Invited talk, *European Conference on Cognitive Science* , Saint-Malo 1995.

TIME AND SPACE IN COGNITIVE SYSTEMS

Christian Freksa and Bettina Berendt

Arbeitsbereich Wissens- und Sprachverarbeitung & Graduiertenkolleg Kognitionswissenschaft,
Fachbereich Informatik, Universität Hamburg,
Vogt-Kölln-Str. 30, 22527 Hamburg, Germany
e-mail: {freksa,berendt}@informatik.uni-hamburg.de

1 What is a cognitive system?

Before discussing the role of time and space in cognition, we will introduce the notion of a 'cognitive system'. The notion of a 'cognitive system' is suitable to overcome certain problems encountered with the notion of an 'intelligent system'.

1.1 The problem of defining intelligence

In classical views, intelligence was considered to be localised in the heart or in the brain of an organism. The view that intelligence is localised in the brain still dominates the discussions about the possible existence of an artificial intelligence: a computer is likened to a human being; the essence of the human being with respect to intelligence is considered to be the mind/brain system.

However, the restriction to consider brain/mind (or hardware/software) as the locus of intelligence has produced paradoxical results. A number of philosophical arguments have been advanced: best-known among these is probably Searle's "Chinese room argument" [Searle 80]. Imagine yourself locked in a room, equipped with a set of instructions (in English) how to manipulate all kinds of combinations of Chinese symbols. Occasionally, somebody slips a set of papers covered with Chinese symbols ('questions') under the door. After you have manipulated them according to the rules, writing down the results of these manipulations on further sheets of papers, you slip these 'answers' back through the door. How could anybody claim that you *understand* Chinese just because you hand back 'the right squiggles'? This parody of the Turing test [Turing 50] has been a highly controversial issue in AI, leading many to abandon of the 'strong AI' claim "that the appropriately programmed computer literally has cognitive states and that the programs thereby explain human cognition" ([Searle 80]: 417).

But even 'weak AI', stating that "the principal value of the computer in the study of the mind is that it gives us a very powerful tool" ([Searle 80]: 417), has increasingly come under attack. This model too assumes that the symbol manipulation process is the essence of intelligence, that an architecture having little or no parallelism and a very narrow channel of interaction with the environment can adequately model 'cognition'. Even taking this view, there has been a marked failure of coming up with a useful definition of intelligence

Figure 1: Cognitive Systems vs. AI-Systems

on the basis of purely computational criteria and a difficulty of matching “intelligent processes” with actually useful performance.

1.2 Cognitive systems vs. AI systems and abstract systems

As a consequence of this and due to the observation of more and more converging evidence from the various branches of Cognitive Science, attention these days is turned more towards investigating brains and computers in connection with their environments and the interfaces that link them together. Specifically, the properties of the sensors and actuators are more and more recognized as crucial for the knowledge structures involved in cognitive processes. Cognitive Science now investigates highly parallel systems with rich multimodal interactions with their environment (e.g. visual, acoustic, haptic). An emerging consensus focuses more on *practical intelligence*, which can loosely be defined as ‘the ability to do the right thing in a given situation’.

This approach implies considering cognitive processes in the context of the situations in which they take place. Cognitive performance is judged on the basis of the ability to solve tasks in the environment.

This typically involves perceiving and interpreting a situation, processing the resulting information, and applying the result of this process to the situation. The cognitive performance can then be judged by an external observer on the basis of a specification of the task and solution spaces (cf. fig. 1).

1.3 Situatedness: extreme and moderate views

In the AI community, however, it has turned out that this realisation cannot simply extend and be embedded into existing research. Some proponents of the *situated* approach envision a paradigmatic shift – the programmatic [Brooks 1991] is called “Intelligence without representation”. Now representations can be viewed as one of the very foundations of AI (as of most of the other Cognitive Sciences too!). But there are more moderate voices as well.[1] These attempt to allow representations formed by a situated cognitive system. It is in this tradition that this paper will argue.

1.4 Cognitive systems use knowledge, not only information

Since we are considering cognitive systems *embedded in their environment*, we also have to re-evaluate our notion of what the overall system uses to achieve intelligent behaviour.

[1] Even from Brooks’ side: [Mataric 1992]).

While in ‘classical AI’ it is admissible to use ‘pure, uninterpreted information’ (like the Chinese symbols in Searle’s example), a ‘cognitive system’ in our sense must have some idea about the connection of this information with structures in its environment. We will call this enriched notion ‘ knowledge’.

