The architectures of seeing and going:

Or, are cities shaped by bodies or minds? And is there a syntax of spatial cognition?

06

Bill Hillier
University College London, UK

Abstract

In my first paper to this Symposium, it was argued that the human cognitive subject played a key part the shaping and working of the city. The key mechanism was the *synchronisation* of diachronically experienced (and usually diachronically created) information into higher order pictures of spatial relations, the guiding form for which was an abstracted notion of a grid formed by linearised spaces. This notion was argued to be both perceptual and conceptual, serving at once as an abstracted representation of the space of the city and as a means of solving problems, such as navigational problems. In this paper, the question addressed is where the notion of the ideal grid comes from, why it has the properties it does, and what it has to do with the real grids of cities, which are commonly of the 'deformed' or 'interrupted' rather than 'ideal' kinds (Hillier, 1996). The answer, it is proposed, lies in the very nature of complex spaces, defining these as spaces in which objects are placed so as to partially block seeing and going, and, in particular, in certain divergences in the logics of metric and visual accessibility in such spaces. The real grid, deformed or interrupted, is, it is argued the practical resolution of these divergent logics, and the ideal grid its abstract resolution. In both resolutions, however, the resolution is more on the terms of the visual than the metric, suggesting that cognitive factors are more powerful than metric factors in shaping the space of the city. The question is than raised: do people have or acquire the concept of the grid, perhaps as some kind of perceptual-conceptual invariance of spatial experience in complex spaces, and do they use it as a model to interact with complex spatial patterns of the urban kind? This possibility is examined against the background of current opinion in the cognitive sciences.

Keywords
Metric integration, visual integration, Space Syntax, perceptual, conceptual

b.hillier@ucl.ac.uk

Introduction and problem definition

In my first paper to this Symposium, it was argued that the human cognitive subject played a key part the shaping and working of the city. The key mechanism was the *synchronisation* of diachronically experienced (and usually diachronically created) information into higher order pictures of spatial relations, the guiding form for which was an abstracted notion of a grid formed by linearised spaces. This notion was argued to be both perceptual and conceptual, serving at once as an abstracted representation of the space of the city and as a means of solving problems, such as navigational problems.

In this paper, the question addressed is where the notion of the ideal grid comes from, why it has the properties it does, and what it has to do with the real grids of cities, which are commonly of the 'deformed' or 'interrupted' rather than 'ideal' kinds (Hillier, 1996). The answer, it is proposed, lies in the very nature of complex spaces, defining these as spaces in which objects are placed so as to partially block seeing and going, and, in particular, in certain divergences in the logics of metric and visual accessibility in such spaces. The real grid, deformed or interrupted, is, it is argued the practical resolution of these divergent logics, and the ideal grid its abstract resolution.

In both resolutions, however, the resolution is more on the terms of the visual than the metric, suggesting that cognitive factors are more powerful than metric factors in shaping the space of the city. The question is than raised: do people have or acquire the concept of the grid, perhaps as some kind of perceptual-conceptual invariance of spatial experience in complex spaces, and do they use it as a model to interact with complex spatial patterns of the urban kind? This possibility is examined against the background of current opinion in the cognitive sciences.

Seeing and going

From an experiential point of view, cities seem to be about *both* seeing and going. Syntactic analysis confirms this by showing they are structured both to make the physical movement of bodies efficient, and to be intelligible to minds. We say that cities are about movement, but also that seeing is the mentor of movement. Concepts like syntactic 'intelligibility' (Hillier et al., 1987; Hillier, 1996a) build on this relation to construct a picture of the ease or difficulty with which we come to understand the shape of a complex space by seeing a part of it at a time through movement within it.

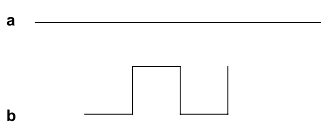

Figure 1

But on the face of it, seeing and going are very different concepts. Movement requires the expenditure of energy to move bodies, and so must be critically related to distance. Seeing is an informational concept, and involves cerebral rather the physical effort. A straight line, for example, has a certain length, and requires so much energy to move along it from one end to the other, but has only one visual field (at this stage, we ignore, conveniently, some might say, the fact that visibility decays with distance, but this is because as we move along the line this adjusts itself automatically and we continue to see the same field only more clearly) and so only one unit of information effort is required (Figure 1a). If we hold length constant and break the same length of line into segments and connect them at angles which prohibit continuous visibility (Figure 1b), we do not add significantly to the energy effort required to move along it, but we do add greatly to the informational effort required.

So defining a complex space as one in which one or more objects are located so that seeing and going are partially blocked, whether objects in a gallery, counters in a department store or islands of buildings in a city, we can index physical accessibility by *metric integration* (in fact, 'metric total depth' in Alasdair Turner's Depthmap software in a version which gives true metric distances), that is the real 'universal distance' from all points to all others, and therefore the energy effort that must be used to go from all points to others; and visual accessibility by *visual integration*, (currently labelled 'total depth' in Depthmap) in the same software, which indexes how many fields of view we have to take into account to see the whole system, and therefore the informational effort required to see all points from all others (Turner, 2001, 2002). Any object placed in a space will both add universal distance to the system and so require additional physical effort to go to all points from all others, that is decrease its metric integration, and at the same time will increase the number of views which must be taken into account to see the whole system from all points to all others, and thus add to the informational effort required, that is decrease its visual integration.

Metric and visual integration: a real case

We can begin by looking at a real case so see how metric and visual integration behave, and what kinds of pictures they give of the complex space formed by a real urban system. Figure 2 is a metric integration analysis of an arbitrary chunk of London centred on the City, and Figure 3 is a visual integration analysis of the same area. Metric integration forms a pattern of concentric rings, while visual integration finds an edge to centre network of lines. The reasons for this are clear enough. Metric integration is based on elements that are uniform and as small as possible (the smaller the more accurate the metric measure) and so must hug the geometric centre of the system. (Hillier, 2001, 2003) Visual integration is based on elements which are as large and as variable as the distribution of objects allows, and

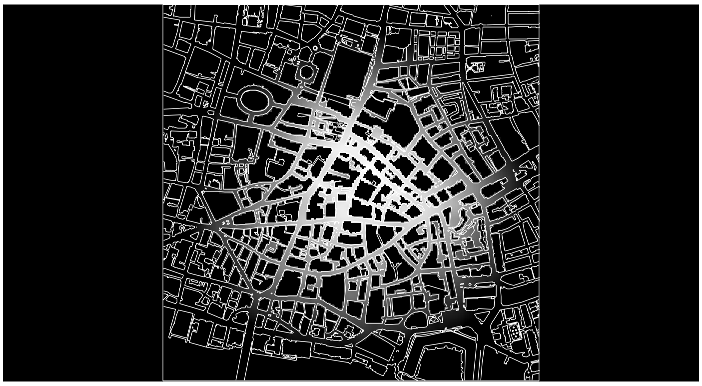

Figure 2

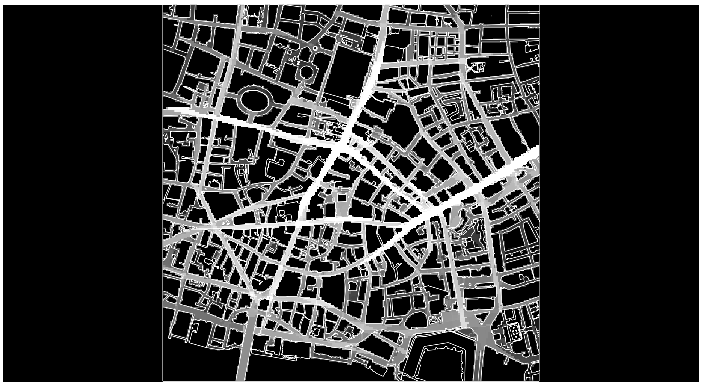

Figure 3

in this case, the disposition of objects in such as to creates a visual integration core in the form of the edge to centre network.

As representations of the spatial and functional properties of the city, it is clear that the visual analysis is much more persuasive. Although the two kinds of analysis find their centres in the same region of the grid, metric integration identifies as its focus a functionally unimportant intersection, and then identifies minor streets and even back alleys close to this centre as being more integrated than much more important streets remote from it. Visual integration, on the other hand, differentiates different types of street close to the focal point, and extend the integration core to cover important streets some way from the visual integration centre. This is essentially why a metric integration analysis cannot predict movement as well as a line analysis (Hillier et al., 1986). It overestimates syntactically insignificant lines close to the geometric centre, and under-estimates significant lines remote from it. It is too tied to the abstract overall geometry of the system, and too little affected by the spatial configuration.

