a = (3,4), b = (4,5), c = (2,5), d = (1,2)$. Equally, if the third constraint is replaced by $c\,\mathrm{m}\,d$, it is easy to see that there is *no* such allocation. A general solution of the constraint satisfiability problem for $\mathcal{A}$ would be an algorithm which will correctly determine for an *arbitrary* set of constraints whether or not it is satisfiable.

An approach which provides partial solutions to this makes use of the idea of *path-consistency*. Suppose we have a set of interval variables, say $i_1, i_2, \ldots, i_n$, together with constraints (i.e., $\mathcal{A}$ relations) for at least some pairs of intervals — a pair of intervals for which no explicit constraint is given will be allocated the 'non-constraint' consisting of all 13 Allen relations. Such a system of relations is said to be path-consistent so long as every solution for any two variables can be extended to a solution for those two variables together any other variable. For $\mathcal{A}$ (though not in general), this property is equivalent to *algebraic closure* [50]: for any three intervals i, j, k, if the $\mathcal{A}$ relation assigned to i and j is R_1 and that assigned to j and k is R_2, then the relation assigned to i and k is contained in the composition of R_1 and R_2 as given in the composition table. Given such a set of constraints, one seeks to enforce path-consistency by systematically removing from the stated constraints any that conflict with the algebraic closure criterion, and iterating this process until either the relation assigned to some pair of intervals becomes empty, or no further changes to the relations occur. In the former case, we can infer that the original set of constraints cannot be satisfied, and we are finished. In the latter case, however, although the constraints have now been made path-consistent, we are still not guaranteed the existence of a model, since unfortunately path-consistency turns out not to be equivalent to satisfiability. The upshot of this is that in order to solve the constraint satisfiability problem in cases where path-consistency is established, further work needs to be done.

How much more work is needed? The path-consistency algorithm of Van Beek [83] runs in cubic time — that is, the number of computational steps required is approximately proportional to the cube of the number of variables in the instance being tested — and is therefore tractable. Vilain and Kautz [85] showed that the constraint satisfiability problem for $\mathcal{A}$ is NP-complete. No tractable algorithms for NP-complete problems are currently known, and it is widely believed that none can exist; if this proves to be correct,[23] then the amount of computational effort needed to resolve the hardest cases of the problem will always increase as an *exponential* function of the number of intervals. For this reason, it is important to identify restrictions to the general decision problem for $\mathcal{A}$ which make it tractable. A number of researchers have investigated restrictions taking the form of *subalgebras* of $\mathcal{A}$, that is, subsets of the full set of interval algebra relations that are closed under the operations of intersection, converse, and composition. A number of maximal tractable subalgebras were identified [59, 19], culminating in an exhaustive enumeration of the tractable subalgebras of $\mathcal{A}$ by Krokhin *et al.* [44].

It is generally agreed nowadays that a model of time based on intervals alone is not wholly satisfactory: it seems that we will always want to talk about both intervals and instants. See [23] for an early critique of the Interval Calculus in this respect.

[23] That is, P$\neq$NP. See [32] for details on NP-completeness and the 'P=NP' problem

SESSION 4: Representation and Modelling

14 Data modelling vs Process modelling

The need to incorporate the temporal dimension has been recognised and repeatedly proclaimed almost throughout the history of GIScience. It is usual to draw a contrast between "traditional GIS functions" comprising the input, storage, retrieval, management, manipulation, and presentation of data, and a range of more advanced capabilities such as analysis, simulation and prediction which engage with the data through some form of theoretical understanding of the real-world systems and processes that the data represent. This contrast has been described in various different ways, all tending to the same (often rather despairing) conclusion that it is high time that the data modelling functions of GIS were integrated with the process modelling requirements of at least a substantial proportion of the users of GIS.

Hazelton et al [37], for example, writing in 1992, complain that the then-current GIS are mainly "map-representation systems" ("data retrieval and sifting engines") whereas what is needed, if GIS are to play a part in e.g., spatial decision-support systems, is "reality-representation systems". They contrast "empirical models" (where the connotation of 'empirical' is that the system is just a repository of observations, with no theoretical considerations or *a priori* reasoning) with "dynamic models", including process models, simulations, predictive models, etc, which (almost) invariably require time as a dimension, and operate with rules based on a theoretical understanding of how the underlying reality works.[24] They claim that the distinction between empirical and dynamic models is clear for 2D GIS, but that if we move to 4D GIS "the dividing line starts to blur"; however, it seems perfectly possible to have "data retrieval and sifting engines" that operate on fully 4D data, and these are then purely empirical, not dynamic models.

Peuquet [64] notes that "current GIS are limited to queries that perform straightforward retrieval of observational data based on relational constraints", and suggests that future spatio-temporal GIS will include a wider range of functions, notably *exploration*, *explanation*, *prediction*, and *planning*. To achieve this, it is not enough merely to incorporate the temporal dimension, which would only result in a "world history model" with similar functions to a traditional spatial GIS; in addition one must add "higher-level knowledge" (including causal relationships) to obtain a "process model".

Raper and Livingstone [69] draw two contrasts which are clearly related to this idea, though neither seems to coincide exactly with the contrast drawn in [37]. First they contrast "geometrically-indexed" and "object-oriented" methods, where the former "makes the coordinate system of the layer into the primary index of the spatial representation and dictates much of the representational structure of the environmental problem of interest", whereas in the latter the spatial representation is structured by the "characteristics and behaviour" of the real-world entities being modelled. It is implicit that traditional GIS uses the former type of method, whereas the latter would suit most scientific purposes better. Second, they draw a contrast between "data models" and "process models", characterising the former as expressing theories "predicting the structure of real world domains in terms of entities and their attributes organised in inter-related sets" and the latter as expressing theories "predicting the nature of the exchange of energy or mass within systems, over time". They note that whereas environmental models are specified as process models, GIS are specified as data models, which leads to a "basic problem" if we want to couple an environmental model with a GIS.