1.5 Cognitive systems employ different levels of explanation

This view of ‘knowledge’ is a decidedly *relative* notion, relative to the concrete situation in which it is to be used to solve a problem. Accepting this means rejecting the dichotomy of “I know it” vs. “I don’t know it”. Instead, it is the admission that we (or indeed any other cognitive system) will never completely understand anything in an absolute sense. Rather, understanding is relative to a given level of knowledge. There will always be different levels of explanation.

As an example, consider entering your office to find a stack of papers on your desk. If you are an orderly person and want to start work, you will consider this to be (explain it as) ‘some papers’ that have to be put somewhere else. If you want to find out whether it was your little daughter, who put them there when you left her alone in the office for a while, or a colleague, you will pay some more attention and see that there are ‘words’ (as opposed to squiggles). If you want to find out what it is your colleague deemed interesting to you, you will read the ‘text’. If you find out it is the draft for a new project proposal, you will also read between the lines and find out it is ‘a plan to revolutionise the whole department’.

And so on. The important point is that it is simply *not necessary* to consider or know these levels of explanation when you’re trying to clean up your desk – in fact, it would even be an *impediment* to your urgent work.

2 Why are time and space important for cognitive systems?

Time and particularly space have been a point of central interest for many branches of Cognitive Science for a long time.[2] Psychologists study spatial cognition – the cognition of visual space (e.g. the perception of images and imagery), the cognition of personal space (e.g. questions of reference frames, visual vs. haptic space) and the cognition of large-scale space (Ecological Psychology’s “Images of the City” [Lynch 1960] and the like).

The neurosciences study spatial cognition – particular neurological deficits of spatial cognition like deficits in the cognition of personal space (neglect), deficits in the cognition of locations (Balint-Holmes syndrome) and objects (cf. also the ‘what/where’ distinction, see [Ungerleider & Mishkin 1982]). Reference systems play a decisive role here too.

Research in these areas is not restricted to the study of *human* spatial cognition, but extends to animals as well [Gallistel 1990].

Linguistics studies spatial language – for example, spatial prepositions. Different reference systems explain what may have seemed inconsistent usage of prepositions before [Retz-Schmidt 1988].

Philosophers have been asking questions about possible structures of time and space for 2000 years.

---

[2]The discussion and particularly the references in this short sketch are intended to be purely exemplary!

An increasing number of interdisciplinary treatments are published – like [Landau & Jackendoff 1993]'s much-discussed article linking " "What" and "where" in spatial language and spatial cognition", to cite but one example.
AI started out with a Computer Science's mostly implicit treatment of time and space (see section 2.1). Realising that this leads to many anomalies, AI researchers then started developing explicit treatments of time and space. It is from this point of view that we shall present our concepts of 'time and space in cognitive systems' here. The article is not intended to be an overview over all existing approaches, but rather the advocation of a particular point of view whose importance is growing.

2.1 Ubiquity of time and space, or: Why are we interested in time and space?

We will consider the implications these general relationships have for a domain of particular importance to cognitive systems. Every action a cognitive system takes is situated in time and space. In humans, these seem to be even more important: spatial and temporal metaphors are used in a huge number of situations and domains [Lakoff 1987]: Do you *follow* our argument?

In computer science, we usually describe individual aspects of space (e.g. length or height) or we abstract from space altogether. Time is implicitly maintained in the dynamics of computer programs. But when considering the input/output characteristics of a computer program, the aspect of time is ignored.
We argue that by decomposing complex spatial structures into individual aspects we lose valuable cognitive power. The reconstruction of space from the decomposition may be excessively expensive. Also, the resulting process of such an approach does not reflect the temporal structure of the corresponding cognitive process.
We therefore try to better understand spatial structures and the power they may give to cognitive processes. In doing this, we want to exploit the inherent constraints of time and space.

2.2 Conceptions of time and space

When we say 'time' or 'space', the features we want to talk about include
– time: duration, simultaneity, speed, acceleration, precedence, concurrency, consequence,
– space: location, orientation, shape, size (height, width, length and their combination), distance, vicinity, neighbourhood.
But how do we think and talk about them?
As the preceding section may have indicated, there are a variety of ways in which time and space can be conceptualised. Each of these rests on implicit or explicit assumptions which may be more or less adequate for a given task.
Let us start with a common sense picture, which could be something like: time is 'an ever growing arrow along which changes take place'; space is 'a collection of entities (places) which stand in unchanging spatial relations to one another and which may or may not have objects located at them'. Implicit in these are the assumptions that the time arrow grows even when no (other) changes take place, that space is there even without objects to fill it, and that spatial relations and changes are observable. For some tasks, however,

it may seem more reasonable to assume that only the events (objects) that fill it *constitute* time (space).