Does this mean then that the visual prevails over the metric in the shaping of cities, and that metric factors can be discarded? We will in due course show that something like the first is the case, but that the second is not. Cities are shaped by a subtle and complex interplay between visual and metric factors, and in key senses the one will be shown to lead to the other. To show this we must first undertake some theoretical explorations into the formal behaviour of visual and metric integration in simple – defining these as convex spaces without objects – and complex spaces.

Metric and visual integration in simple spaces

So let us first look at simple convex spaces without objects. We can anticipate, of course, how metric integration will behave: it will form concentric rings from centre to edge. Figure 4 shows the effect on metric integration of varying a convex shape from square to increasingly elongated. As we expect (Hillier, 1996a), the square is the most metrically integrated shape overall, with an internal distribution from centre to edge and from centre edge to corner, and the most elongated shape the least integrated overall, but in each case integration goes from centre to edge and from centre edge to corner.

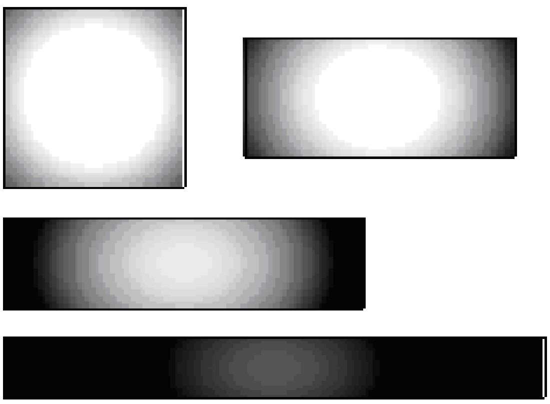

Figure 4

What about visual integration? There is of course a problem. All points can see all others so the four shapes seem to be undifferentiated, either in mean differences between them or differences in locations within the shape. But intuitively to see spaces easily we need to be at their edges or corners where the whole space can be within a single field of view. Is this intuition mistaken? It is not, and we can see this by looking at what happens visually as we move from the centre to a corner location, by using a crude device which I call a 'metric j-graph', much inferior to Alan Penn's geometrical j-graphs (Penn, 2003), but simpler to illustrate the point that needs to be made here. We cover a shape with a uniform tessellation, then take a root cell in the tessellation and align all cells face-wise contiguous with it one layer above, then those contiguous with the 'one deep' set on the next alignment above and so on until the limits of the system are reached. The result is a shape which looks like a j-graph without its links, and whose shapes indexes the 'manhattan', or 'taxicab' distance from that points to all others. The results would be similar with true distances but the 'metric j-graph' would be more time-consuming to construct.

In the centre of the space the visual field in 360 degrees, so the shape of the metric j-graph is as in Figure 5a. The four cells at the top are the corners, and those at the widest point those where the shape of the graph most closely approximates the circle. On the right is the equivalent metric j-graph for the corner location. The j-graph has become much deeper, reflecting the greater mean distance to be travelled to all other points from a corner location, but by the same token, the j-graph has also become much narrower, that is, it has moved into a more focused field of vision.

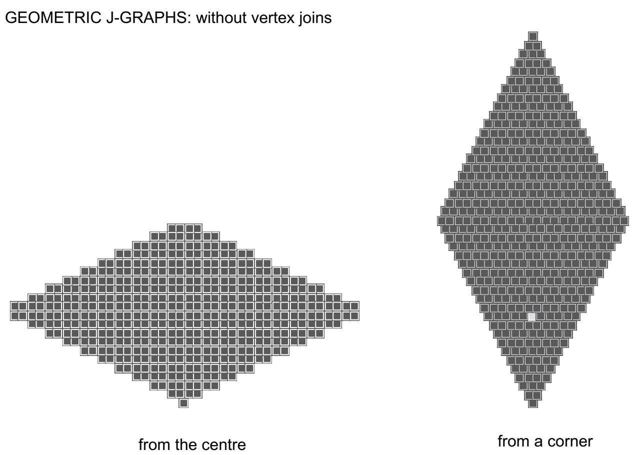

Figure 5

This will always necessarily be the case in a simple convex space. By moving an observer from the centre towards the edge, we are necessarily both increasing the distance from that point to all others, that is we are making the system more metrically segregated from that point, but at the same time we are *narrowing* the system by making the cells more asymmetric from the observer, and thus focusing the field of view and making it easier to see the whole system. Seen another way, more area is being brought into single point of view, though at the cost of making it farther away. This is the fundamental relation between seeing and going in a simple convex space. The more metrically segregated our position, that is, the more energy effort we need to go to the whole system, then the more visually integrated the system, that is the more we see at once and the less informational effort we need to see the whole system.

A key implication of this is that a notion of the directionality of vision is built into the fundamental relation between the metric and visual properties of spaces: the more segregated our locations, the more what we see is focused into a single view. Since we define visual integration in terms of the number of view we have to take into account to see the whole of a system, then it is clear that in this sense a corner location is more visually integrated than a central location where we must take more views into account. This makes linear spaces highly interesting, since the more elongated the shape the more the focusing of vision will be the case, and the more distance will be overcome by vision. This limit of this tendency is of course the line itself. We will explore the consequences of this in due course.

06.6

Metric and visual integration compared: some simple cases

But let us look first at what happens in complex spaces, beginning with the simplest kind, those with a single object. One effect, of course, is that as soon as we put an object in the space we can measure visual integration since the system is no longer uniform. How then do metric and visual integration relate to each other in such spaces? Figure 6 shows the effect on metric integration of moving an object from the corner first to centre edge and then to centre. The systems are processed together so the colours (or dark to light shading) represent the same numerical values in each case. We can easily see that as we move the object towards the centre the degree of metric integration in the system falls. Figure 7 then shows the effect on metric integration of changing the shape of a central object while maintaining the area of the object constant. As theory predicts (Hillier, 2001, 2003) the elongated object decreases metric integration much more than the square object. Note that the focus of metric integration with the elongated shape is split between the two ends of the shape.

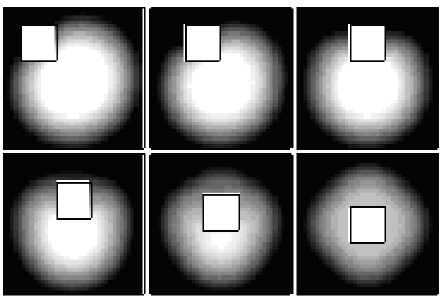

Figure 6

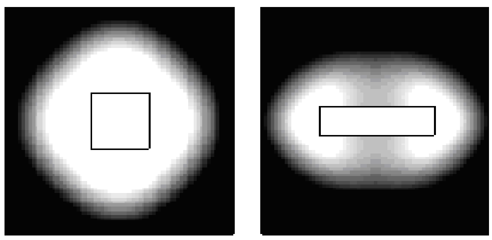

Figure 7

Figures 8 and 9 then show the effect on visual integration of the same series of changes in block position and shape. Once again, as we move the square block from corner to centre edge and then to centre the visual integration of the shape decreases, and when we change the shape of the central object from square to elongated holding area constant, visual integration is again decreased. However, in terms of the internal *distribution* of integration *within* the space, the two measures are almost inverses. While metric integration is maximal in the metric centre of the object, visual integration is highest on the periphery, and would be in the corners if the central object were circular or octagonal rather than square (the visual integration from the corner is reduced compared to the sides because the amount of blocking from the object is defined by the length of its diagonal rather than that of its sides). So, in this simple case, metric and visual integration are *directly* proportional to each other in terms of their *degree*, but *inversely* related in terms of their *distribution* within the space. Figure 10 summarises these effects.