Raper and Livingstone are not explicit about the relationship between the two contrasts drawn here, although it seems to me that they are clearly distinct and to some extent independent. One might characterise raster and vector GIS as geometrically-indexed data models and object-oriented data models respectively; and there can be both geometrically-indexed process models such as cellular automata (see §16.2) and object-oriented process models such as agent-based systems (see §16.3).

[24] They attribute this distinction to Burrough et al [12], who, however, use "process" for "dynamic".

There is a general presumption that process models are intrinsically temporal, but this does not imply that data models cannot be temporal too. Brown et al [11] draw attention to Raper and Livingstone's [69] distinction between spatial data models and spatial process models, describing them as representing *form* and *process* respectively. They point out that even "extentions to existing spatial data models to provide richer representations of time" (such as [90, 62, 93, 46]) do not include "explicit representations of the processes (i.e., sequences of operations) by which spatial features change and move, focusing instead on representing the changing structure of the phenomena" — which limits these systems to being data models rather than true process models.

Takeyama and Couclelis [77] similarly refer to, and seek to address, the problems of integrating mathematical modelling and GIS, noting that "the map layer is a natural metaphor for representing static spatial patterns and structures, but it cannot express rules, regularities and processes which are more appropriately represented by equations".

In 1998, we find Bregt and Bulens [10] complaining that "current GIS applications . . . tend to concentrate on the *storage*, *management* and *presentation* of spatial data, not the *analysis* and *modelling* applications" [my italics], whereas on the other hand, "[m]any current [process] models lack proper tools for data input and management, and have poor presentation facilities", thus clearly indicating the need to integrate the two types of system. It is striking that almost the same complaint is still being made in 2009, witness Torrens [81], who points out that data models "determine how GIS *stores*, *organizes*, and *displays* various types of information for different purposes" [my italics], but that "[i]n GIS there is usually no dedicated model of the *processes* that govern dynamics, adaptation, and evolution of a system". He claims, though, that "the next generation of geographic information systems will be driven by process models".

Bivand and Lucas [8] list the primary functions of existing GIS as data input processing, data storage and retrieval and database management, data manipulation and analysis, and display and product generation. The problem is how to link such systems to process models, where the latter may cover such diverse domains as hydrology, land surface and subsurface models, ecology, landscape ecology, and atmosphere and ocean models, all of which, critically, involve evolution through time.

It is striking that in the sources discussed above the term 'analysis' is sometimes recruited on the side of data modelling and sometimes on the side of process modelling. The truth is that both kinds of modelling require various forms of analysis. Bregt and Bulens [10] refer to the "analytical capabilities" of GIS as including overlay, interpolation, and transformation, and it is certainly possible to think of these as ways of analysing data as such without reference to any underlying theories of what the data represent.

In [26], I introduced a distinction between *information systems* and *modelling systems*, where the former functions as a repository of information endowed with sufficient computational capability to allow the information to be retrieved and presented in various different ways, whereas the latter should incorporate sufficient theoretical grounding to enable us to predict the future, explain the past, or explore 'what if' scenarios. This is, again, essentially the distinction between data modelling and process modelling.

What are the implications of all this for the way we handle the temporal component of geographical information, in particular processes and events? As already noted, it cannot be said that process modelling has an exclusive monopoly on the temporal dimension. There has been a good deal of work on incorporating time into data models, e.g., the considerable amount of work on temporal databases [45, 76] (which cover a much wider range of concerns than just the geographical), and such work continues to be done. However, the functions of a temporal data model are essentially the same as those of a non-temporal one: input, storage, manipulation, analysis, combination, and presentation of data, which in this case has a temporal dimension. By contrast, process models, as the name suggests, are, almost invariably, intrinsically temporal, and one cannot easily envisage the development of a process model without some kind of theory of time, process, and event, however rudimentary, to support it.

In the following sections I shall describe some examples of spatio-temporal data modelling (§15) and spatio-temporal process modelling (§16). The coverage is inevitably rather incomplete but is designed to give a flavour of the main types of work that have been done in these areas. Further details may be obtained by following up the references given.

15 Spatio-temporal data modelling

15.1 ESTDM (Peuquet and Duan, 1995)

Peuquet and Duan [62] introduce an 'event-based spatiotemporal data model' (ESTDM) to provide a framework for analysing temporal geographical data. ESTDM uses a field-based spatial representation, but the temporal part is event-based. The model consists of a *base map* giving a snapshot of the field at some initial time, and, associated with this, a list of events.[25] Each event is located at a time, and contains information about all changes that have taken place since the previous event. The changes are organised into components, each of which records, for a particular field value, all the locations which acquire that value as a result of the change. The data model has been designed for storage and retrieval efficiency.

Peuquet and Duan note, amongst potential variations on this formalism, 'the representation of changes in polygon boundaries or other features instead of locations' — in other words, they suggest that ESTDM could be adapted to handling object-based rather than field-based geographical data; but they give no indication of how this might be accomplished.

15.2 GEM (Worboys and Hornsby, 2004)

Worboys and Hornsby [88] note that although "objects and events are fundamentally different", they exhibit many "structural similarities"; for this reason, in creating their geospatial event model (GEM) they aimed to "use as much as possible of the methods of object-orientation, designed for modelling static entities, [for] event modelling". Thus both objects and events are first-class entities in their system, allowing a rich set of relationships between events and events, objects and objects, and events and objects. Thus objects *participate* in events, which in turn *involve* the objects that participate in them. Both objects and events are *situated* in 'settings', which may be spatial, temporal, or spatio-temporal. Both objects and events may be classified hierarchically, may have associated attributes, and may be aggregated into composites or disaggregated into parts. Thus the philosophy of object-orientation may be applied equally to both types of entity.

This kind of approach is well-adapted to providing an appropriate conceptual structure to a spatio-temporal database. Relationships between classes of object or event can capture general truths about the domain, which enable inferences to be made concerning specific instances of those classes. It should be noted that, in common with other spatio-temporal data models, although the model portrays a dynamic reality, the model itself is in a certain sense static; this is not to say that the model cannot undergo change, but that changes to the model reflect changes in our state of knowledge about the information being recorded, not changes in the reality portrayed by the model.

