Aspects of Space

Marcus Kracht
Universität Bielefeld, Germany

Recommended Citation

MARCUS KRACHT
Universität Bielefeld

ASPECTS OF SPACE

ABSTRACT: It is argued that spatial expressions come together with an encoding of the space called "aspect", which changes as we climb up the syntactic tree. The changing nature of aspect is necessary in order to simplify the meanings of elements. What appears to be a rather peculiar property of an element will be perfectly natural once we acknowledge that the elements compute on the space viewed in a particular way. Coordinates are always rooted in the landmark, for example. Thus, for the purpose of the distinction between static and dynamic it is not the "absolute" motion of the figure that counts, but the motion relative to the landmark.

1. INTRODUCTION

Model theoretic semantics (MTS) is in disarray. For some it has lost relevance since it only deals with arbitrary models (in addition to the real one, whatever that may be), for others it is utterly trivial and fails to address the hard questions of semantics; and for yet others it is just a sort of *glasperlenspiel*. As much as I wish to agree with these points the problem is that the concerns over MTS will not go away by ignoring MTS as a result. We are still facing the problem of saying what the meaning of a complex expression is and how we get it by manipulating the meanings of its parts. And we still have to admit that "snow is white" is true if (and only if) snow is (in fact) white.

However, advocates of model theoretic semantics have become complacent. Rather than addressing these problems head on they have simply stayed on course, trusting in the superiority of their approach. Worse still, when it comes to recent additions to our stock of primitives, like events and situations, nobody bothers to tell us what the models look like, so that all we are left with is formulae without any indication how the model actually makes them true. For events cannot be independent of the individuals and the actions that these individuals perform, and these dependencies between the elements of various categories capture essential properties of them that the models need to respect. Instead of trying to improve the situation, type theoretic semantics has become a free for all: Invent a new type, add it to the primitives – and you are all set.

This is unsatisfactory. If anything, we would like to expand our semantic vocabulary and rest assured that we can not only write down formulae but also say something significant about the models.

The cognitive turn in linguistics has also added another kind of model to worry about: the mental model. That which is in our heads. It is unfortunate that some people like to insist that this is the only kind of model we need to study. For, if there is any sense in talking about reality, we ought to be interested in knowing what that reality actually is. There can be a host of different mental models, at the end of the day disagreements can only be settled by turning to outside reality.

Of course, one may object that each one of us sees that reality differently. Moreover, some areas of semantics cannot be said to be about the outside world as it is here and now, for example as propositional attitudes, and – to some extent – intensionality. Still, there are areas where all of this seems not to get in our way. Space is one such area. We seem to be firmly committed to the view not only that there are some intersubjectively testable intuitions about space that can tell us whether you or I are wrong about certain claims, but also that there is a science whose success is beyond doubt, which aims to uncover the real nature of space. That science is called physics. If physics is simply a big blunder and there is no objective space then I would like to know what it is that explains how physics can deliver with such precision.

In other words: we know in enough detail what the real world looks like when it comes to space. Cognitive linguistics and other empirical

sciences on the other hand have, based on this knowledge, also shown how that space is coded in our minds. Maybe we do not know enough yet. But, as I shall argue here, we know so much already that it is possible to show how model theoretic semantics and linguistics as well as the cognitive sciences can learn from each other in a nontrivial way. In particular, I will show that the meaning of a spatial expression involves a sophisticated algorithm of (re)coding the space that makes claims of compositionality quite nontrivial.[1]

2. PRELIMINARY REMARKS

I shall be concerned here with the meaning of locative PPs like

(1)   `on the mat`

Standardly, they involve two objects, called here "figure" and "ground", where these have their usual meaning. The PP locates the figure in terms of the ground at the time of the event.

The argumentation used here is as follows. Expressions can be translated into a mental model (or cognitive map) or into a real model (as MTS uses it, but with emphasis on physical reality). I shall scrutinize what the semantics can legitimately say if it takes the expressions to denote something in the real world. So, model here means at least the real physical model. Second, we shall use the linguistic structure of these expressions to arrive at a componential analysis that is supposed to reveal a cascade of cognitive processes that change what I call the *aspect* of the space. This allows us to say how the real model is mapped into the cognitive model. Semantics is then derivative: by translating expressions into cognitively informed maps, the objective content can be found by undoing this cascade of codings.