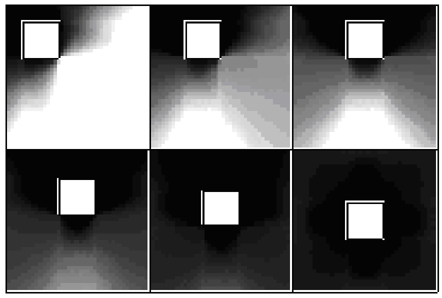

Figure 8

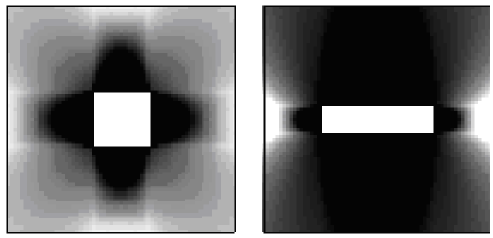

Figure 9

We can see then that the visual integration of a point is affected by two things: where the object is in the space and where the point – or observer perhaps – is in relation to the object. This is clarified in Figure 11. On the top line, we locate the observer in the corner, and move the object from centre to corner. The black area of unseen points shrinks as the object moves away, and the area of seen points increases proportionately, with the boundary between the two given by a line drawn from the root point. On the bottom line, we move the observer from the vertex of the object to the corner of the space. Once again, the area of seen points grows and that of the unseen points shrinks. The common denominator is the distance of the observer from the object. It is this that determines the ratio between the seen and unseen areas, and this determines visual integration.

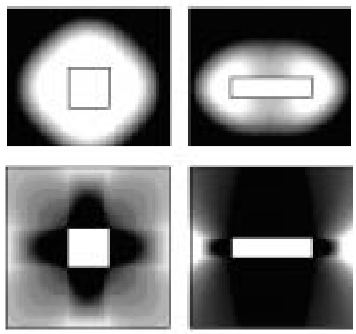

Figure 10

The priority of the edge for visual integration is then in this case not the same as in the simple space, but depends on *metric* factors, namely the total number of points, or total *area*,

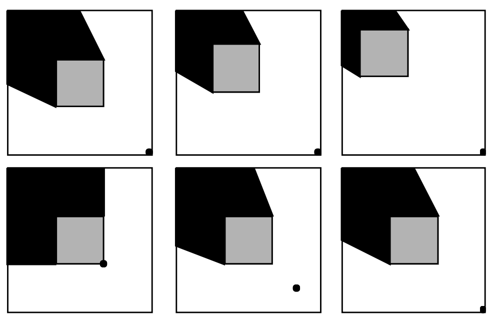

Figure 11
The top line shows that the propotion on unseen space shrinks as the object moves from centre to corner.
The bottom line shows that the propotion of unseen space shrinks as the observer moves from centre to cprner.
The key variable is the distance between observer and object.

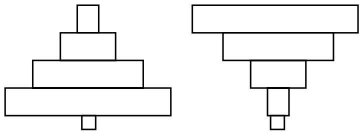

a visually integrated (left) and visually segregated (right) point showing total area at each level of depth from the point

Figure 12

at each level of depth from the root point. This has an important implication. While the visual *depth* between a visually integrated and visually segregated point will be the same in either direction, the visual integration values are different because the large and small areas have changed places depending which point we are looking from. If we represented total area as a strip of a certain length, we could again represent the visual integration value metrically by a technique analogous to j-graph. An integrated point would look like a pyramid, with a wide strip at the first level and narrower strips at succeeding levels, while a segregated point would look like and inverted pyramid, as in Figure 12

06.8

Looking closely at what happens inside visual integration

To the syntactic analyst, it is of interest that lines featured in both of these accounts of visibility, in the case of simple spaces as the limiting shape for directional visibility, and in complex spaces as the means by which the boundaries were drawn which yielded the visual integration value. However, apart from the obvious fact that for a visual relation to exist between two points, visual integration seems to be based on the analysis of the relations between point isovists, and to have little to do with lines. But looking more closely at what happens inside visual integration, this will turn out not to be the case. What visual integration does initially (before factoring for the metric areas at each level of depth) is to calculate the visual depth between points, and it does so by implicitly identifying virtual line like elements, which it then treats as discrete elements, discarding all intrinsic information whether metric, geometric or configurational. This is surprisingly similar of course to what axial analysis does, but here it is done between pairs of points, and implicitly as well as algorithmically. What follows has close analogies to Peponis et al (1998), though differences in the way it is put allow perhaps a wider theoretical embedding of the consequences of the ideas, so this particular re-invention of the wheel is set out in the form in which it was conceived, rather than trying to adapt itself to the excellent text of Peponis and his colleagues.

Formally speaking the visual integration calculation is made by filling the complex space with an arbitrarily fine square tessellation, then taking the centre of each cell in turn, a depth one connection exists if another point is in the isovist of the selected root point, and a connection at depth two wherever a point not in the first isovist from the root point is in the isovist of one of the points in the first isovist, and so on until all points in the system are covered. On the face of it, this process seems to have little to do with lines, apart from the fact that each connection between points implies a line. But there are a huge number of such implicit lines in the system, so to think of them seems a move in the direction of additional – and therefore unnecessary – complexity.

But this does not quite settle the matter. The initial value that is calculated in visual integration, and the one that is eventually factored by the metric area at each level, is the visual depth between any two pairs of points, that is the number of isovists that must be passed through to see one from the other. But what does this actually mean? We can see this graphically by looking at the complex space shown in Figure 13, where the point highlighted bottom left is our root, and the point top right the target point whose visual depth from the root we are calculating. In the figure we have calculated visual point depth from the root point, so the system is coloured from dark blue to red (or shaded dark to light) from the root outwards in a series of depth layers. The dark blue area around the root point is the depth one isovist, depth two is then the 'isovist of the isovist' of all the points in the depth one isovist, and so on. With this representation, it is easy to see how the sequence of isovists 'gets to' the destination point, and in this case it does so in two different ways in six steps, so the visual depth is 6.

In effect, the visual depth of any point to any other is established by finding routes from one to the other, and visually the 'route' is as clear as a j-graph. The route seems to be made by isovists, but it is not quite the case. If we look at the relation between any two levels in the isovists along the route - say between the first and second - the relation that is part of the route from root to destination is not made by the whole isovist but only by a subset of its points. This will always necessarily be the case. If we then consider the relations across three levels rather than two – say the first and third levels in Figure 13 - the relation between the first and third layers is made by a space defined between two subsets of points, one in the first layer and the other in the third layer. This section of the route from root to target will be, necessarily, a convex strip or wedge-shape connecting these two subsets of points, and the same will be the case for all relations between triples of levels. To the extent that the complex space is densely filled with objects, these linking strips will more and more approximate lines, and indeed the limiting case will be a sequence of lines.

Figure 13

These connecting elements are not of course real, but virtual, but the fact that they are connecting element means that they are then treated as discrete elements, in that all their intrinsic properties are discarded and only the fact of making the connections is taken into account. The fact that all information bar making the connection is discarded about these virtual elements implies that only the fact of making the simplest linear connection is taken into account. To all intents and purpose, the virtual elements are treated as virtual lines. This sounds, of course, analogous to axial analysis, which calculates only the extrinsic topological relations between line element, but on reflection we can see that technically speaking it is its dual. In axial analysis, the lines are the nodes and the intersection the links of the graph. Here the points are the elements and the virtual lines are the links. Indeed, it is precisely because the virtual lines are the links in the graph that they are treated as a discrete elements and their intrinsic properties ignored. But what visual integration is doing in effect is identifying the least sequences of virtual lines from root to target, and find these algorithmically, if implicitly, and then treating these lines as discrete units and calculates only their extrinsic properties[1]. The fact that the initial representation is the dual of an axial line system should not distract our attention from the similarity of what is going on in the two forms of analysis.

The paradox of the line

So visual integration analysis is in an important sense all about lines, and lines, moreover, treated extrinsically in terms of their topological relations, with no reference whatsoever to metric or geometric factors. As soon as this is clear, a possible relation to metric properties comes into view. A line is not only a direct visual connection between two points: it is also the shortest distance between them. Any break in a line or making two lines into one will cause metric and visual integration to co-vary positively in terms of overall degree. The co-variance will not be linear, since a marginal change in alignment may make a step change visually but a marginal change metrically. Even so, there will be change in both and it will be in the same direction. It follows that by discarding intrinsic metric information about lines and treating them as line topologies, visual integration comes to reflect at least some metric properties of the system**.**

We can clarify this through what we might call the paradox of the line. In the Euclidean world which we inhabit in everyday life, the most obvious thing that can be said about the line is that it is the shortest distance between two points. However, in syntax we know, and have just reminded ourselves of, something else about the line: considered configurationally in terms of the universal distance from all point to all others, that is, in terms of its *intrinsic* properties from all points to all other, the line has the *largest* sum of distances for any way of arranging that set of points in a shape. Seen in terms of its intrinsic configurational properties as a simple space, then, the line is the ‘longest’ not the shortest distance. It is *extrinsically* that it is the shortest distance. This is the paradox of the line. We can also say that metrically speaking, the centre of the line is the most integrated location and the ends of the lines the least, so the two points between which the line is the shortest distance are also intrinsically the farthest from all other points along the line. The paradox, it seems, has some twists and turns. Seen in terms of its *intrinsic* properties, the line is the inverse of its extrinsic properties.