15.3 Map Algebra and its extensions

Map Algebra, introduced by Tomlin in 1983, is a codification of the various arithmetical, boolean and other operations used in GIS for performing overlays on raster data. Raster data is typically organised into *layers*, each of which consists of a two-dimensional grid of raster cells, to each of which is assigned

[25] This essesntially conforms to the 'base map with overlays' idea which is given by Langran and Chrisman [47] as one of three possible "images of cartographic time".

a thematic value, with different layers corresponding to different themes. In conventional map algebra [80] there are three types of operations:

- *Local* operations, which take as input one or more data layers and generate as output a new layer in which the value assigned to each cell is a function of the values assigned to the corresponding cells in the input layers.
- *Focal* operations, which take as input one or more data layers and generate as output a new layer in which the value assigned to each cell is a function of the values assigned to the cells in some specified neighbourhood of the corresponding cells in the input layers.
- *Zonal* operations, which take as input two layers, a value layer and a zone layer (which assigns each cell to one of a discrete set of *zones*), and returns, for each zone, a value determined by the values assigned to the cells in that zone by the value layer.

By combining operators one can build up more complex operations by means of which sophisticated analyses may be performed on the raw data.

Mennis *et al.* [53] pointed out that it is straightforward to extend the conventional map algebra to three dimensions, operating on data *cubes* rather than layers. It is natural to interpret the third dimension as spatial, representing altitudinal variation, but alternatively it can be interpreted as representing time. As Mennis *et al.* point out, this must be done carefully since 'there are important differences between a third spatial dimension (Z) and the temporal dimension (T)' — notably the directed nature of time ('inevitably, time marches on'), and the different units for time and space measurement. However, this does not prevent the specification of a powerful array of map algebra operations for manipulating data layers of different dimensionality, and they illustrate some of these operations in a case study of climate-related vegetational change in the south-eastern USA. In a subsequent paper [54], their prototype implementation is substantially extended, to produce *multidimensional map algebra* (MMA), which they describe as 'a general data processing language for spatio-temporal data'.

Frank [21] likewise notes that there is a close analogy between rasters and time-series: a time series is in effect a "temporal raster". In the terminology of [25], it is a *temporal field*, exactly analogous to a spatial field. Like Mennis *et al.*, Frank suggested extending the map algebra with analogous operations on time-series, thereby providing a means for GIS to handle temporal data in ways that are exactly conformable with well-established methods for handling spatial data — providing, in Frank's words, 'consistent operations for maps, time series, stacks of maps, etc.'.

15.4 Analysis of Movement Patterns

There has been a good deal of interest recently in methods for handling data concerning moving objects. This reflects the widespread availability of location-aware mobile data-collection and communication devices which mean that such data is becoming ever more easily obtainable. The problem is what to do with it. Typically we are concerned with large collections of data of the basic form "Object o is in position p at time t"—or, if object identities are not recorded, simply "There is an object in position p at time t". A collection of (o, p, t) data for a fixed individual o represents a sampling from the geospatial lifeline of o; for a fixed time t, we have a snapshot of the collective motion at that time; and for a fixed position p we have a history of the what happens at p. One wants to be able to extract from this data information about trends and patterns, the formation of coherent groupings of individuals (collectives) and the interactions within and between such groupings. All of this is in the first instance a data modelling problem; subsequently, of course, process modelling issues can be brought to bear (e.g., in order to simulate such collective phenomena using agent-based models—see §16.3).

The first step in the analysis of movement patterns is to classify them in some way, distinguishing different characteristic types. An interesting and detailed early attempt is proposed by Abler *et al.* [1, Chapter 8], who classify different forms of "movement geometry" in terms of the dimensionality (point, line, area, or volume) of both the source and destination. This leads to a table of sixteen different kinds. In a similar spirit, the "classification of geographic movement" given by Yattaw [92] also takes the form of a cross-classification, in this case of four "spatial characteristics" (i.e., the dimensionalities of point, line, area, and volume again) against three "temporal characteristics", namely continuous, cyclical, and intermittent, yielding twelve different classes.

A classification based on rather different principles is given by Thériault *et al.* [78], whose taxonomy for the evolution of "sets of geographical entities" (SGEs) has a top-level division into (I) changes related to the appearance or disappearance of individuals, (II) changes related to the movement of individuals, and (III) changes related to the movement of the entire SGE.

More recently, Dodge *et al.* [18] have attempted a comprehensive classification of movement patterns, taking account of such factors as the movement variables (e.g., position, direction, speed, acceleration), the number of objects involved, the path type, influencing factors (e.g., from the environment or other agents), and scale and granularity. They divide patterns into "generic" patterns such as concentration and dispersion, which relate to the abstract form of the movement without reference to the particular nature of the objects that are moving, and "behavioural" patterns such as courtship, foraging, and migration which are characteristic of particular types of objects. By its nature, this latter type is presented as an unstructured list, consisting as it does of a wide variety of different behaviours reflecting the highly diverse nature of the various collections of moving objects we find in the world; by contrast, the former type (generic patterns) is rather more systematically organised, though again there seems to be a tendency for the classification to take the form of bundles of unstructured lists with less formal organisation than in the cross-classifications discussed above.

The next step, after classifying patterns, is to devise methods for detecting and recognising them, a formidable problem given that such patterns may be hidden amongst vast quantities of unstructured positional data. This is a problem of *data mining*.

Laube *et al.* [48] introduce an analytical tool which they call REMO (for "RElative MOtion"). The idea is to extract from the collection of (o, p, t) data-points a matrix recording some motion attribute (e.g., position, heading, speed) of each individual at a set of regularly spaced times. Depending on the granularities of the raw data and the time-dimension of the REMO matrix, this may require interpolation or aggregation to be applied to the data. For the REMO matrix thereby obtained, methods are provided for detecting various significant patterns, notably *constancy*, in which one individual maintains the same value of the motion attribute over a specified number of time steps, *concurrence*, in which a set of individuals have the same value of the attribute at a given time step, and *trend-setting*, where one individual (the trend-setter) exhibits constancy over a period, at the end of which a set of individuals exhibit concurrence on the same value. These ideas were illustrated using data collected from a team of footballers during a match, and the motions in "ideological space" of political preferences in different districts of Switzerland over a four-year period.