The main source is Kracht (2008), where a detailed account of the model theory is given. The present paper will not be as technical. This is because I now think it is actually more important to understand the philosophical and physical background, and that that can be explained without a lot of mathematical machinery. The present discussion may actually serve as a motivation to read the paper just quoted.

I shall not deal much with syntax or morphology, though I shall make some occasional remarks on these matters.

The methodological issues can be used to analyze other expressions as well (verbs of stasis and motion, spatial adjectives etc), but I did not want to venture into that territory. It would make the exposition less coherent (and hence more confusing).

Also, much of what I say can equally be said about time. Yet, time will figure only marginally in this discussion.

3. MOTIVATION

There is a modern mythology around how space functions that is incompatible with both physics and cognition, though as we shall see it is incompatible in different ways. And this is the mythology of an absolute space, with coordinates to locate objects, in the way we are being taught in high school. This space appears to us nowadays – thanks also to modern technology – so much more natural than everything else, though that naturality is more apparent than real. I shall argue here that it has very little to offer for semantics.

The idea of an absolute space is that space is "at rest", providing a container (called *aether*) in which objects are being placed and can move around. Fix an origin in this aether, and some coordinate frame, and points can henceforth be equated with triples of numbers (or quadruples if you add time), also called *vectors*. So, the sentence

(2)   `The cat is near the mat.`

is true if and only if the set of differences $\vec{x} - \vec{y}$ where $\vec{x}$ is a vector (= triple of numbers) from the set of cat-points and $\vec{y}$ a vector from the set of mat-points, contains some vectors of relatively small length (however that may be spelled out in detail).

But we may ask: for whom could such a definition be useful or even explanatory? Not for humans, since they do not know how to calculate these vectors. Not for physicists, since they will insist that the coordinates associated with points are not unique, not even the differences (they are not immune to rescaling, for example). Not for cognitive scientists, since they will have the same trouble as the others identifying such "absolute" coordinates.

Note. The vector semantics of Zwarts & Winter (2000) can easily be reformulated in terms of points in a space. That the authors use the word "vector" where I prefer "point" is actually inessential for their aim.

4.   ABSOLUTE SPACE

Physicists will say that a physical theory must be invariant under translation. By that they mean that if the laws are expressed in terms of coordinates then moving the origin of the coordinate frame somewhere else must not change the expression of the laws. In fact, rather than expressing location in terms of absolute position we can only express it in terms of relative position. But relative to what? The aether theory would say: in relation to the space itself. Modern physics says: in relation to other objects only. To see the point of the exercise, try identifying your location in a space that is absolutely empty. (To appreciate the scale of the problem, ask yourself how you can find out where you are in a desert.) There is actually nothing that you can do, since locating a point in space can be done *only* by means of objects. And we know of no object that is fixed in space; more exactly, we have found no evidence and no diagnostic that would point to the existence of such objects. What goes for humans, however, goes for the physical objects as well. There is no reality to space. It is constituted by matter (and energy) alone. As strange as that sounds.

    The physicist Ernst Mach had a similar problem with respect to rotation. He asked: When there is no other body in space except for a single sphere, then how can you tell it is rotating? Even worse: how can *the sphere* tell it is rotating? And his answer was that it can't. There is nothing (=no thing) by means of which the sphere could diagnose its own rotation. It might just as well assume it is not rotating at all. Hence the forces that act on it when rotating cannot be attributed to its rotating in absolute space. And so we get the following principle, attributed by Einstein to Mach, maybe erroneously so, see Norton (1995). (The formulation is the one given in Wikipedia.)

> [Mach's Principle]
> Any phenomena that would seem attributable to absolute space and time (e.g. inertia, and centrifugal force) should instead be seen as emerging from the large scale distribution of matter in the universe.

We know the rest: the notion of absolute space was dealt a blow with the Michelson-Morley Experiment in 1887. It shows essentially that absolute space does not exist. For it establishes that no matter in what direction we send a light beam, its velocity is the same. Assuming that the earth is in motion (at 30km/s around the sun), one would have to expect a difference of up to the speed of the earth. But none was observed. So, all coordinate frames are just arbitrarily selected. Coordinates are not intrinsic properties of space points nor of objects. And so on.[2]

    What is true in physics is true a fortiori for humans. We can have no sensors that detect what coordinates we inhabit at a given moment – this is a physical impossibility. Hence to ask for a procedure like the one given above – taking triples of numbers for locations, calculating differences thereof and so on – is asking for the impossible. On the other hand, there does seem to be a clear difference between someone who is moving and someone who is not. How can that be?