Another distinction concerns the question whether points should be used as primitives or intervals (for time) or regions (for space). If we talk about Rocquencourt being south-west of Hamburg, we are most likely thinking of two points on a map of Europe. If, in a different situation, we say that you have to follow road X through Rocquencourt to reach a particular destination, Rocquencourt has to have an extension.

Also, it is not a priori clear whether a discrete, a dense or a continuous representation of time and space may be more adequate for a given task: If we want to be able to reason about arbitrarily small changes, a dense representation seems a better choice; if we want to say that two objects touch each other, we do not want anything to get in between them, so a discrete representation seems preferable. If on one level of consideration a touching relation and on another leve arbitrarily small changes seem to be appropriate, yet another structure may be required. Lastly, a continuous representation (e.g. R^2) is often implied – for a better correspondence with models from Physics?!

There seems to be mainly one point on which agreement can be found: time is generally conceived of as directed (and therefore irreversible), whereas space is not.

2.3 Abstract or concrete space?

Two disciplines which have dealt extensively with time and space are Mathematics and Physics.[3]

Mathematics defines an abstract space as a set of points, which are extensionless. It is a structure specified by a set of axioms which have to be satisfied by the points. Euclidean geometry builds a system of concepts on the basics of points, lines and planes : distance, area, volume and angle. More abstract than this metrical view are topological concepts [Klein 1939]. Its elegance and conceptual simplicity have made these abstract views of space very influential in Spatial Reasoning (cf. section 3.2).

Physics talks about a concrete space spanned by three positive orthogonal axes. Usually, one of these axes is the gravitational vector. Consequently, this model of space is local (the gravitational vector's orientation changes with a change in global spatial position). Physical space is always positively extended, and one can move in each spatial direction. Physical space and its elements are related to other physical quantities in many ways: movement relates time and (spatial) distance, molecular structure relates mass and spatial extension, gravity relates weight and mass, etc. These relations are particularly important for metaphorical concepts of space. Moreover, the 'locality' and the extensive use of reference frames (which are usually aligned with gravitation) are predominant concepts in many perception-related treatments of space. [Taylor & Tversky 1992, Paillard 1991]

3 Cognitive systems conceptualise time and space on different levels – dynamically and flexibly

In this section, we want to describe in more detail how different conceptualisations of time and space can be classified, and what advantages the *qualitative* approach has (sections 3.1, 3.2). We also want to describe *conceptual neighbourhoods*, a particular approach to the inherent imprecision of time and space in cognitive systems, which reflects the nature

[3] This section closely follows the discussion in [Freksa & Habel 1990].

Figure 2: *Tall* as a fuzzy quantity

of these notions transmitted by perception and is thus particularly adequate for cognitively oriented AI (section 3.3). Lastly, we want to argue the case for considering a problem at different granularities (which for the moment can be regarded as 'different levels of resolution', section 3.4), and integrate the treatment of neighbours in the horizontal dimension (conceptual neighbours) and neighbours in the vertical dimension (which are on different levels of granularity) (section 3.5).

3.1 Qualitative vs. quantitative

3.1.1 What do we mean by 'qualitative' and 'quantitative'?

The usual view of time and space is what we would call 'quantitative': "The talk started at 15 hours, 20 minutes and 30 seconds", "Claudia is 180 cm tall". However, if, in a cognitively motivated approach, we ask how people (or other cognitive systems) know about space, we will see that it is primarily through perception and only secondarily through abstract notions. This implies that reference values are taken entirely from *within the domain*; no external scale is imposed on the domain. Knowledge about time and space is then qualitative ("the talk started right after the coffee break", "Claudia is taller than David") rather than quantitative. This requires that knowledge acquisition be qualitative or that quantitative information be preprocessed.
Now 'qualitative' is often thought of merely as 'less precise'. However, this is not necessarily the case: If Claudia is – at a higher level of quantitative resolution – in fact 180.3 cm tall, while David is 179.8 cm tall, quantitative statements at 'cm'
resolution will yield values that are equal, while a direct comparison (see above) will reveal that she is taller. So in order to compare her to David using a quantitative scale, we would have to choose a 'mm' resolution – but this implies wasted computational resources if we want to compare her to Danny, who we know is something like 150 cm tall.
Language shows particularly well that *'qualitative' and 'quantitative' denote two different views* of – potentially the same – descriptors. As an example, consider 'tall' again. This descriptor may represent a fuzzy quantity (fig. 2) or a crisp quality (*tall* vs. *short* vs. *very tall*).
The emphasis we put on qualitative representations reflects our earlier stress on the dependence of knowledge (and reasoning) on the situation: Quantitative representations refer to external scales, which are situation-invariant. Qualitative representations refer to internal distinctions, which are context-sensitive:

- A rich domain enables a rich description; a poor domain implies a simple description.
- Context is required for generating and interpreting descriptions.
- The domain is represented using an inhomogeneous resolution: high resolution at the quality boundaries, low resolution elsewhere.