Andrienko and Andrienko [4] are also interested in extracting patterns from large data sets recording the movements of collections of individuals. They propose that the primary interest is in what they call the "dynamic collective behaviour" (DCB), but they note that this can be described from two orthogonal points of view, which they call "aspectual behaviours". From one point of view, the DCB can be thought of the aggregation of the "individual movement behaviours" (IMBs) of the entities constituting the collective under study — in other words, something resembling a bundle of geospatial lifelines — while from the orthogonal point of view the DCB can be thought of as the evolution of the "momentary collective behaviour" (MCB) of the collective, where each MCB can be regarded as a snapshot showing the distribution of movement characteristics (position, heading, speed, etc) of the individuals at a given time. The two aspectual behaviours will exhibit different pattern types, which means that for a full understanding of the DCB, one must use both. (Compare the discussion of snapshots and histories in [25].)

16 Spatio-temporal process models

16.1 Eulerian vs Lagrangian process models

Brown *et al.* [11] draw attention to an important distinction between two approaches to spatial process modelling which, borrowing terminology from fluid dynamics, they call the *Eulerian* and *Lagrangian* approaches. The Eulerian approach is concerned with tracking the variation in properties, and the processes that give rise to this, at fixed locations; as such it provides a description of *change*. The Lagrangian approach, by contrast, is concerned with tracking the change of location of elements under study as they move through space, and as such, it provides a description of *movement*.

This distinction can be illustrated using the example of road traffic. On the one hand—and this is the Eulerian view—we might concentrate on the road network, characterising each location within the network in terms of the density and speed of traffic passing through it; the evolution of the system can then be determined using rules which, for example, might derive the attributes of a location at one time-step from the attributes of it and its immediately neighbouring locations at the preceding time-step. We could describe this as the view of the traffic obtained by a series of stationary observers distributed across key points around the network (what I called the bystander perspective in §6).

On the other hand—taking now the Lagrangian view—we could concentrate on the traffic itself, characterising each vehicle by its location and speed. The evolution of the system is then determined using rules governing how the location and speed of a vehicle at one time-step is determined by its location and speed at the previous time step and how it interacts with other vehicles (as well, presumably, as some statistical model of the travel plans of individual drivers). This corresponds to the participant perspective of §6.

As pointed out in [11], the Eulerian and Lagrangian models provide *dual* perspectives on a single underlying reality: in principle, from a complete Eulerian description one could derive a complete Lagrangian description, and vice versa. In reality, of course, our descriptions are never complete, and for different applications, one or the other approach may be more satisfactory, e.g., because it more naturally captures the way in which we are accustomed to conceptualise the reality being modelled.

In the next two sections we describe the most popular representatives of the Eulerian and Lagrangian approaches in current GIScience: cellular automata for the former, and agent-based models for the latter.

16.2 Cellular automata

A cellular automaton (CA) is a mathematically defined structure which evolves in accordance with purely local interactions, i.e., there is no action at a distance. The ingredients of a standard CA are

- An underlying space consisting of discrete *cells* with an *adjacency* relation. Most usually, this is a rectangular grid or raster.
- A set of discrete *states* which each cell may assume — in the simplest case, there are just two states, which may be characterised as "on/off", "alive/dead", "active/inactive", etc.
- A definition of the *neighbourhood* of each cell, i.e., which other cells can have an immediate effect on it. Where the underlying space is a rectangular grid the two commonest neighbourhoods used are the *Von Neumann neighbourhood*, which consists of the cell itself and the four orthogonally adjacent cells, and the *Moore neighbourhood*, which in addition includes the four diagonally adjacent cells (see Figure 13).
- A *transition rule*, which determines the state of a cell at time $t + 1$ as a function of the states of the cells in its neighbourhood at time t.

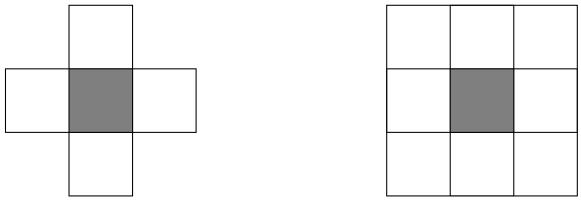

Figure 13: Commonly-used neighbourhoods for Cellular Automata: Von Neumann neighbourhood (left) and Moore neighbourhood (right)

A well-known example of CA is Conway's *Game of Life*, which uses an unlimited rectangular grid, two states ('dead' and 'alive'), Moore neighbourhoods, and the following transition rule:

> A live cell remains live if and only if it has either two or three live neighbours;
> A dead cell comes alive if and only if it has exactly three live neighbours.

This rule is remarkable for giving rise to extremely rich behaviour, and the resulting CA has been shown to be equivalent to a universal Turing machine — i.e., by suitable choice of initial states, it can in principle simulate any computation whatever.

The application of CA to geography can be traced back to Tobler [79], who was in turn inspired by the diffusion models of Hägerstrand [35]. Cellular automata have been used for simulating various geographical processes, notably urban development [87, 5, 49] and wildfires [13, 60]. However, in order for such simulations to have much chance of realism, it is generally necessary to modify the basic definition of CA in some way. Proposed modifications include

- Introducing "top-down" global constraints to reflect, e.g., the power of central government to control overall land development patterns [49]
- Allowing the underlying space to be non-uniform, e.g., with different cells having differently configured neighbourhoods or different transition rules [14, 61].
- Allowing the underlying space itself to change, e.g., by using Voronoi cells determined by objects which themselves may move or change shape as a result of the transition rules [74].
- Making the transition rules stochastic, so that the state of cells in a cell's neighbourhood only determine its subsequent state with a certain probability [6].
- Allowing some form of action at a distance (e.g., using an inverse-square law to model the influence of a cell on cells at different distances)—this breaks with the fundamental principle of cellular automata, i.e., local interactions only, and hence the resulting models have been called *cell space models* [6]. Action at a distance may be limited to small distances, thereby preserving the spirit of locality to some degree: for example, in the model for urban growth proposed by [87], cells can only be influenced directly by other cells up to 6 cell-widths distant, but within this neighbourhood the influence is dependent on distance from the centre.