How does this compare with the visual behaviour of the line? In terms of the inversion of visual and metric properties for simple spaces what we saw earlier, the line is most easily seen from its two ends. In this it is comparable to metric integration from its two ends, that is extrinsically. How then does it compare to metric integration *intrinsically* as a simple space in terms of the visual relations between all points and all others within it? The intrinsic metric properties were calculating the distance from all point to all others in a convex shape. We have not yet seen how this can be done for visual integration, though we did show the priority of edge and corner locations for directional vision. We need some more theory.

Intrinsic visual integration

We are familiar with the effect on metric integration of moving a cell placed between two end cells from the centre to periphery (Hillier, 2001, 2003). (Figure 14) Metric integration along the line increases (that is, the mean distance between pairs of points, or ‘universal distance’, decreases) as we move the blocking cell from centre to edge of the line between the two end cells. (Hillier, 2001, 2003) This followed from the ‘centrality principle’ for partitions set out in (Hillier, 1996a).

INTER-ACCESSIBILITY from all points to all others

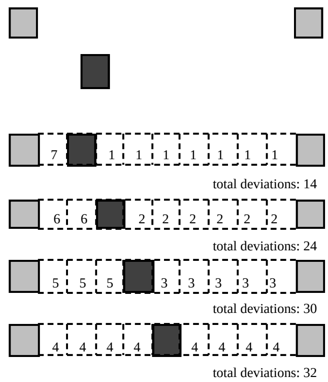

Figure 14

How then does visual integration behave? Figure 15 shows a simple linear space with a partition that is progressively moved from centre to edge. Regardless of where we put the partition, the total visible area in the two cells taken together remains of course the same. But if we take a ‘configurational’ point of view, and consider the area, or, equivalently, the number of other points,

MEAN AREA SEEN FROM ALL POINTS so mean number of other points visible from each point

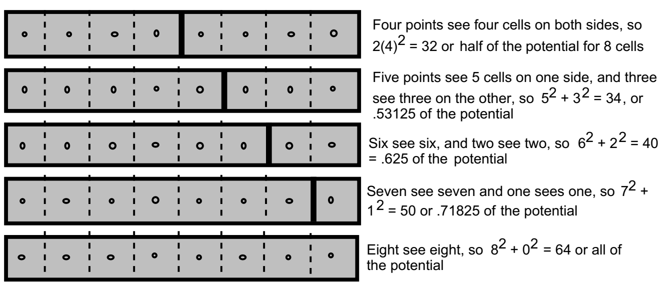

Figure 15

that can be seen from each point within each of the two cells, we find that the mean *increases* as we move the partition from centre to edge. The reasons for this are shown in the figure, and are quite similar to those for the effect on metric integration. If we have a line of n cells and define area as the number of cells that can be seen from a point (though it could equally be the points making up those cells), then with a centrally placed partition four cells can see four others on each side, so the total for all points is $2(n/2)^2 = 32$ and the mean is 4. If the partition is moved one cell along, then on one side five will see five and three will see three, so $5^2 + 3^2 = 34$, and so a mean of 4.25. Since $n^2 + n^2$ must always be a smaller number that $(n-1)^2 + (n+1)^2$, and $(n-(1+m)^2+(n+(1+m)^2$ must always be smaller than $(n-1)^2 + (n+1)^2$ it follows that the mean area, and therefore the mean number of points, seen from all points must increase as the partition moves from centre to edge.

06.12

The total co-visibility of points from others in the system is not then a constant like area, but a variable depending on where we place the partition. If we hold total area constant, there is more visibility from points to all others in a large space plus a small space than in two similar sized spaces of equal total area, and this will hold whether or not the spaces are connected to each other. This is not immediately obvious to reflection, but it is experientially and mathematically clear. We can say experientially, because it implies that people spread evenly throughout the space see more others in a large and small space, than in two evenly sized space. It follows that we can also say that if an individual who moves in the space will also see more points over time in a large and small space than in the same area divided equally. It is not just the *sense* that points are more co-present with each other in a large and small space than in evenly sized spaces. For mobile individuals or groups of individuals it is a mathematical *fact*.

Once we have the principle, we can see that this also applies to directional vision. If from a central space we have two (or four, it does not matter) fields of view of equal size, as we approach the corner we increasingly have one large and one small space, and this means that the potential co-presence of points – and therefore of other people - from a near corner location is greater than for the central location, and greater from a corner than from a near corner location, since we will see the whole space in one view. Intrinsic visual integration is then very much the inverse of metric.

We must be clear what is going on here, and how it differs from the visual integration we have seen do far. What we are doing here is summing visibilities from points, or sets of points, and so we are dealing with space as a set, but not as a pattern, of visual experiences available from particular points. In that sense the measure is perceptual in a way that the patterns brought to light by visual integration are not. It is about what it is like to be in a space or a set of spaces, but it does not give a picture of the relations between those spaces or their relation to the rest of the system which elude direction perception.

But having said this, we can apply this measure to movement and therefore to routes by seeing them as a sum - but not as a structure – of experiences. Given a journey of a certain length, we see more space on a journey composed of a long and short lines than of even length lines, and more on one made up of a single line than any made of a sequence of lines. We can show this through simple sequences of numbers representing line lengths. For example, suppose we have a journey of length 8 made up of two lines of 4, so in terms of what we see along the route we will have from successive points a view of 4-4-4-4-4-4-4-4 other points, giving a total of 32. If we have lines of 5 and 3 we have 5-5-5-5-5-3-3-3 = 34, so we actually see space more on the way through time, and so on for 6-6-6-6-6-6-2-2 = 40, 7-7-7-7-7-7-7-1=50 to 8-8-8-8-8-8-8-8 = 64.

If we limit what we see directionally, so that what we see is just ahead of us and so decreases as we move along each line, the outcome is similar. For two lines of length 4, we have 4-3-2-1-4-3-2-1=20 and for lines of length 5 and 3 we have 5-4-3-2-1-3-2-1=21 and so on through 6-5-4-3-2-1-2-1=24 and 7-6-5-4-3-2-1-1=29 to 8-7-6-5-4-3-2-1 =36. These calculations will of course be the same whichever direction we take. Of course it could be said that as we go along the line, we are receiving smaller and smaller amounts of the same information, whereas when we change lines we receive new information. This however suggests a natural way of link linking space to the mathematics of information theory (Shannon and Weaver, 1948): the length of lines controls the amount of information theoretic redundancy, or structure, in the system, and the number of changes the amount of information, or unpredictability. We will return to this theme in due course.

This then is a pervasive, though hidden, effect of all our experience of moving in space. We actually see more in moving through large and small spaces that equal sized spaces of the same total area, and it has different effects on our sense of how space is structured. We find similar effects with time. Suppose we have a morning during which we want to work on, say, a paper, but we have a couple of tutorials which must be fitted in somewhere. Where should we put the tutorials? Intuitively, it is quite likely that we would seek to make the working periods as long as possible so as to sustain concentration. We would do this by clustering the tutorials either at the beginning or the end of the morning. The opposite would be to time the tutorials at even intervals throughout the time period, since this would seem to create the greatest amount of interruption. But why exactly do we think this, when whatever we do the total time available for work and tutorials is the same. The answer lies in the same simple maths that increases 'co-presence' of points as we move the partition from centre to edge. When we say we want as long a period as possible of sustained effort, what we are formally saying is that we want to maximise the co-presence of as many points in time as possible within the same time frame, and this is what is achieved by clustering the tutorials at the beginning or end. In this case, the place of movement in space is taken by passage through time. Spacing the tutorials evenly throughout the period, so that each tutorial was in the centre of the time period of its two neighbours, would minimise the co-presence of points in time within the same time frame. There is, quite objectively, more time co-presence in a long and short time period that in two equal length periods of the same total duration, and this is how we experience it through the passing of time.