Figure 3: Qualitative relations in concrete free space

What are the relations we can distinguish qualitatively? Fig. 3 shows the relations that can be distinguished in 1D (topological) space. For time, which is directed, the mirror images of the first 5 relations have to be added (e.g. 'after" is the mirror image of "before") to yield altogether 13 relations. These relations were popularised in this context by [Allen 1983], who intended them to be used for a description of relations between time intervals, and have since been used in various formalisms to deal with time and space (for an overview, see [Freksa & Röhrig 1993]).

3.1.2 Why is it important to represent space as an integrated structure?

Transfer of this scheme to space, however, has not been straightforward: We usually interact with higher than one-dimensional spaces. Initial approaches simply decomposed n-dimensional space into projections to n one-dimensional spaces. For a complexity $O(c)$ of the one-dimensional case, this gives rise to representational and computational complexity $O(c^n)$. Also, the inherent spatial interrelation between the spatial dimensions disappears. Also, cognitive space is of lower dimensionality than Cartesian space. Subsequent approaches have tried to take spatial invariances into account. They have chosen different decompositions, e.g. a distance/orientation-decomposition [Zimmermann & Freksa 1995]. In these approaches, the number of dimensions determines the number of reference objects. This avoids a drastic increase in complexity, but a shape problem persists. Therefore, it has become clear that global and local spatial information have to be separated (i.e. the location and shape of objects). For a given task, the cognitive system has to *abstract* from certain spatial properties.

There is a strong interrelation between spatial operations and spatial dimensions: If the location of an object varies by a small amount, the variation of location is small *for each individual spatial dimension*. A given variation in location is reflected by an accumulation of variations in individual coordinates, i.e., the spatial dimensions are strongly interdependent with respect to a given relation in space. The main motivation for representing space as a integrated structure is for directly making use of the built-in constraints rather than composing substructures and adding constraints.

3.1.3 Abstraction in space

But despite the desideratum to maintain an integrated spatial structure, it is desirable to abstract from spatial aspects that do not affect the solution of a given task or from aspects that can be handled independently.

There is substantial evidence that biological systems split up various aspects of space into different 'maps' for independent treatment and later integrate the partial solutions.

Figure 4: Abstraction from specific quantities

If splitting up into Cartesian dimensions is not the solution, what kind of abstractions can we use?
Abstraction from specific quantities to yield qualitative relations appears to be one of the first knowledge condensation steps in natural cognitive systems. As an example, consider fig. 4: Everyone will immediately recognise that both pictures 'show the same thing'. (If the pictures are given at different times, people might not even notice the difference.)
Conceivably, we can abstract from every single spatial feature mentioned in section 2.2:

- Abstraction from location yields shape information.
- Abstraction from orientation yields shape information (provided there is a way of describing shape in an orientation-invariant manner).
- Abstraction from shape yields a representation of location and possibly approximate size or maximal size (provided there is a way of describing location in the absence of shape).
- Abstraction from distance yields a representation of direction/orientation.

It is worthwhile to reiterate here the central advantage of qualitative reasoning: a question like "Are A and B equal?" as such is *meaningless*. It only acquires meaning if it is seen in relation to a task or problem. This relation defines the criteria necessary for the decision. Consequently, only necessary decisions are made.

3.2 Geometry: from topological to metrical relations

Mathematics defines one particular, well-studied system of progressive abstraction from aspects of space [Klein 1939]. If, as is commonly assumed, geometry is the underlying model of 'real spaces' and also of 'mental spaces', we can order different levels of abstraction. Full metrical specifications are the most restrictive. Its properties are invariant with respect to a certain set of transformations: if we rotate or translate a spatial arrangement, its metrical description will remain the same. If we abstract from these properties, the resulting properties of the space are invariant with respect to a larger set of transformations. Continuing this relaxation process, we finally arrive at topology: now we can stretch and squeeze the spatial arrangement whichever way we want[4], its topological relations will remain the same.
In Psychology, the view that spatial cognitive development starts with topology and gradually progresses towards full metrical knowledge has been widely accepted since [Piaget & Inhelder 1956]'s

[4]Of course there are limits here too, but for practical purposes, they are not of too much interest.

topological analysis of the development of spatial representation.[5] This 'sequence' is also assumed to hold in adults when they learn new environments. Later studies [Lynch 1960, Siegel & White 1975] have elaborated on this view, defining the elements of more and more refined knowledge: a transition from knowledge of single *landmarks* (outstanding objects or places) via knowledge of *routes* between landmarks (specifying the topological relation of 'connectedness') to map-like *'survey'* knowledge (more or less distorted metrical relations).