By exploiting the correspondence between the functions on neighbourhoods used for focal operations in Map Algebra (see §15.3) and the neighbourhood-based transition functions of a cellular automaton. Takeyama and Couclelis [77] integrated cellular automata with a generalised form of Map Algebra called Geo-Algebra, with the specific intension of facilitating the incorporation of dynamic spatial modelling within GIS. They described this work as "an aspect of a wider effort to integrate mathematical modelling and GIS"—compare our discussion in §14.

16.3 Agent-based models

As noted by [9], agent-based modelling (ABM)[26] is “a mindset more than a technology”, the key element being the description of a system “from the perspective of its constituent units”. Like cellular automata, therefore, ABM represents a “bottom-up” approach to modelling, and indeed, cellular automata could themselves be regarded as agent-based, the agents in this case being the individual cells. As normally understood, however, the agents in ABM are generally mobile, not static like the cells in a CA, and for this reason it seems more natural to contrast ABM and CA rather than unite them: as pointed out by Torrens and Benenson [82], it was precisely the static nature of the cells in CA that turned the attention of geographers to ABM.

Agent-based models are typically used for simulating large-scale collective phenomena which emerge from the interaction of many mobile autonomous individuals each of which responds to its environment (including other individuals) in accordance with a set of rules. The rules governing agent behaviour may be simple or complex, although there is particular interest in those cases (analogous to Conway’s Game of Life in the CA field) where complex emergent behaviour arises from relatively simple rules. A classic example of this kind is Craig Reynolds’ “boids”, in which animal flocking is simulated using agents whose individual behaviour is governed by just three principles: *separation* (to avoid collisions with flock-mates), *alignment* (to move in a similar way to flock-mates), and *cohesion* (to move closer to group of flock-mates).[27]

The work of Schelling [72] on dynamic segregation models is an early example of what is in effect agent-based modelling, showing how segregation of different social groups can occur purely as a result of each individual exerting a preference to be located close to members of the same group, without any overall “top-down” control. This really belongs to the prehistory of ABM; the modelling was all performed by hand without any recourse to automated computation. Subsequently, ABM has been applied to many areas of geographic relevance, including urban development [65, 16], traffic simulation [20], wayfinding [70], and the rôle of humans in landscape evolution [86].

Torrens and Benenson [82] note that cellular automata and multi-agent systems are ‘popularly confused in the geographic literature’, and propose a new model which integrates both concepts into a truly geography-based framework. They borrow the notion of an automaton from theoretical Computer Science as an entity characterised by a set of states S and a set of transition rules expressed by a function $T : S \times I \rightarrow S$, where I is the set of possible inputs: hence, if an automaton is in state S_t at time-step t, and receives input I_t, its state at time-step $t+1$ will be $S_{t+1} = T(S_t, I_t)$. This formulation is sufficiently general to cover both the individual cells of a cellular automaton and the individual agents in a multi-agent system, and hence can be used to model both *agents* and *features* and combine them within a single complex system.

A Geographic Automata System (GAS) [82] consists of a set of automata classified into some fixed collection of types; each automaton is characterised, at any time-step, by its *state*, its *location*, and its *neighbours*. For each type of automaton there are three sets of rules, determining an automaton’s new state, location, and neighbours as a function of its current state, location, and neighbours. Mobile and fixed automata can thus both be accommodated, the latter being characterised by rules determining that their location never changes — on the other hand their state can change in response to the movement patterns of the mobile automata, and this means that the CA-like aspects and ABM-like aspects can, in principle, be fully integrated into a single model. Torrens and Benenson systematically examined the models found in the literature on urban modelling and concluded that all of them can be easily reformulated as a GAS (see their table 2).

[26] Also known as multi-agent systems (MAS).

[27] See http://www.red3d.com/cwr/boids/ for demonstration applet, and [71] for a technical paper.

16.4 Process Calculi

Process calculi (including process algebras) are formalisms originally designed for modelling concurrent computational processes, but Worboys [91] has advocated their use for modelling real-world geospatial processes too. In the context of process calculi, a process is regarded as *the behaviour of a system*, typically viewed in an abstract, idealised way as constituting a sequence of discrete events or actions. A given system may exhibit different such sequences depending on the nature of its interactions with other systems; the definition of the system's behaviour (that is, its process definition) will be such as to cover all the different particular event-sequences that it may exhibit.[28]

A process algebra allows one to define complex processes by specifying how they are built up out of simpler ones. The basic operations are

- *Sequential composition*: The process $a; b$ consists of the process a followed by the process b.
- *Parallel composition*: The process $a \parallel b$ consists of processes a and b operating concurrently.
- *Alternative composition*: The process $a + b$ consists of a non-deterministic choice between processes a and b.

Thus, for example, if you decide to toss a coin to determine which route to take at a junction, your behaviour may be described as the process $heads; right + tails; left$, i.e., non-deterministically either you obtain heads and then turn right or you obtain tails and then turn left. In a similar way, Worboys specifies the (considerably more complex) protocol for traffic passing through a four-way stop.

The π-calculus [55] is an extended form of process algebra which includes amongst the basic operations the output and input of messages along specified channels, as follows:

- *Input*: The process $c(x)$ consists of receiving an input on channel c and assigning it to the variable x.
- *Output*: The process $\bar{c}(v)$ consists of sending the value v along channel c.

Then the process

$$(heads; \bar{c}(heads) \parallel c(x); right) + (tails; \bar{c}(tails) \parallel c(x); left)$$

models the situation where one agent tosses the coin and communicates the result to the second agent, who then turns left or right depending on the information received. Worboys uses π-calculus to specify the protocol for a vehicle passing through a toll booth, and, more ambitiously, to show how the motion of a vehicle through a sequence of positions over a sequence of times can be specified as a process.