    Before I move on let me say that the mathematical theory of space as exposed here actually *does* assume an aether. It says that there is, say, a structure $\mathbb{O} = (O, C)$, where O is a set of *points*, and C some relational structure on these points. I shall occasionally reveal what C must minimally contain in order to get the semantics off the ground. For example, we will have a notion of *distance*. Given two points o and o', $d(o, o')$ is a positive real number. (There are issues of calibration, to which I shall return.) I think that even if physical space is different, cognitively we do think of space as an inhabited aether. The problem is how it gets fixed in practical terms.

5.   MORE OR LESS ABSOLUTE SPACE

The idea is that motion is constituted by a change in distance from a certain object. Not just any object, of course. For us, the earth is a suitable reference point. We take it to be at rest and calculate everything

with respect to it. To do that we take landmarks, whose position we deem to be fixed relative to it (and consequently relative to each other), like mountains, rivers, trees and houses. The stars and the sun could be used, but care must be exercised as their relative position is changing slowly over time. There is thus a hierarchy of *stability* that we exploit to tell whether someone or something is *moving* or *at rest*.

In terms of Cartesian coordinates we need an *origin*, a point that the location of all other points makes reference to. Ideally, we would like that point to be fixed once and for all. But that is not how it works. It is unrealistic to assume a single absolute origin for everyone at all times. This is impractical. To be able to tell whether (2) is true we do not really need to know whether either of the two, the cat or the mat, are moving. We do not need to be able to calculate coordinates if only we can calculate *distances*. This, it seems, we can do. The notion of nearness does not answer to positions but only to distances.

> [Step 1]
> In a locative expression, the location of the figure is calculated only with respect to the location of the ground. We take the latter as the *origin of space*.

What does that mean? Consider the expression "near the mat" in (2). It involves the ground object "mat". In deciding whether the figure is in the relation of nearness, we may as well consider the mat as the origin of a coordinate frame and calculate everything with respect to it. This origin is quite a temporary one: it is agreed on by the PP "on the mat" and used within the sentence. Outside of it, any other origin may be used.

At this point the way we picture the space, what I call its *aspect*, has changed. We no longer see it as a collection of points with some properties, we see it as a collection where one point, the center c of the mat, is being distinguished. Our new structure is (O, c, C), where c is the *origin*. But more is going on. Instead of seeing our points as points we may also now refer to them using c. We shall change the way we address them. We *recode* them. The point p will be recoded as the *vector* $\overrightarrow{(c, p)}$. This vector is a particular set of pairs of space points, namely those pairs that have same length and direction as $\overrightarrow{(c, p)}$. This is the familiar notion of a vector used in physics (e.g. the force vector). In

a Euclidean space these are the vectors that are being used in *translations*. Vectors have a structure: they can be added, for example. Thus, as we set the origin, we can address each point p by its vector $\overrightarrow{(c, p)}$. The function $f_1 : p \mapsto \overrightarrow{(c, p)}$ is one-to-one (= injective). Its image is a structure $\mathbb{V} := (V, +, -, \vec{0}, \hat{C}, c)$, where $+, -$ are the operations of addition and inverse of vectors, $\vec{0} := f_1(c)$, and $\hat{C}$ is the image of the relations of C under the function f_1. Notice that by corollary we can form multiples $n \cdot \vec{v}$, where $\vec{v}$ is a vector and n an integer. Assuming a suitable topology on the vector space so that the operations become continuous, we can define a multiplication $r \cdot \vec{v}$, where r is a real number. Thus we have a real vector space of dimension 3, which I also denote by $\mathbb{V}$. Before moving on let us note that I have added c to the structure. Otherwise we could not calculate the aspect.

Typological Note. Spatial systems make use of obvious landmarks such as: rivers ("upstream", "downstream") and shorelines ("seawards", a general feature of Austronesian spatial talk, see the collection (Bennardo 2002). Others use the sun (east is where the sun rises, west is where it sets) or winds (e.g. the monsoon).