Some approaches in AI have explicitly modelled a corresponding hierarchy of kinds of spatial knowledge, using 'route' and 'survey' knowledge (where the transition between the two forms of knowledge poses the same problems as in Psychology, however) [Kuipers & Levitt 1988].

Most approaches focus on a subset of spatial aspects. In the domain of Qualitative Spatial Reasoning (for a review, see [Freksa & Röhrig 1993]), a lot of effort has been spent on trying to find out 'how far one can get' with concepts that are as simple as possible. A few examples shall serve as an illustration:

- Randell,Cohn and Cui [Randell *et al.* 1992] focus on topology and derive all their relations from a single primitive 'is connected to'. They are able to model processes like phagocytosis and exocytosis – a unicellular organism's way of surrounding, engulfing and then digesting food particles and the expulsion of waste material from cells initiated by white blood cells, respectively. (It is obvious that for questions like these, aspects like *inclusion* are important, while direction and distance can usually be ignored.)

- Schlieder started from the question of how to describe a view (like a view from a mountain top, a tower, etc.). Obviously, the *arrangement* of the landmarks under consideration plays a decisive role here – 'left' or 'right' makes a difference. He is exploring how much expressive power is added to topology by knowing the arrangement ($\approx$ the mathematical 'sense') of landmarks [Schlieder 1993].

- Freksa and Zimmermann [Freksa 1992b, Freksa & Zimmermann 1992, Zimmermann & Freksa 1995] consider questions of route-planning, including finding shortcuts, finding home, etc. Here, *orientation* comes into play: 'Where is the church with respect to the filling-station?'.

These examples should illustrate our central point again: for different questions, different aspects of space are required (i.e. different aspects of space can be abstracted from). 'Intelligence' and 'knowledge' are, again, all about finding the right level of abstraction in a given situation.

3.3 Conceptual neighbourhoods

If space is important in a given task, it is because things that are 'close' have more to do with each other than things that are 'far'. Space is an organisation principle. So if we are provided with this organisation, we can employ processes more effectively.

As an example, consider planning a train journey. We can do this with the help of train timetables, which represent routes one-dimensionally. But it is far easier to do it with the

[5] It should be noted, however, that their notion of 'topology' is not exactly the same as that of Mathematics [Mandler 1988].

help of a map, which represents the same routes two-dimensionally. Now it is obvious that the two representations are equivalent (in [Palmer 78]'s sense). The difference is that the 2D representation is 'intrinsic' in the sense that it represents the constraints of the represented space implicitly. In the 1D representation, we have to note them and/or think of them explicitly: If our destination is north of our starting point, we have to know that considering a journey via a city south of it will (usually) not make much sense. And we as cognitive systems obviously have processes which can use this intrinsic 2D representation (cf. research on imagery).
Neighbourhoods reflect this organisation:

1. If an object occupies a given region, neighbouring regions are prime candidates for being occupied by the same object.

2. If a region is occupied by a given object in a given state of affairs, only neighbouring regions are candidates for becoming occupied by that object in the next state after a (small) movement of that object.

3. When an error occurs in locating objects at a given place, neighbouring places are the prime candidates for carrying the object.

For temporal events, 1., 2. and 3. can be applied accordingly.
Now this idea can be transferred to the domain of relations: certain spatial/temporal relations are conceptually closer to one another than others. Relations which differ by a minimal difference are conceptual neighbours.[6]
The special properties of conceptual neighbours are much like the ones of spatial neighbours.
Conceptual neighbourhood allows us to *order* concepts of time and space. This will prove particularly useful for a flexible use of concepts, as illustrated in the next section.

3.4 Granularity

There is another feature of descriptions, which stands orthogonal to the ones mentioned in the preceding sections. This is granularity, the 'resolution' of the description [Hobbs 1985]. With different levels of granularity, we associate different properties. An example are maps. We would not expect to find city streets on a road map, whereas with a tourist guide map, we would even be disappointed if it didn't specify the names of the buildings along these streets.
But this is not just a matter of 'looking at the map through a magnifying lens': contrary to common belief, spatial relations typically *are* affected by granularity transformations. We must therefore model the effects of switching granularity levels, e.g. angles, distances, shapes and neighbourhood relations. For an illustration, see figs. 5 and 6.