Because of their genesis in theoretical computer science, process calculi are best adapted for modelling well-defined processes operating according to rules that are known in advance. If a system of agents can be specified in this way, then the process calculus can be used to simulate or reason about its behaviour, thereby serving the goals of prediction or explanation. In the context of GIScience, the most likely settings for applying such calculi are those in which, as with the examples given by Worboys, the behaviour of the relevant entities is governed by a set of well-defined protocols which admit of no exceptions. Such settings are most likely to be encountered in the context of human geography rather than physical geography. It is an open question how far this kind of approach can take us towards modelling compex real-world scenarios in a useful way.

[28] For a readable general historical introduction to process algebras, see J. C. M. Baeten, 'A Brief History of Process Algebra' at http://www.win.tue.nl/fm/0402history.pdf (Accessed 29/07/09).

References

[1] R. Abler, J. S. Adams, and P. Gould. *Spatial Organization: The Geographer's View of the World*. Prentice-Hall International, 1971.

[2] James Allen. Towards a general theory of action and time. *Artificial Intelligence*, 23:123–54, 1984.

[3] James F. Allen and Patrick J. Hayes. A common-sense theory of time. In Aravind Joshi, editor, *Proceedings of the Ninth International Joint Conference on Artificial Intelligence (IJCAI'85)*, pages 528–31. International Joint Conferences on Artificial Intelligence, Inc., 1985.

[4] N. Andrienko and G. Andrienko. Designing visual analytics methods for massive collections of movement data. *Cartographica*, 42(2):117–138, 2008.

[5] M. Batty and Y. Xie. Possible urban automata. *Environment and Planning B: Planning and Design*, 24:175–192, 1997.

[6] Michael Batty. Geocomputation using cellular automata. In S. Openshaw and R. J. Abrahart, editors, *GeoComputation*, pages 95–126. Taylor & Francis, London, 2000.

[7] J. F. A. K. Van Benthem. *The Logic of Time*. Kluwer, Dordrecht, 1983. Second Edition, 1991.

[8] Roger Bivand and Anne Lucas. Integrating models and geographical information systems. In S. Openshaw and R. J. Abrahart, editors, *GeoComputation*, pages 331–363. Taylor & Francis, London, 2000.

[9] Eric Bonabeau. Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences of the USA*, 99:7280–7287, 2002.

[10] A. K. Bregt and J. Bulens. Integrating GIS and process models for land resource planning. In H. Heineke et al, editor, *European Soil Bureau Research Report No. 4*, pages 293–304. Laboratory of GeoInformation Science and Remote Sensing, Wageningen University, The Netherlands, 1998.

[11] Daniel G. Brown, Rick Riolo, Derek T. Robinson, Michael North, and William Rand. Spatial process and data models: Towards integration of agent-based models and GIS. *Journal of Geographical Systems*, 7:25–47, 2005.

[12] P. A. Burrough, W. van Deursen, and G. B. M. Heuvelink. Process models and GIS: A marriage of convenience or a blossoming partnership? In *Proceedings of GIS/LIS'88*, pages 598–607, Falls Church, Va., 1988. ASPRS, ACSM.

[13] K. C. Clarke, J. A. Brass, and P. J. Riggan. A cellular automaton model of wildfire propagation and extinction. *Photogrammetric Engineering and Remote Sensing*, 60(11):1355–1367, 1994.

[14] Helen Couclelis. Cellular worlds: a framework for modeling micro-macro dynamics. *Environment and Planning A*, 17:585–596, 1986.

[15] Helen Couclelis. People manipulate objects (but cultivate fields): Beyond the raster-vector debate in GIS. In A. U. Frank, I. Campari, and U. Formentini, editors, *Theories and Methods of Spatio-Temporal Reasoning in Geographic Space*, volume 639 of *Lecture Notes in Computer Science*, pages 65–77, Berlin, 1992. Springer-Verlag.

[16] A. T. Crooks. Exploring cities using agent based models and GIS. In D. Sallach, C. M. Macal, and M. J. North, editors, *Proceedings of the Agent 2006 Conference on Social Agents: Results and Prospects*, pages 125–132. Argonne National Laboratory and University of Chicago, 2006.

[17] Donald Davidson. The logical form of action sentences. In Nicholas Rescher, editor, *The Logic of Decision and Action*. University of Pittsburgh Press, Pittsburgh, Pa., 1967. Reprinted in D. Davidson, *Essays on Actions and Events* (Oxford: Clarendon Press, 1980).

[18] S. Dodge, R. Weibel, and A.-K. Lautenschütz. Taking a systematic look at movement: Developing a taxonomy of movement patterns. In *The AGILE Workshop on GeoVisualization of Dynamics, Movement and Change, Girona, Spain, 5th May 2008*, 2008.

[19] Thomas Drakengren and Peter Jonsson. A complete clasification of tractability in Allen's algebra relative to subsets of basic relations. *Artificial Intelligence*, 106:205–219, 1998.

[20] K. Erol, R. Levy, and J. Wentworth. Application of agent technology to traffic simulation. In *Proc. Complex Systems, Intelligent Systems and Interfaces, Nîmes, France, May 1998*, 1998.

[21] Andrew U. Frank. Map algebra extended with functors for temporal data. In J. Akoka et al., editor, *Perspectives in Conceptual Modeling: ER2005 Workshops CAOIS, BP-UML, CoMoGIS, eCOMO, and QoIS*, pages 194–207, Berlin, 2005. Springer.

[22] Christian Freksa. Temporal reasoning based on semi-intervals. *Artificial Intelligence*, 54:199–227, 1992.

[23] A. P. Galton. A critical examination of Allen's theory of action and time. *Artificial Intelligence*, 42:159–88, 1990.

[24] A. P. Galton. *Qualitative Spatial Change*. Oxford University Press, Oxford, 2000.

[25] A. P. Galton. Fields and objects in space, time, and space-time. *Spatial Cognition and Computation*, 4(1):39–67, 2004.

[26] A. P. Galton. Dynamic collectives and their collective dynamics. In D. M. Mark and A. G. Cohn, editors, *Spatial Information Theory*. Springer, 2005.