The intrinsic properties of lines

This means that the line has important intrinsic visual properties. In terms of all round visibility, of course, the line is uniform: each part can see each other part. But if we consider its intrinsic properties in terms of directional visibility – and above all else a line is directional - it is structured. On a pure line there are two visual fields to be taken into account from the centre (as opposed to the postulated four in a two dimensional strip), and one from each end, with point co-presence on the line increasing as the observer moves from centre to end. The line is therefore intrinsically more visually integrated, and requires less visual information, than any convex two-dimensional space, and the more a convex strip approximates a line, then the stronger this effect will be. Practically speaking, we could say that the longer and narrower the convex linear strip, the more the lateral views needed to see all round will be small spaces to set alongside the long linear spaces, and this will then maximise visual integration. In other words, the more a convex strip approximates the line, the more intrinsically visually integrated the shape in terms of its whole area as well as from its two ends. Up to certain practical limits imposed by the focusing of vision, narrowness as well as length increases visual integration[2].

We can say then that for metric integration the intrinsic behaviour of the line – having the greatest distance from all points to all others – is inversely related to its extrinsic behaviour – being the shortest distance between its two end points - while for visual integration the intrinsic and extrinsic are positively related: it has maximum intrinsic point co-presence along the line, which increases with increasing length, and maximum visibility from end to end, and therefore optimises extrinsic visibility, and again this increases with length. We can pretty well say theoretically that for intrinsic all to all (and directional) visual integration the longer the line the more integrated it is, while for all to all metric integration the less integrated it is. At the same time in terms of extrinsic properties metric and visual integration agree that the line is optimal.

Intersection

Both metric and visual integration, then, prioritise the line extrinsically, that is in terms of relations between discrete spatial elements. The longer the line, the more extrinsically efficient it will be both from a visual and metric point of view, but the more intrinsically inefficient it will be from a metric point of view. Intrinsically, longer lines increase visual integration but decrease metric integration, and visual integration integrates from edges while metric integrates from centres. How can these divergences be resolved?

Both, in fact, are resolved at once by the simple device of line intersection. As soon as a pair of lines intersect, intrinsic visual integration is refocused from the edge of the system to its centre, that is at the point of intersection, and this is also the point of highest metric integration in the combined system. Since the paradox of the line arises from the fact that metric integration decreases with increasing length of line, by cutting the line in two we in effect halve the effect, at least with respect to the combined system of two lines. Line intersection then the simple device which can be used to create and control the pattern the pattern of both metric and visual integration, and to do so in such a way as effect on one will tend to positively co-vary with effects on the other. Since any pattern of intersection will form the topology of the grid, and since any complex space that becomes densely filled with discrete objects will construct such a pattern, it follows that some kind of topological grid formed by more or less linearised spaces – though as yet without an angular structure - but see below – is inherent in the nature of complex spaces, and serves as a means of defining the pattern of metric and visual integration together. Since a fully linearised grid is, other things being equal, the form that optimises both metric and visual integration extrinsically (in that any breaking of lines will reduce both), then it follows that the linear grid arises conjointly from two properties of complex spaces that most closely interact with human behaviour: that is, their metric and visual properties.

The ‘other things being equal’ clause, is necessary because there are in fact two ways in which the linear grid can be improved on in terms of both metric and visual integration. One is by the creation of open spaces; and the other is by reducing the size of blocks. Each of these of course locally in real grids, the former by the creation of ‘squares’ and open spaces, the second is by local grid intensification, as found in local ‘live centres’ (Hillier, 2000). But these occur only as local features of grids. The global structure of the grid is given by the pattern of intersections between more or less linearised spaces. This is why it is possible to say that the grid in general is implied by the interaction of metric and visual integration in complex spaces, and that the structure of real grids is given by their the topological patterns formed by their intersection, and in particular, their degree of alignment or non-alignment into longer or shorter sequences, and so longer and shorter lines.

The perceptual and the conceptual

Having arrived at the grid, however, as a way of resolving relations between metric and visual integration by intersecting lines, we immediately find another issue: our two different versions of visual integration – the point based directional version and the all round, all-points-to-all-others version – tell us rather different stories about the grid. We can perhaps, an indicated earlier, think of the first as *perceptual*, since it refers to visual fields from points or sets of points seen additively rather than as a structure, which would only in limiting cases be available from a single point, and the second a *conceptual* since it refers not to the experience of visual fields from particular points, but to the structure of relations from all points to all others.

Let us begin with the perceptual version. If we think of an intersection as a point isovist, and analyse it for its intrinsic, or perceptual, visual integration, then point co-visibility is maximal where the isovist is maximally asymmetric, that is when one branch of the isovist is longer than the others, This implies that perceptually an intersection of a long and short line in a kind of L-shape where the point of intersection is close to the end of each line, is optimal. This seems initially alien to the concept of a grid, but in reality the relation between long and short line seems to be the fundamental building block of the city, whatever its geometric form. At whatever scale we look, we find that the city is made up of a few long lines and many short lines, and it is this relation which gives rise both to ease of direction finding in the city, and to the overall structure of the urban grid (Carvalho, 2003; Hillier, 2001, 2003). We can say then that the point based perceptual version of visual integration leads us what is empirically a pervasive spatial relation of the city: the intersection between a long and short line.

What then of all-round visual integration, that is, the measure that seems to give such a persuasive account of the overall structure of the grid, and does to in striking contrast to metric integration which identifies little more than the geometric shape of the overall system? First, we can say that because all round visual integration proceeds in all directions simultaneously by identifying virtual lines leading to all points, integration will be optimal where long lines move in all directions from a point, giving a preference for an X-shaped intersection. To the extent then that routes to all points are linear and intersect in the X-relation, integration will be increased.

This suggests that the two kinds of visual integration seem to relate to the two kinds of movement that are found in urban systems: (Hillier, 1999). Point to point movement, such as is found from edge to centre, is directional, and tend to be supported by the L-shaped intersection where one direction is prioritised, and all to all movement is omni-directional and tends to be supported by the X-shaped intersection. Looking at axial maps, we see a greater occurrence of the L-shaped relation in outer areas where the overall grid structure is less resolved, and movement is more point to point, and a greater occurrence of the X-relation in the centre and in local centres, where movement is from all points to all others.

From the perceptual to the conceptual

However, both versions of visual integration seem to play a role in the process of 'description retrieval', as outlined in my previous paper (Hillier, 2003b), by which individuals come to form a picture of a complex space by moving through its parts. We have already suggested that this is an information theoretic process in which line length controls the amount of information theoretic redundancy, or structure, in the system, and the number of changes the amount of information, or unpredictability. If this were the case, then there are a number of implications.

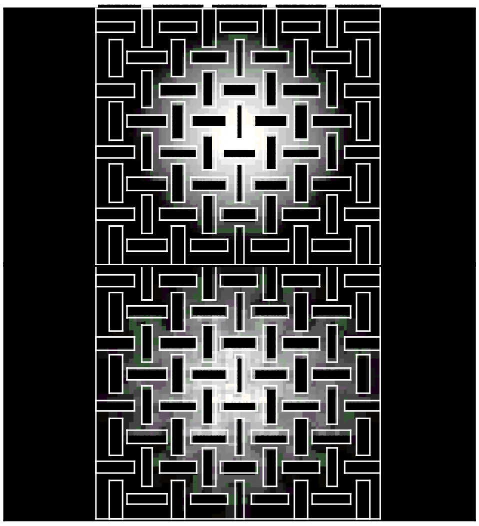

Figure 16

The first is that cities, like language, work as intelligible systems because there is a balance between structure and unpredictability. In the case of cities, the place of linear structure in language is taken by movement through space, which is similarly linear through time. This would suggest that the deformed and interrupted grids that are found in real cities are language-like in this respect: they balance the structure obtained from longer lines with the information from changes in direction. This seems to be supported by our notion of 'labyrinthine' space. Labyrinthine tends to mean that lines are uniform and short, and lack the longer lines to structure the system. It is of interest that we can create a labyrinthine complex space by simply arranging blocks so as to reducing the length of lines and make them as uniform as possible. This will have the effect of compressing visual integration down to the level of metric. For example, in the discrete T-shaped pattern of block aggregation shown in Figure 16, the visual integration pattern in the bottom figure is virtually identical to the metric in the top figure.