6. CALIBRATING THE SPACE

In the first step we have allowed the origin of space to be adjusted locally. This sidesteps the calculation of absolute coordinates. However, we will still have to calculate relative coordinates. This will require two more things: we need to fix directions and we need to fix unit vectors in these directions.

We do the second first. Consider again the notion of nearness. Suppose we are being asked whether the moon would be near the earth if it were 100 km away from earth. Presumably we would say yes. But a cat that is 100 km away from a mat is definitely not near it. So, nearness actually does not answer to absolute distance but only to relative distance, which is dimensionless, that is, a number. One solution would be to calculate the size of the ground object, the size of the figure and then see whether the distance is a large or small multiple of these numbers. (I will not say much about the exact details of how exactly "near" is computed. Something like the larger of the diameters of figure and ground seems to go into it, minimally.) The *diameter* of a region r, $d_\varnothing(r)$, may

in first approximation be set to the supremum of the numbers $d(o, o')$, $o, o' \in r$. On the other hand, $d(r, s) = \inf\{d(o, o') : o \in r, o' \in s\}$, which is zero for closed sets if and only if the intersection is nonempty.[3] So, for nearness we look at the ratio $d(r, s)/\max\{d_\varnothing(r), d_\varnothing(s)\}$. This is a real number. So we have a *scale*.

This only works for regions. There is another concept of scale, call it *external*, where rather than looking at the diameters of the objects, we look at some external measure (humans, etc). Thus, I may say that I am near your house just because going there will not take long, not because the distance is a small multiple of the size of the house. But comparatively the distance might be large if only the bicycle and the house are considered.

[Step 2]
The number 1 is assigned a set of vectors K, called the *unit ball*. The members are called *unit vectors*. The length of a given vector $\vec{v}$ is that number λ such that the multiple $1/\lambda \cdot \vec{v}$ is a unit vector.

This means that given $\vec{v}$, we find a $\vec{u} \in K$ such that $1/\lambda \cdot \vec{v} = \vec{u}$. This $\vec{u}$ is unique (by assumption on K) and therefore λ is unique as well. Hence $\vec{v} = \lambda\vec{u}$, so the length counts how many multiples of the unit vector $\vec{u}$ we need to take to get $\vec{v}$.

When we have set up a scale we can talk of *lengths*. Each vector is given a length, which is a real number. The aspect of the space has changed again. We now have, in addition to the group structure on the vectors, also a function assigning length. Thus we have what is called a *normed vector space* or a *metric vector space*. (Norms give lengths, $||\vec{v}||$ is the length of $\vec{v}$. Metrics give distances: $d(\vec{v}, \vec{w})$ is the distance between $\vec{v}$ and $\vec{w}$, and is identical to $||\vec{v} - \vec{w}||$.) If the metric is rotation invariant we may just as well single out just one vector as a unit vector instead of the entire set K. (Think of the prototype meter in Paris.)

7.   SETTING UP THE COORDINATE FRAME

In addition to concepts like "near" we also have those that point in a certain direction, like "north", "left", "seaward", and so on. When we

say

(3)   Edinburgh is north of London.

we position Edinburgh in relation to London as the origin of a coordinate frame. (Mind you, we can do the converse too by saying that London is south of Edinburgh...) Doing this however requires singling out a particular vector called "north". It is unique, since it must be of length 1 and point northward. Similarly,

(4)   John is to my left.

asserts that John is found by following from myself as origin in the left-hand direction.

Indeed, so many spatial PPs involve such expressions that we can say that the most important step (judging also from the literature on spatial language) is setting up the coordinate frame. We do that by positioning what is called a virtual observer in the origin. This observer defines the six cardinal directions, *right*, *left*, *front*, *back*, *up* and *down*. Obviously, only three of them are needed. Now that we have the center c and the unit ball K, we choose three vectors, $\vec{e}_1$, $\vec{e}_2$ and $\vec{e}_3$ to represent the three cardinal directions, called *front*, *right* and *up*. We assume them to be linearly independent. Once this is done, for each vector $\vec{v}$ there are unique numbers γ_1, γ_2 and γ_3 such that

(5) $\vec{v} = \gamma_1 \cdot \vec{e}_1 + \gamma_2 \cdot \vec{e}_2 + \gamma_3 \cdot \vec{e}_3$

The triple $(\gamma_1, \gamma_2, \gamma_3)$ is the coordinate vector of $\vec{v}$. The map $f_2 : \vec{v} \mapsto (\gamma_1, \gamma_2, \gamma_3)$ is a bijection from $\mathbb{V}$ onto $\mathbb{R}^3$. (By now, $\mathbb{V}$ and $\mathbb{R}^3$ have assumed the structure of normed vector spaces.)