3.5 The horizontal and vertical dimensions

The preceding sections have shown what levels of conceptualisation cognitive systems can move between in order to achieve a dynamic and flexible adaptation to different

[6]Conceptual neighbours were introduced by [Freksa 1992a], who defined them as: "Two relations between pairs of events are *(conceptual) neighbors*, if they can be directly transformed into one another by continuously deforming (i.e. shortening, lengthening, moving) the events (in a topological sense)." ... "A set of relations between pairs of events forms a *(conceptual) neighborhood* if its elements are path-connected through 'conceptual neighbor' relations."

Figure 5: Non-monotonicity of intersection angles at different levels of granularity

Figure 6: The length of a coastline depends on the measuring unit – non-asymptotically

problem-solving situations. Drawing these thoughts together allows a new look at concept hierarchies.

Consider fig. 7. It depicts a hierarchy of different concepts of size (human 'height'). In this hierarchy, we can distinguish two dimensions of neighbourhood between concepts: horizontal and vertical neighbourhood. 'Horizontal neighbourhood' refers to competing concepts on the same level of granularity [Freksa & López de Mántaras 1982], while 'vertical neighbourhood' refers to compatible concepts on different levels of granularity [Zadeh 1978, Hobbs 1985].

How do people use these neighbourhoods?[7]

1. The speaker may be sure which descriptor he wants to use. Assume he wants to say that "Claudia is tall".

   If he is asked to specify what he means, he could either say

   (a) "Claudia is tall compared to other women." This statement would define an anchor-point on the horizontal level – to indicate which concrete objects the size scale is taylored to.

   (b) "Claudia is tall, but not very tall." This statement would define an anchor-point on the vertical level – to indicate which other concepts *this* "tall" is differentiated from; the level of granularity is fixed.

2. The speaker may be uncertain which descriptor to use. Assume he thinks Claudia is of normal height, but he is not sure he can really judge it.

   (a) Is Claudia rather tall or rather small? He is not sure. With this level of information, it is safe to say that she is *not* huge (and not tiny either): only conceptual neighbours are candidates for an adequate description of Claudia.

[7] For more discussion of neighbourhoods, see [Freksa & Barkowsky 1995].

Figure 7: Concepts of size

   It is safe to ‘compromise’ and to say that she is “of normal height”: uncertainty is resolved by moving to the vertical neighbour at the higher level of granularity. (But one cannot move to the non-neighbour small.)

   (b) He thinks that Claudia is of normal height, but that she is on the border to being tall.

   He should then use a concept like “rather tall”: need for a more detailed description of borderline cases is fulfilled by moving to the vertical neighbour at the lower level of granularity.

3. These considerations concerning the speaker also help the listener: If she hears that “Claudia is tall”, she may not know exactly what the speaker thinks of, but (provided she has an idea of the order of magnitude of female height)

   (a) know that Claudia cannot be small (horizontal neighbourhood),

   (b) no matter at what granularity level she thinks this “tall” resides, she cannot be far off the mark, since all the “tall”s are vertical neighbours.

It should have become clear that these hierarchies – which are not trees! – allow a very flexible and efficient choice of descriptor that fits a given situation best. This stands in stark contrast to an ever-constant level of resolution with a well-defined quantitative meaning.

4 How do these theoretical designs respond to computational / AI desiderata?

It is no coincidence that a computational realisation of these theoretical positions is a step towards fulfilling many of the desiderata that have emerged in modern AI:

- Greater economy of representation and processing is achieved by using qualitativeness.

This definition of *'efficiency'* is of course one of the central aims of Computer Science as a whole, and in this respect at least, AI is (or should be) a true subdiscipline of Computer Science.

- Systems are robust. On the one hand, they can take advantage of *incremental knowledge* (this might, for example, bring them onto a finer level of granularity, which allows a more precise answer to a query – but there had been an answer before too!). On the other hand, they show *graceful degradation*: If there is less information, or less precise information, they will move to a coarser level of granularity and still be able to provide an answer.

  This responds to one of AI's oldest problems: 'graceful degradation' has been one of the buzzwords of AI ever since its inception, but as it turned out, classical AI systems are not capable of it: They either know something, or they don't. As illustrated in section 1.5, this is contrary to our definition of 'cognitive systems'. Our approach represents one of today's ways of responding to this challenge.

- Because there will be an answer at every 'information state', an architecture built on our principles will also be capable of *anytime computing* – it will at any stage of its computation be able to give an answer 'based on what is known so far'.