[27] 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.

[28] A. P. Galton. Operators vs arguments: The ins and outs of reification. *Synthese*, 150:415–441, 2006.

[29] A. P. Galton. Experience and history: Processes and their relation to events. *Journal of Logic and Computation*, 18(3):323–340, 2007.

[30] A. P. Galton. Spatial and temporal knowledge representation. *Earth Science Informatics*, 2009. DOI 10.1007/s12145-009-0027-6.

[31] A. P. Galton and M. Worboys. Processes and events in dynamic geo-networks. In M. A. Rodríguez, I. F. Cruz, S. Levashkin, and M. J. Egenhofer, editors, *GeoSpatial Semantics: Proceedings of First International Conference, GeoS 2005, Mexico City, Mexico, November 29-30, 2005*, pages 45–59. Springer-Verlag, 2005.

[32] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, New York, 1979.

[33] Pierre Grenon. Spatio-temporality in basic formal ontology: SNAP-SPAN, upper-level ontology, and framework for formalization, Part I. Technical Report 05/2003, Universitšat Leipzig Institute for Formal Ontology and Medical Information Science, 2003.

[34] Pierre Grenon and Barry Smith. SNAP and SPAN: Towards dynamic spatial ontology. *Spatial Cognition and Computation*, 4(1):69–104, 2004.

[35] T. Hägerstrand. *Innovation Diffusion as a Spatial Process*. University of Chicago Press, Chicago, 1967.

[36] Shyamanta M. Hazarika and Anthony G. Cohn. Qualitative spatio-temporal continuity. In Daniel R. Montello, editor, *Spatial Information Theory: Foundations of Geographic Information Science*, pages 92–107. Springer-Verlag, Berlin, 2001. Proceedings of International Conference COSIT'01.

[37] N. W. J. Hazelton, F. J. Leahy, and I. P. Williamson. Integrating dynamic modeling and geographic information systems. *URISA Journal*, 4(2):47–58, 1992.

[38] Mark Heller. *The Ontology of Physical Objects: Four-dimensional Hunks of Matter*. Cambridge University Press, Cambridge, 1990.

[39] Kathleen Hornsby and Max J. Egenhofer. Qualitative representation of change. In Stephen C. Hirtle and Andrew U. Frank, editors, *Spatial Information Theory: A Theoretical Basis for GIS*, volume 1329 of *Lecture Notes in Computer Science*, pages 15–33, New York, 1997. Springer-Verlag. Proceedings of International Conference COSIT'97.

[40] Kathleen Hornsby and Max J. Egenhofer. Identity-based change operations for composite objects. In T. Poiker and N. Chrisman, editors, *Proceedings of the 8th International Symposium on Spatial Data Handling*, pages 202–213. International Geographical Union, 1998.

[41] Kathleen Hornsby and Max J. Egenhofer. Identity-based change: a foundation for spatio-temporal knowledge representation. *International Journal of Geographical Information Science*, 14(3):207–224, 2000.

[42] Kathleen Hornsby and Max J. Egenhofer. Modeling moving objects over multiple granularities. *Annals of Mathematics and Artificial Intelligence*, 36:177–194, 2002.

[43] R. A. Kowalski and M. J. Sergot. A logic-based calculus of events. *New Generation Computing*, 4:67–95, 1986.

[44] A. Krokhin, P. Jeavons, and P. Jonsson. The tractable subalgebras of Allen's interval algebra. *Journal of the ACM*, 50(5):591–640, 2003.

[45] Gail Langran. A review of temporal database research and its use in GIS applications. *International Journal of Geographical Information Systems*, 3(3):215–232, 1989.

[46] Gail Langran. *Time in Geographical Information Systems*. Taylor and Francis, London, 1992.

[47] Gail Langran and Nicholas R. Chrisman. A framework for temporal geographic information. *Cartographica*, 25(3):1–14, 1988.

[48] P. Laube, S. Imfeld, and R. Weibel. Discovering relative motion pattersn in groups of moving point objects. *International Journal of Geographical Information Science*, 19(6):639–668, 2005.

[49] X. Li and A. G. Yeh. Modelling sustainable urban development by the integration of constrained cellular automata and GIS. *International Journal of Geographical Information Systems*, 14(2):131–152, 2000.

[50] Gérard Ligozat and Jochen Renz. What is a qualitative calculus? A general framework. In Chengqi Zhang, Hans W. Guesgen, and Wai-Kiang Yeap, editors, *PRICAI 2004: Trends in Artificial Intelligence, 8th Pacific Rim International Conference on Artificial Intelligence, Auckland, New Zealand, August 9-13, 2004, Proceedings*, volume 3157 of *Lecture Notes in Computer Science*, pages 53–64. Springer, 2004.

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

[52] Damir Medak. Lifestyles. In Andrew Frank, Jonathan Raper, and Jean-Paul Cheylan, editors, *Life and Motion of Socio-economic Units*, pages 139–153. Taylor and Francis, 2001.

[53] J. Mennis, J. Leong, and R. Khanna. Multidimensional map algebra. In *Proceedings of the 8th International Conference on GeoComputation*, 2005.

[54] J. Mennis, R. Viger, and C. D. Tomlin. Cubic map algebra functions for spatio-temporal analysis. *Cartography and Geographic Information Science*, 32(1):17–32, 2005.

[55] R. Milner. *Communicating and Mobile Systems: The π-calculus*. Cambridge University Press, Cambridge, 1999.

[56] Marc Moens and Mark Steedman. Temporal ontology and temporal reference. *Computational Linguistics*, 14:15–28, 1988. Reprinted in Mani, I., Pustejovsky J, and Gaizauskas, R, *The Language of Time: A Reader*, Oxford, 2005.

[57] A. P. D. Mourelatos. Events, processes, and states. In Philip Tedeschi and Annie Zaenen, editors, *Tense and Aspect*, pages 191–212. Academic Press, New York, 1981.

[58] Philippe Muller. Space-time as a primitive for space and motion. In Nicola Guarino, editor, *Formal Ontology in Information Systems*, pages 63–76. IOS Press, Amsterdam, 1998.