One small wrinkle must be addressed. In order for this to work we need some additional structure, a *scalar product* $\langle \cdot, \cdot \rangle$ compatible with our norm. Such a structure is called an *inner product space*. Scalar products are bilinear forms, that is, maps from $V \times V$ to $\mathbb{R}$ which are linear in both arguments. They make it possible to define angles between vectors, and, importantly, *orthogonality*. We say $\vec{v}$ is *orthogonal* to $\vec{y}$ if $\langle \vec{x}, \vec{y} \rangle = 0$.

This is the second step. The function f_2 is an aspect of the vector space $\mathbb{V}$. Thus, by composition, $f_2 \circ f_1$ is an aspect of the original space $\mathbb{O}$. At this point we can address the points as triples of numbers.

Typological Note. The problem of setting up coordinate frames has been recognized for a long time. I only refer to Levinson (1996) and Levinson (2003). The idea that there is a special morpheme, called Ax-Part, responsible for hosting the frame has been proposed by Svenonius, see e. g. Svenonius (2008). Expressions for (absolute) cardinal directions need denote orthogonal directions, see Palmer (2002). However, this just means that the expressions do not share the same coordinate frame, that is, a single language can use different methods for constructing coordinate frames even for absolute directions.

8. POSITIONING THE FIGURE

The space is set up: we have put on two lenses in sequences through which we can take a view. We are now ready to position the figure. Look at the sentence

(6)   The cat is on the mat.

Here, as with (2), we have the mat as the origin of the space. In addition, to get at the meaning of "on", we make use of the cardinal directions from the coordinate frame. We select the direction "up". Moreover, the distance of figure from ground is 0 (they are touching). Otherwise we should have used "above". Actually, I am undecided whether the direction to use is always $\vec{e}_1$, with the others being ancillary at best, or whether we take the idea of the virtual observer seriously and let $\vec{e}_1$ be "front", so that "up" is $\vec{e}_3$. For the purpose of this paper I choose the second option.

The point of the coordinate frame can now be put quite succinctly. After all recoding we only need to look at the coordinate γ_3 for points of the figure. If its infimum is 0, the sentence is true, and false otherwise. The recoding has made complex computations unnecessary.

Typological Note. Above AxPart we find another head, the localizer. There are case systems that show clearly the presence of a distinct head; Avar, Finnish, Hungarian, Tsez case systems are of that kind (see Kracht (2002)).

9. TRACKING THE FIGURE

When Michael Schumacher is trying to get ahead of Mikka Häkkinen in a race we cannot really assume that Mikka is standing still. Neither, of course, is Michael. After all this is a race. But at the same time, if he is unsuccessful, we would assert

(7)   Schumacher stayed behind Häkkinen.

This suggests that nothing is changing in terms of position. This apparent contradiction has a simple explanation.

Positioning the figure is done at a particular time point. Events usually take some time to unfold, so we do a positioning of the figure at each time point during the entire event. What is crucial is that the ground is the origin of the coordinate frame, regardless of whether or not it is in motion. Thus, motion of the figure is described relative to that of the ground (remember Mach's Principle). We may or may not be aware of the fact that the ground is in motion (cognitively speaking, not physically), but the linguistics of spatial PPs prescribe that they record motion relative to ground.

(7) is not expressing change in the sense of the local PP "behind Häkkinen". During the event time, Schumacher is locked into that position. At the same time we can assert that he is racing past the visitors at 250 km per hour. Still, the truth of "be behind Häkkinen" remains unchanged.

Again, the trick is done by the aspect. It is by using the aspect as induced by the PP that evaluation of the sentence becomes trivial. Just let the coordinate frame be such that Häkkinen's car is the center, and $\vec{e}_1$ points in direction of its motion (= "front"). Then $-\vec{e}_1$ points towards Schumacher *throughout event time*. That's it.

Typological Note. Niikanne (2003) notes that in Finnish there are three types of Ps: those that require motion of the ground (*edellä* "in front of"), those that require the ground to be at rest (*edustalla* "in the front of"), and those that are indifferent (*edessä* "in front of").