The second implication is that the *metric* factor of line length is the decisive variable in defining how easy or difficult is it to retrieve a description of the system. Where the same total length of line is divided into a few long lines and many short ones, as we typically find in cities, then the moving observer sees more space over the time spent in movement than if the lines are of even length, and the information obtained from the longer lines is more redundant and therefore more structural. Metric factors are thus critical to the way in which the city constitutes a field of information with both structure and unpredictability. The degree to which we can arrive at an abstract *conception* of the system as a whole from series of *perceptions* of its parts depends then on is metric organisation, or perhaps more accurately, on the topological organisation of metric properties.

06.18 The third implication is that if the series of perceptions made by the moving observer are the information in the system, and the conception which we arrive at through description retrieval is about the structure of the system, and if the structure of the system is a function of the organisation of line lengths, then it would seem to follow that the distribution of all round visual integration in the system is in an important sense a picture of its structure. It would seem to follow then than by giving a view of the system which does not prioritise directions but proceeds indifferently from each point to all others, and by this means identifying an overall structure in the grid, the 'conceptual' visual integration measure in some sense modelling *the means by which the perceptual is synchronised to become the conceptual* in urban space, that is the means by which an accumulation of perceptions is aggregated into a picture of the whole system which is at once perceptual and conceptual. Precisely because it describes the topology of line lengths in the whole system, and not just the accumulation of perceptions of the parts, it describes the objective structure of the system as whole which description retrieval seeks to grasp. One could perhaps add a rider to this. Since the critical information that visual integration extracts from the system of space is linear, and since we know that two dimensional variation can affect the pattern of visual integration quite strongly, it may be that the simpler least line map, or axial map, is a simpler and truer representation of the properties of the grid that give rise to the conceptual picture of the grid as a whole.

And from the real to the ideal grid

But whichever is the case, this picture, because it is both conceptual and built up from perceptions, and so retains both perceptual and conceptual dimensions, also has another important property. It is in its nature *allocentric*, that is, it is not tied to a particular point of view from which the system is seen, but works as a representation of all points from which the system can be seen. By being synchronised from the perceptions from which it was derived, and becoming an all at once concept of the whole, the notion of the structure of a grid seems to be projected outside ourselves and, as it were, laid on the real world. The concept of the grid is 'out there' as a representation of a system independent of our point of view on it, even though the perceptions that make it up are not.

This is a critical property from the point of view of spatial cognition, since allocentric representations of space have, as we will see below, become a key theme in the study of space in the cognitive sciences in recent years. At a common-sense level too allocentric models seem necessary to our ability to navigate grids other than from and to a single point. In learning our urban surroundings we not only learn routes from one point to others, though this may well be all that happens in the initial stages, but also we begin to form some kind of picture of the relations between destinations.

However, the notion of allocentricity also lead in another interesting direction: the ideal grid. We have already seen that a grid which is perfectly linear optimises both metric and visual integration. However, if we analyse it for all round 'conceptual' visual integration we find that it is homogenous: all intersections and all alignments of them have the same value (though 'intrinsic analysis of intersection isovists will prioritise edges and corner rather than centres). This lack of differentiation creates difficulties for the perception-based information theoretic process we have described by which a conceptual structure is retrieved from differences in the visual integration pattern, but it does have interesting implications for allocentricity. Precisely because the ideal grid does not differentiate parts visually, and prioritise one location and another, it seems to offer gives a general abstract model for allocentricity. The ideal grid is perhaps ideal because it has ideal allocentricity. In this sense may offer itself as the abstract allocentric model for all grids prior to their acquisition of structure.

Does this have implications for spatial cognition?

It is clear than that visual and therefore cognitive issue play a critical role in the shaping of the city. However, at each stage of our argument we noted that the visual was dependent on the metric: in the computation of visual integration; in the notion of perceptual (or intrinsic) visual integration of a space from a point or series of points; and in the way in which we pass from the perceptual to the conceptual in forming a picture of the city. Although we can say then that metric integration does not in itself give a plausible structure in the city and visual integration does, we must also say that visual integration does so by adapting metric information to its own purposes at every stage. It is not then true to say that cities are shaped by visual and not metric factors, but by the interaction of the visual and the metric. But because the intervention of the metric is, while pervasive, on the terms of the visual it is perhaps justifiable to say that cities are *more* shaped by the visual than the metric, and therefore more by minds than bodies.

If then visual and the cognitive are then dominant in the formation of cities, then clearly this raises vital question about spatial cognition. For our purposes here the key questions seem to be these. If cities are indeed shaped by human individuals in the light of laws governing metric and visual integration, do they in some sense ‘know’ or acquire knowledge these laws, not formally of course, but in the sense that practical human activity often reflects physical laws? If so, do they know or acquire the notion of the grid, and perhaps even the ideal grid, since it seems to be implied by the spatial laws operating in the complex spatial situations we deal with everyday. If this were the case, then we might expect human cognition to use the conception of a more or less regular grid – through as yet without any specific angular organisation (we will see below where this might come from) – as a kind of reference point for dealing with complex spatial situations such as those found in cities[3].

To explore this possibility would of course require a research programme which would cover both the empirical study of behavioural regularities under different kinds of grid conditions, but also the neural modelling of how such a concept of spatial knowledge might be encoded in the brain. The latter is entirely beyond the skills available within our discipline, and must await comment – if indeed it attracts comment – from our cognitive neuroscience colleagues. The former would require a cross-disciplinary research programme involving both cognitive psychologists and syntax specialists.

Pending such enquiries, and also perhaps to facilitate them, we will now review three kinds of evidence which seem to lead in the direction of this conjecture: first, evidence from everyday spatial behaviour and interpretation of spatial situations which suggest that people have abstract and even theoretical knowledge of the laws of metric and visual integration; second, indirect evidence from syntactic studies, and one significant piece of direct evidence which suggest that the linearisation of space and grid approximating forms play a key role in spatial cognition as reflected in movement patterns and direction finding; and third, evidence from several sources in neuroscience studies that suggest such a model is far from incompatible with current theoretical positions within neuroscience on how spatial cognition works in complex artificial environments.

Do people know the laws of metric and visual integration?

First, we look for evidence in everyday spatial behaviour that an allocentric and abstract, or even quasi-theoretical, understanding of the laws of metric and visual integration underlie everyday spatial behaviour and interpretation of spatial situations. In (Hillier, 2002, 2003) I told a story of my two year old grandson placing balloons on strings at head height in the centre of an adult interaction space where it would maximise visual and metric interference to the adult conversation, an adult moving them to a corner where the effect would be diminished, and then my grandson moving them back to the centre. This was interesting, because what the child seemed to be doing was, although psychologically ego-centred in the sense that he was drawing attention to himself, clearly spatially *allocentric*, since the placement of the balloons suggested awareness of the relation between a set of positions independent of his own momentary position. Barbara Tversky (Tversky, 2002) has suggested, however, that the child's behaviours could be interpreted more simply as the child moving the balloons back to where he had chosen to place them. My reply would be that even if this were the case – and having seen it happen I suspect it was not - the *adult* behaviour still implies an allocentric appreciation of these laws, in that the adult reacted to the general nuisance of the balloons, not a nuisance in relation to himself, since he was standing at the time.

However, without trying to resolve this particular case, let us look at a richer set of examples which seem to show evidence of a more general allocentric awareness of the visual and metric laws of space: the social logic of table shapes. We have already seen that in general as we go from a square shape to a very elongated shape we increase universal distance in the shape and we changed its pattern. The more we elongate the shape, the farther we have to go to get from all points to all other points, but the more visually integrated the shape becomes from its most asymmetric locations, namely the two ends. There are two aspects to this: the pattern of integration of the interior; and that of the edge. The more elongated the form, of course, the higher the ratio of edge to interior, so a long thin form maximises the edge to interior ratio, so holding area constant you can have more people around a long thin table than a square or round table for that amount of surface area. These are simple facts of the spatial world 'out there', although easily discovered by anyone who seeks both to seat people at a table and also to place a large number of objects on the table.