5 Conclusion: how can time and space be dealt with in cognitively oriented AI implementations?

It seems reasonable to take stock at this point and to ask how the wealth of theoretical knowledge about time and space can be validated with the methods of Artificial Intelligence. We shall conclude this paper by mentioning predominant current approaches which explicitly take into account our main tenets:

- that intelligence can only be understood when thinking of cognitive systems in relation with their environment through their sensors and actuators,

- that cognitive systems employ knowledge at different levels, use these different 'pieces of knowledge' dynamically and flexibly, and that they can take advantage of neighbourhoods between concepts to find the knowledge fitting best to a given situation.

An important branch of modern Robotics is grounded in [Braitenberg 1984]'s thought experiments in 'synthetic Psychology'. Starting from the observation that 'analysis is difficult, but synthesis is easy', he describes 'creatures' made up of simple wiring between a small number of sensors and actuators. These could for example be light sensors and motors, mounted on a Lego robot. The creatures will respond to their environment and 'behave' in ways that look surprisingly complex *to us*. He shows how tempting it is to ascribe 'love', 'hatred', 'aggression' and a host of other dispositions and properties to them. It is the *interaction with their environment* that makes their behaviour looks so complex, that makes them seem so smart.
By now, the school variously called 'situated action', 'subsumption architecture', 'behavior-based robotics' and so on has actually built robots working along these principles. Some of these projects also deal explicitly with space [Mataric 1992].

However, using physical robots implies a lot of its own problems, many of whom are of the nasty engineering kind. Simulating 'creatures' in the computer allows more time to be spent on the questions one started out with – always at the risk of creating yet another 'blocksworld' whose results will not scale up to 'real' worlds.[8] Modern approaches, described as 'Artificial Life', usually put a lot of emphasis on biological principles, particularly ethology, ecology and evolution. See [Cliff 1994] about why its proponents think that the blocksworld risk is negligible.

For our own 'experiments in synthetic Psychology', we use the *realator*. This agent also lives in the computer. Our main focus is on using *qualitative* reasoning, building only on the information available through its sensorical and motorical capabilities [Barkowsky *et al.* 1994]. This platform allows us to test many of the predictions made in this paper.

References

[Allen 1983] Allen, J.F., Maintaining knowledge about temporal intervals, *Communications of the ACM* **26**, 1983, 832–843

[Barkowsky *et al.* 1994] Thomas Barkowsky, Bettina Berendt, Steffen Egner, Christian Freksa, Thiemo Krink, Ralf Röhrig and Antje Wulf, The Realator – How to Construct Reality –, *11th European Conference on Artificial Intelligence – Workshop on Spatial and Temporal Reasoning*, 1994, 19–26

[Braitenberg 1984] Braitenberg, Valentin, *Vehicles: Experiments in Synthetic Psychology*, Cambridge / Mass.: MIT Press 1984

[Brooks 1991] Brooks, Rodney A., Intelligence without representation, *Artificial Intelligence* **47**, 1991, 139–159

[Cliff 1994] Cliff, Dave, AI and A-Life: Never Mind the Blocksworld, *Proceedings of the 11th European Conference on Artificial Intelligence*, 1994, 799–804

[Freksa 1992a] Temporal reasoning based on semi-intervals, *Artificial Intelligence* **54**, 1992, 199–227

[Freksa 1992b] Freksa, Christian, Using Orientation Information for Qualitative Spatial Reasoning, *Proceedings of the International Conference GIS - ¿From Space to Territory: Theories and Methods of Spatio-Temporal Reasoning, LNCS*, 1992

[Freksa & Zimmermann 1992] Freksa, Christian & Kai Zimmermann, On the Utilization of Spatial Structures for Cognitively Plausible and Efficient Reasoning, *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*, 1992

[Freksa & López de Mántaras 1982] Freksa, Christian and López de Mántaras, R., An adaptive computer system for linguistic categorization of "soft" observations in expert systems and in the social sciences, *Proceedings of the* 2^{nd} *World Conference on Mathematics at the Service of Man*, 1982

[Freksa & Habel 1990] Freksa, Christian and Habel, Christopher, Warum interessiert sich die Kognitionsforschung für die Darstellung räumlichen Wissens?, in Freksa, Christian and Habel, Christopher (eds.), *Repräsentation und Verarbeitung räumlichen Wissens*, Berlin etc.: Springer 1990

[Freksa & Röhrig 1993] Freksa, Christian & Ralf Röhrig, Dimensions of qualitative spatial reasoning, *Proc. QUARDET* 1993

[Freksa & Barkowsky 1995] Freksa, Christian and Barkowsky, Thomas, On the Relation between Spatial Concepts and Geographic Objects, to appear in Burrough, P. and Frank, A., *Geographic Objects with Undetermined Boundaries*, London: Taylor & Francis 1995

[8] The derisive term 'blocksworld' derives from a – particularly 1970s – fashion of building simple worlds which could be fully understood and described, even by the imaginary robot. The first such program actually did deal with blocks. After much initial enthusiasm, it was understood that the approach, which presupposed full understanding, would not work for physical robots.