[59] B. Nebel and H.-J. Bürckert. Reasoning about temporal relations: A maximal tractable subclass of Allen's interval algebra. *Journal of the ACM*, 42(1):43–66, 1995.

[60] A. Ohgai, Y. Gohnai, and K. Watanabe. Cellular automata modeling of fire spread in built-up areas—a took to aid community-based planning for disaster mitigation. *Computers, Environment and Urban Systems*, 31:441–460, 2007.

[61] D. O'Sullivan. Exploring spatial process dynamics using irregular cellular automata. *Geographic Analysis*, 33(1):1–18, 2001.

[62] Donna Peuquet and Niu Duan. An event-based spatiotemporal data model (estdm) for temporal analysis of geographical data. *International Journal of Geographical Information Systems*, 9(1):7–24, 1995.

[63] Donna J. Peuquet. A conceptual framework and comparison of spatial data models. *Cartographica*, 21(4):66–113, 1984.

[64] Donna J. Peuquet. It's about time: A conceptual framework for the representation of temporal dynamics in geographic information systems. *Annals of the Association of American Geographers*, 84(3):441–461, 1994.

[65] J. Portugali. *Self-Organization and the City*. Springer, berlin, 2000.

[66] Arthur N. Prior. *Past, Present, and Future*. Clarendon Press, Oxford, 1967.

[67] Arthur N. Prior. *Papers on Time and Tense*. Oxford University Press, Oxford, 2003. New edition, edited by Per Hasle, Peter Øhrstrøm, Torben Bräuner, and Jack Copeland.

[68] Jonathan Raper. *Multidimensional Geographic Information Science*. Taylor and Francis, London and New York, 2000.

[69] Jonathan Raper and David Livingstone. Development of a geomorphological data model using object-oriented design. *International Journal of Geographical Information Systems*, 9:359–83, 1995.

[70] Martin Raubal. Ontology and epistemology for agent-based wayfinding simulation. *International Journal of Geographical Information Science*, 15(7):652–665, 2001.

[71] Craig W. Reynolds. Flocks, herds, and schools: A distributed behavioural model. *Computer Graphics*, 21:25–34, 1987.

[72] T. C. Schelling. Models of segregation. *American Economic Association Papers and Proceedings*, 59:488–493, 1969.

[73] John R. Searle. Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3):417–457, 1980.

[74] W. Shi and M. Y. C. Pang. Development of voronoi-based cellular automata—an integrated dynamic model for geographical information systems. *International Journal of Geographical Information Science*, 14(5):455–474, 2000.

[75] Theodore Sider. *Four-dimensionalism: An Ontology of Persistence and Time*. Oxford University Press, Oxford, 2001.

[76] R. T. Snodgrass. Temporal databases. In A. U. Frank, I. Campari, and U. Formentini, editors, *Theories and Methods of Spatio-Temporal Reasoning in Geographic Space*, pages 22–64, Berlin, 1992. Springer.

[77] Masanao Takeyama and Helen Couclelis. Map dynamics: integrating cellular automata and gis through geo-algebra. *International Journal of Geographical Information Science*, 11:73–91, 1997.

[78] M. Thériault, C. Claramunt, and P. Y. Villeneuve. A spatio-temporal taxonomy for the representation of spatial set behaviours. In M. Böhlen amd C. Jensen and M. Scholl, editors, *Spatio-Temporal Database Management*, pages 1–19. Springer, 1999.

[79] W. R. Tobler. Cellular geography. In S. Gale and G. Olsson, editors, *Philosophy in Geography*, pages 379–386. D. Reidel, Dordrecht, 1979.

[80] C. D. Tomlin. *Geographic Information Systems and Cartographic Modeling*. Prentice Hall, Englewood Cliffs, NJ, 1990.

[81] P. M. Torrens. Process models and next-generation geographic information technology. *ArcNews Online*, Summer 2009. http://www.esri.com/news/arcnews/summer09articles/process-models.html.

[82] P. M. Torrens and I. Benenson. Geographic automata systems. *International Journal of Geographical Information Science*, 19(4):385–412, 2005.

[83] Peter van Beek. Reasoning about qualitative temporal information. *Artificial Intelligence*, 58:297–326, 1992.

[84] Zeno Vendler. Verbs and times. In *Linguistics and Philosophy*, chapter 4. Cornell University Press, Ithaca, 1967. Reprinted in Mani, I., Pustejovsky J, and Gaizauskas, R, *The Language of Time: A Reader*, Oxford, 2005.

[85] M. Vilain and H. Kautz. Constraint propagation algorithms for temporal reasoning. In *Proceedings of AAAI-86*, pages 377–382, San Mateo, CA, 1986. Morgan Kaufmann.

[86] J. Wainwright. Can modelling enable us to understand the rôle of humans in landscape evolution? *Geoforum*, 39(2):659–674, 2008.

[87] R. White and G. Engelen. Cellular automata and fractal urban form: a cellular modelling approach to the evolution of urban land-use patterns. *Environment and Planning A*, 25:1175–1199, 1993.

[88] M. F. Worboys and K. Hornsby. From objects to events: GEM, the geospatial event model. In M. Egenhofer, C. Freksa, and H. Miller, editors, *Geographic Information Science: Proceedings of the Third International Conference (GIScience 2004)*, pages 327–344, Berlin, 2004. Springer.

[89] Michael Worboys and Matt Duckham. *GIS: A Computing Perspective*. CRC Press, Boca Raton, Florida, USA, 2004. Second Edition.

[90] Michael F. Worboys. A unified model of spatial and temporal information. *The Computer Journal*, 37(1), 1994.

[91] Michael F. Worboys. Event-oriented approaches to geographic phenomena. *International Journal of Geographical Information Science*, 19(1):1–28, 2005.

[92] N. J. Yattaw. Conceptualizing space and time: A classification of geographic movement. *Cartography and Geographic Information Science*, 26(2):85–98, 1999.

[93] May Yuan. Representing complex geographic phenomena in GIS. *Cartography and Geographic Information Science*, 28(2):83–96, 2001.