Let us then look at the 'social logic' of table shapes considering two things: first where different kinds of people sit at differently shaped tables in different circumstances; and second, at where the camera typically tends to locate to take picture of the table, since it is argued that this will always give emphasis to the key metric and visual relationships – implying, by the way, that cameramen and directors know the laws of space just as much as those sitting at tables. Human beings of course sit around the perimeter of the table, but their metric and visual integration with others is determined by the whole table, not just its periphery. So let us look at perimeter locations in the context of the whole table shape in terms or visual and metric relations and consider the effects of varying the shape of the table. The key argument here is that in interpreting the social 'meaning' of table shapes we are implicitly using knowledge of the laws of metric and visual integration.

We start top left in Figure 17: the simple case of a circular table. This is the most integrated shape possible, and the only one in which all points on the periphery are equally metrically and visually integrated. We therefore associate round tables with equality, as in the 'knights of the round table', though with the caveat that in a spatially equal situation any inequalities in status will tend to become powerful
06.22 even without being reinforced by privileged edge locations. In the square shape, which is also metrically integrated, some inequalities are introduced into the situation in that the mid points on the side are more metrically integrated than the corners, but have less visual integration. The centre-edge locations are then better if you want to have one conversation at the table. In the corners you are more likely to talk to your neighbour, or survey the scene from a distance. But every hostess knows that you do not place the key conversation makers in the corners but in the centre edge locations if you have a square table. However the square table is still broadly egalitarian compared to other shapes, and in both the circular and round cases the position of the camera would be diagonal, since any end-on view would suggest a symmetrical distribution or people around a focal person and this would go against the egalitarian ideology (not the practice) of round tables. So we do not see King Arthur head-on, with knights on either side, since this would suggest a spatial focality which was alien to the table shape and its social logic.

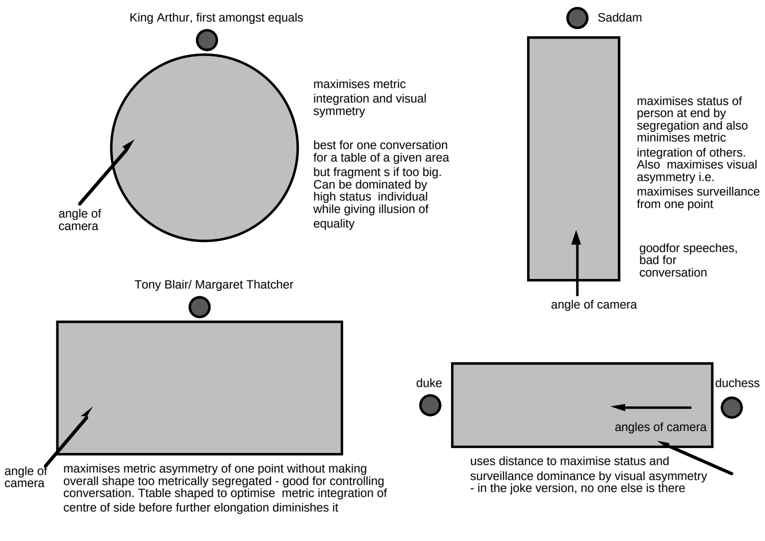

Figure 17

However, this is the only way we are ever permitted to see Saddam: in a head-on view at the end of a long thin table, with cabinet member disposed either side of him. Note that we never see the other end of the table. This maximises the status of the person at the end by maximising metric segregation from others while also maximising visual asymmetry i.e. it maximises surveillance from one point and minimises it from all others. One could note that this is the best shape for

Not all leaders opt for this social logic. Tony Blair – and Margaret Thatcher previously - is always pictured diagonally at the centre of the long side of a broad rectangular or oval table. We can see why by referring back to Figure 4 and exploring the changing shapes a little more. As we move from a square to a more and more elongated shape, the actual metric integration value of the centre of the long sides first increases, because the decrease in metric integration of the whole object is at first compensated by the fact that the centre of the whole object is moved closer to the edge, then decreases as the effects of elongation become stronger than these initial effects. There is therefore a definite ratio of the long to the short sides of a rectangle where the metric integration of the centre of the long edge is optimised. More experimentation (or a theorem) will be needed to pinpoint exactly what this ratio is, but a first approximation suggests it is in the region of 9:4. With a broad rectangular table approximating the optimal ratio, the advantage of the centre edge locations is increased as the mid points on the long side become more integrated than any in the square, but the short sides have become more segregated. At the same time the visual integration has become much less good from these centre-edge locations, especially when compared to the centres of the short sides. However, this is where Tony Blair and Margaret Thatcher sit, because it is metric integration that gives advantage in controlling a conversation, and this is also where a hostess with a table this shape will place her key conversation making guests, while retaining the end location for either herself or high status low conversation persons since this give advantages for surveillance rather than conversation. These effects will of course still be present in an oval table which conserves something like the optimal ratio, though with increases both in metric and visual integration from all peripheral points.

If we elongate the rectangular shape more, then and the whole periphery becomes more segregated, and the advantage of the centre of the long side disappears. So we find that the duke and duchess sit at either end of a very long table to emphasis their status through metric segregation and maximise visual surveillance through visual integration. The joke is of course that even when there are no other guests, the duke and duchess still sit at the two ends where they cannot converse. But jokes are of course about structure, and only work if we know what the structures are. And of course we do. In this long thin case, unlike the one-ended Saddam case, the camera angle is likely to be diagonal in order to see both ends.

There are other aspects to the social logic of table shapes. For example the long thin refectory table by minimising both metric and visual integration for those along the sides minimises the opportunity for anything but the most localised interaction; or the school version of the same table where the monitor sits at the end to establish status and maximise surveillance. But already we can see that in all these cases space is a social strategy of knowledgeable agents. Patterns of metric and visual integration are lawful properties of shapes, and because they are so, their potentials can be exploited by knowledgeable individuals as social strategies, in much the same way as we make use of the intuitively know the laws of physics in throwing an object to that it lands in a certain place.

Is there evidence that people use an allocentric cognitive grid in wayfinding?

Next we consider indirect evidence from syntactic studies that grid regularity may play a role in spatial cognition in that where the grid is overly broken up and de-linearised the relation between local integration structure and movement, which is robustly approximately in most 'normal' urban environment patterns, breaks down. In the series of studies reported in 1987 (Hillier et al., 1987), the first published reporting of a systematic relation between grid structure and movement, it was argued there where environments became unintelligible, as measured by the degree of agreement (indexed by the r^2 value) between the connectivity and integration values of the lines making up the axial map of the grid, the relation between movement and local space structure was no longer found, but replaced by a relation reflecting the depth of lines from the surrounding grid (and so partially reflected in the global integration pattern of the wider context, since internally this reflected the layers of depth in the local system). The geometric means by which syntactic unintelligibility was created involved the breaking of alignments in the local grid, more or less on the lines outlined theoretically in (Hillier, 2001, 2003) where it is shown that if at least some line lengths are not increased to a certain proportion of the diameter of the system as it grows larger, then unintelligibility is the inevitable result.

In the light of these results, an experiment was constructed by Ruth Conroy-Dalton which deserves greater fame than it has so far achieved. During her PhD (Conroy, 2000) which sought to answer the question how far movement in immersive virtual worlds resembled that in real worlds, she took two theoretical exemplars that had been previously constructed to illustrate the notion of an intelligible and unintelligible urban environment, as in Figures 18 (top left and right), and turned them into navigable 3D immersive worlds. The difference between the two worlds was simply that in the intelligible world, blocks were arranged so as to create urban-type linear continuity between spaces, whereas in the unintelligible world the blocks were slightly moved so as to break the linear continuities. Simply breaking these alignments seems enough to make the latter environment appear unintelligible and even labyrinthine even in plan.

However, Conroy-Dalton's concern was with how people would navigate this environment internally. Accordingly, she constructed an experiment in which 30 subjects were asked to start outside the world around the mid-point on the left, then to find their way to the monument (which was always invisible due to the height of the buildings, until the points where it became wholly visible) in the main open space (half right) and then to return to their starting point by another route. The results are shown in Figures 18 (bottom left and right). They show that while in the intelligible environment the traces of subjects' navigations were largely confined to a relatively narrow envelope between the origin and destination points, those in the unintelligible world were spread all over the system, including parts well beyond the destination. While only a single study, Conroy-Dalton's results strongly support the conjecture that our cognition of the urban environment depends in some way on its linear, and even on its grid, organisation.