There is a separate head that talks about the change of truth. Fong (1997) deals with such heads in terms of phase quantification. See Kracht (2008) for an expansion of that idea. Languages with heads devoted to expressing phase quantification include Finnish and Hungarian, see Kracht (2002).

10. MODEL THEORETIC SEMANTICS AGAIN

So now everything seems to fall in place. As we move upwards in the hierarchy, space gets coded up in different ways, and the meanings of the newly added heads just take these encoded meanings as inputs. However, there is a problem to solve. As I noted in Kracht (2008), we shall need a double accounting to keep track of locations. One layer contains information about the objects under the encoding scheme, the other contains information about the encoding function itself.

This essentially means that simplicity is gained by outsourcing the complications into some clever encoding. This can hardly be called a solution. Moreover, the truth conditions can only be effectively determined if both the representation *and* the encoding are simultaneously known.

I have as yet no answer to this complaint. However, I wish to point out that when it comes to our mental models as well as our way of accessing real space we are always dependent on a certain coding, and this coding is changing quite frequently. When pointing at a red dot on a map saying: "This is Berlin." we are not actually pointing at the city of Berlin; we can be said to point at it only under the tacit assumption that there is a code relating points on paper to points in real space. Although we may not even be aware of what exactly that coding is, the reference is nevertheless successful, at least for practical purposes. We can see that this is what is going on by looking at situations where things go wrong. We may think, for example, that New Zealand is close to Australia since the map visually tells us so. For on the map Australia has the size of France. To correct this, we must realize that the two maps are of different scales. Of course, this assumes that we are talking about an absolute notion of nearness, that is, *distance* rather than *nearness*. When we mean nearness in relative terms, scale does not matter.

Not only that: we are unable to look at space as it really *is*, we picture locations by means of internal maps that we can resize, turn around and so on. These maps show a coherence also with respect to each other. The coding functions discussed above treat these maps as something that gives names to objects. The more we code, the more complicated it gets. Viewed from inside, it is not the addresses that get complicated, it is the outside world. When Michael Schumacher sits in his car, he can see Mikka right in front of him. He has no trouble track-

ing his location. Only the viewers are experiencing some difficulties. From their perspective, to decide whether Michael makes an attempt at overtaking Mikka requires good eyesight and attention. One must not lose track of either car since one needs to calculate some relative differences in their positions.

Ultimately, we need to acknowledge that Ernst Mach was right. When I return to the beautiful lecture hall of the university of Riga on the next day, the reason that I confidently say that I have been in that location before is not that I can recognize the space points in the aether. It is not that I have an ability to somehow track my physical motion in the metrical space time. We know that I can't. What I really want to say therefore is that I recognize *that lecture hall*. I remember the location by way of remembering the things that inhabit it. Just as I remember events not by the time stamp they get but by the order they have in the overall chain of events. That I was here yesterday is again not saying that some absolute portion of time passed between then and now of which I possess direct evidence. It is saying that there was night in between them.

We get something of a Kantian theme here. Outer space is a mental construct: we construct locations by means of objects that inhabit them. That it may ultimately be easier to use that construct as a container for dumping our mental furniture is just a mental convenience. However, from the standpoint of evaluating linguistic expressions, we need to go all the way. Essentially, we can evaluate expressions *not thinking of them as codes of anything* but thinking of them as objects in their own right. We thus have no single model that explains everything (transcendental space, if you like), but a progression of them, between which we can shift in systematic and predictable ways.

11. ASPECT

The term "aspect" has a long history in linguistics. So why use it for something entirely different? The answer is that it is actually *not* different. As far as I can see the same problems of positioning an event arise in the domain of time. We speak of "progressive aspect" when some event "is still going on". But what does that mean? From a model theoretic perspective that makes no sense: events are not going on, they

have a temporal trace. Just like objects aren't simply present, they have a location.

What is true of locations is true of time points as well: they are not absolute, only relative. There is no unique "zero" (choose between *ab urbe condita* (birth of Rome), birth of Christ, first Olympic games, January 1, 1970 (the beginning of "the epoch" of Linux), and many more). That an event is going on can be said only relative to some other event (Klein 1994). Thus, when speaking of time, the same process of picking a reference point and scaling the measure of duration can be observed. And since they go under the name "aspect", I think calling this phenomenon for space "spatial aspect" is justified.