[Gallistel 1990] Gallistel, C.R., *The Organization of Learning*, Cambridge/MA: MIT Press 1990

[Hobbs 1985] Hobbs, J.R., Granularity, *Proceedings of the Ninth International Joint Conference on Artificial Intelligence* 1985, 432–435

[Klein 1939] Klein, F., *Elementary mathematics from an advanced standpoint. II. Geometry*, New York: MacMillan 1939

[Kuipers & Levitt 1988] Kuipers, Benjamin J. and Levitt, Tod S., Navigation and Mapping in Large-Scale Space, *AI Magazine*, Summer 1988, 25–43

[Lakoff 1987] Lakoff, George, *Women, Fire, and Dangerous Things*, Chicago: University of Chicago Press 1987

[Landau & Jackendoff 1993] Landau, Barbara and Jackendoff, Ray, "What" and "where" in spatial language and spatial cognition, *Behavioral and Brain Sciences* **16**, 1993, 217–265

[Lynch 1960] Lynch, K., *The Image of the City*, Cambridge/MA: MIT Press 1960

[Mandler 1988] Mandler, Jean M., The development of spatial cognition: on topological and euclidean representation, in Stiles-Davis, Joan, Kritchevsky, Mark & Bellugi, Ursual (eds.), *Spatial Cognition: Brain Bases and Development*, Hillsdale/New Jersey: Lawrence Erlbaum 1988, 423–432

[Mataric 1992] Mataric, Maja J., Integration of Representation Into Goal-Driven Behavior-Based Robots, *IEEE Transactions on Robotics and Automation* **8**, 1992, 304–312

[Paillard 1991] Paillard, Jacques (ed.), *Brain and Space*, Oxford: Oxford University Press 1991

[Palmer 78] Palmer, Stephen E., Fundamental aspects of cognitive representation, in Rosch, Eleanor and Lloyd, B.B. (eds.), *Cognition and categorization*, Hillsdale / New Jersey: Lawrence Earlbaum 1978, 259–303

[Piaget & Inhelder 1956] Piaget, Jean & Inhelder, Bärbel, *The child's conception of space*, London: Routledge & Kegan Paul 1956

[Randell *et al.* 1992] Randell, D., Cohn, A.G. & Cui, Z., Naive Topology: modeling the force pump, in Faltings, B. & Struss, P., *Recent Advances in Qualitative Reasoning*, Cambridge/MA: MIT Press 1992

[Retz-Schmidt 1988] Retz-Schmidt, Gudula, Various Views on Spatial Prepositions, *AI Magazine*, Summer 1988, 95–105

[Schlieder 1993] Schlieder, Christoph, Representing visible locations for qualitative navigation, in Piera Carrete, N. & Singh, M. (eds.), *Qualitative Reasoning and Decision Technologies*, Barcelona: CIMNE 1993, 523–532

[Searle 80] Searle, John R., Minds, Brains, and Programs, *The Behavioral and Brain Sciences* **3**, 1980: 417–424

[Siegel & White 1975] Siegel, A.W. and White, S.H., The development of spatial representations of large-scale environments, in Reese, H.W. (ed.), *Advances in child development and behavior*, Vol. 10, New York: Academic Press 1975

[Taylor & Tversky 1992] Taylor, Holly A. and Tversky, Barbara, Spatial Mental Models Derived from Survey and Route Descriptions, *Journal of Memory and Language* **31**, 1992, 261–292

[Turing 50] Turing, Alan M., Computing Machinery and Intelligence, *Mind* **LIX**, no. 2236, 1950: 433–460

[Ungerleider & Mishkin 1982] Ungerleider, L.G. and Mishkin, M., Two cortical visual systems, in Ingle, D.J., Goodale, M.A. and Mansfield, R.J.W. (eds.), *Analysis of Visual Behavior*, Cambridge/MA: MIT Press 1982

[Zadeh 1978] Zadeh, L.A., Fuzzy sets and information granularity, in Gupta, M.M., Ragade, R.K. and Yager, R.R. (eds.), *Advances in Fuzzy Set Theory and Applications*, Amsterdam: North-Holland 1979

[Zimmermann & Freksa 1995] Zimmermann, Kai und Freksa, Christian, Qualitative Spatial Reasoning Using Orientation, Distance, and Path Knowledge, to appear in *Journal of Applied Intelligence* 1995, also available as Report No. 42, Graduiertenkolleg Kognitionswissenschaft, Universität Hamburg, November 1994