12.   CONCLUSION

When faced with spatial expressions, we are tempted to follow the traditional path of Cartesian coordinates: locate points via coordinate triples, and use those in the computation of local relations. However, such a procedure is only useful when such coordinate triples are readily available. The fact is that they are not. Bypassing the need to compute coordinates of both figure and ground is the key to understanding the peculiar nature of PPs, for example, that they do not answer to absolute motion of the figure, only to relative motion with respect to the ground.

Notes

[1]I wish to thank Jurgis Skilters and Michael Glanzberg for their enthusiasm and for providing me feedback on this paper. Also, many thanks András Kornai and to the participants in the workshop, in particular Yoad Winter and Peter Gaerdenfors, for their critical evaluation.

[2]I mention here in passing that the matter is far from settled. Relativistic physics uses the notion of an inertial system, which is a system that is in constant motion. In an accelerated system, time slows down, for example. But how can a system tell it is inertial (see Esfeld (2012))? Thanks to Michael Glanzberg for bringing up this point.

[3]I take regions to be path connected closed sets. Contact between them requires overlap. The complications of this view are discussed in Smith (2007) and need not concern us here.

References

Bennardo, Giovanni (ed.). 2002. *Representing Space in Oceania—Culture in Language and Mind*. Pacific Linguistics. Canberra: Australian National University Press.

Brown, Cecil H. 1983. 'Where Do Cardinal Direction Terms Come From?' *Anthropological Linguistics* 25: 121–161.

Esfeld, Michael (ed.). 2012. *Philosophie der Physik*. Frankfurt am Main: Suhrkamp.

Fong, Vivienne. 1997. *The Order of Things: What Directional Locatives Denote*. Ph.D. thesis, Stanford University.

Klein, Wolfgang. 1994. *Time in Language*. London: Routledge.

Kracht, Marcus. 2002. 'On the Semantics of Locatives'. *Linguistics and Philosophy* 25: 157–232.

——. 2008. 'The Fine Structure of Local Expressions'. In Anna Asbury, Jakub Dotlacil, Berit Gehrke & Rick Nouwen (eds.) 'The Structure of Local P', 35–62. Amsterdam: John Benjamin's.

Levinson, Stephen C. 1996. 'Frames of Reference and Molyneux's Questions: Crosslinguistic Evidence'. In Paul Bloom, Mary A. Peterson, Lynn Nadel & Merrill F. Garrett (eds.) 'Language and Space', 109–169. Cambridge, Mass.: MIT Press.

——. 2003. *Space in Language and Cognition. Explorations in Cognitive Diversity*. No. 5 in Language, Culture and Cognition. Cambridge: Cambridge University Press.

Niikanne, Urpo. 2003. 'How Finnish Postpositions See the Axis System'. In Emile van der Zee & John Slack (eds.) 'Representing Direction in Language and Space', No. 1 in Explorations in Language and Space, 191–208. Oxford: Oxford University Press.

Norton, John D. 1995. 'Mach's Principle before Einstein'. In Julian Barbour & Herbert Pfister (eds.) 'Mach's Principle: From Newton's Bucket to Quantum Gravity', 9–58. Boston: Birkhäuser.

Palmer, Bill. 2002. 'Absolute spatial reference and the grammaticalisation of perceptually salient phenomena'. In Giovanni Bennardo (ed.) 'Representing Space in Oceania', 107–157. Canberra: ANU Press.

Smith, Sheldon. 2007. 'Continuous Bodies, Impenetrability, and Contact Interaction: The View from the Applied Mathematics of Continuum Physics'. *British Journal for the Philosophy of Science* 58, no. 3: 503–538.

Svenonius, Peter. 2008. 'Projections of P'. In Anna Asbury, Jakub Dotlacil, Berit Gehrke & Rick Nouwen (eds.) 'The Structure of Local P', 63–84. Amsterdam: John Benjamin's.

van Lambalgen, Michiel & Hamm, Fritz. 2005. *The Proper Treatment of Events*. Oxford: Blackwell.

Zwarts, Joost & Winter, Yoad. 2000. 'Vector Space Semantics: A Model–Theoretic Analysis of Locative Prepositions'. *Journal of Logic, Language and Information 9* 169–211.