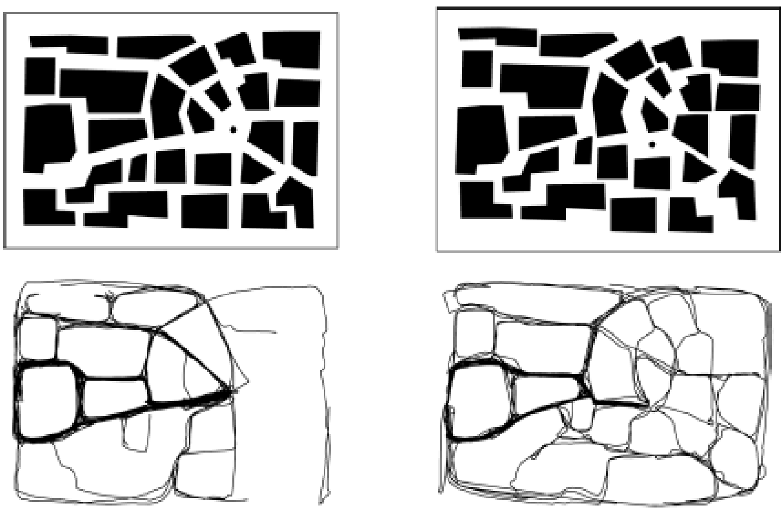

Figure 18

These findings offer preliminary, though strongly suggestive, evidence for the conjecture that a linearised grid serves as a cognitive model for negotiating spatial complexity of the urban type. Similar evidence seems to come from the simple facts of direction-giving. Although it is customary within our subject area to assign direction finding to landmarks and their inter-relations, this seems obviously wrong (perhaps a product of the lack of a formal language for the formal description of spatial patterns) when we consider every day behaviours in appropriate detail. For example, such common or garden instructions as: ‘carry straight on down this road (even if it is not really straight), then fork right in about 200 metres, and take second on the left at the petrol station and then third right opposite the Dog and Duck’ clearly implies some kind of general – though often highly inaccurate – geometrical picture of the local grid with the landmarks in a clarifying role - since they would make no sense without the underlying geometry to link them. Equally clearly, the same direction-giver would be able to give comparable directions from and to other locations in the same grid, and this shows clearly that the spatial representation being used is transformable for different ego-centred position, and therefore allocentric. Indeed, it is because the model is both allocentric and relatively abstracted from concrete details (though these undoubtedly serve to hold it in place), that it can be used in this way as a basis for configurational - all points to others - problem solving – which is what direction-giving really is (see Hillier, 1997, 1999 for a further discussion).

Between the body and the universe: one view from cognitive science:
It has to be acknowledged, however, that this line of thought seems at first sight to be to some degree at variance with some recent opinion in the cognitive sciences. In her review at the Third Symposium (published in more extended version in Tversky et al., 1999, to which reference is made here), Barbara Tversky described the spatial knowledge of the ‘navigational space’ in which we move as a ‘cognitive collage’ in preference to a ‘cognitive map’, because it seemed to be ‘not Euclidean’ and ‘represented qualitatively, in terms of elements and the coarse spatial relations among them’. This seems initially hard to reconcile with the position being advocated here, but closer examination of the argument suggest it may not be. In their paper, Tversky et al describe ‘three spaces of spatial cognition’: the space of the body, the space surrounding the body, and the ‘navigational space’ in which we move. They also make extensive reference to a fourth level: the vital role in spatial cognition played by universal and macro-geographical frameworks, most notably the sense of global orientation that come from the cardinal points of the compass. Properly speaking, navigational space seems to come between the body and the universe, that is between the spatial models we project from out bodies into our immediate surroundings, and the macro-awareness we have of large scale all-round directionality.

Now according to Tversky et al the first space between the body and the universe, that around the body, has a clear geometry, in that ‘people construct a mental spatial framework from the extensions of the three axes of the body (front-back, left-right, and up-down) and associate objects to it’. So does the macro-picture that we derive from the cardinal direction, and Tversky et al show how powerful this is cognitively, since it leads people to misrepresent the relation of Europe to the

US as east-west, and North to South America as north-south, when neither is the case. In the Tversky et al model, however, neither the small not the large geometry seems to penetrate the *structure* of the navigational space between the body and the universe to any extent. One a priori reason for expecting that some degree of penetration might occur would be that if we move a person with geometric body axes from point to point within a space where overall directionality is governed by a macro-geometric model of the universe, then part of the outcome would seem to be some kind of conceptual grid, without dimensionality, it is true, but a grid nonetheless. Indeed, in the concept of the grid which has been derived from the metric and visual laws of space earlier in this paper, the angularity which was missing from this grid seems to be most naturally supplied by Tversky et al's bodily and universal geometries.

It was partly to test this possibility that I took the opportunity offered by my recent move to the City of London to conduct a year long experiment in finding my way on foot from the Barbican to the Space Syntax Laboratory in UCL, deliberately not consulting maps, until after several months of trial and error and failing to find a way that satisfied both my cognitive model of the area (I know London extremely well) and my stop-watch which over several repeats of the same route I expected to give fairly reliable information on total route length. Route finding was essentially a matter of conceptualising the position of UCL in relation to my starting point and then using my almost complete knowledge of the intervening grid to find the shortest route. What essentially happened is the more I tried to follow Ruth Conroy's principle that at every point one tries to minimise the angle of deflection from the conceived target (I believe she is quite correct in this in terms of what people seek to do, although as we will see in certain kind of local grid conditions it can be precisely this that misleads people), the longer it seemed to be taking me, and the more I routed myself around the southerly edge of the grid. as I conceived it, the less time I was taking. Conversely, the more I tried to navigate what I took to be the northern edges of the intervening grid, the longer it took me.

After months of error, and unable to reconcile my intuitions with my timing, I decided to make a careful study of the map. A number of things immediately became clear, some very much in accordance with the Tversky et al model, though others not obviously so. First, in my cognitive map of the area, I clearly conceptualised the overall direction from the Barbican to UCL as south-east to north-west, when in reality, and to my astonishment as a knowledgeable agent in London, it is almost due west. I had also conceived of the alignment of the Holborn-Oxford Street axis as being east west, when it is closer to west-south-west to east-north-east, and of the alignment of Gower Street as north-south, when it is closer to north-north-west to south-south-east. Both of these seem cases of the type of error that Tversky et al describe as resulting from the imposition of an external reference frame – the cardinal directions - on the area to be navigated.

However, it was also clear that other errors followed from this, notably that the cognitive model I had imposed on the area, perhaps partly under the influence of my understanding or the cardinal directions, took the form of a consistent Euclidean grid, and this turned out to be even more misleading than the cardinal directions error itself. Most importantly, I had completely misrepresented two key intrinsic facts about the grid: that different parts of the grid were offset against each other by being differently oriented; and that different lines in the grid has different angular connections when compared to each other. I had in fact used the cardinal directions model to construct not a collage but a Euclidean misrepresentation of the local grid, and it was this tidied up grid that was the clear source of most of my errors.

It was, for example, the errors in my model of the larger scale that led me to believe that in navigating from the Barbican to UCL I should continually try to 'climb' the grid to avoid seemingly orienting myself more and more away from my destination. However, this would not have mattered had the structure of the grid been consistently Euclidean as it was in my mental model. It was the fact that the grid shifted orientation from one part to the next that made me unaware that in climbing the grid to maintain the conceived orientation towards UCL all I was doing was adding distance because one grid converged on the other. The greatest single error in the mental map with assuming that Holborn and Theobalds Road were more or less parallel, when in fact they were rapidly converging – a fact that was fully familiar to me through extensive past travel in the area but for some reason did not register on my model.

My problem in short was that my mental model was Euclidean, but the environmental reality was not. Nor did it seem that my errors lay in the topology of the model, since that was more or less correct. My errors were decidedly geometrical. Moreover, although my model was wrong, it was a the same time clearly allocentric, as it involved continuous experimention with the local microstructure of the grid, and attempts to move to sub-destinations from sub-origins throughout the experiment. So my allocentric model was both the means by which I could address the environment to solve problems within it, and also erroneous. In fact a model that is allocentric, Euclidean, and erroneous, is precisely what some opinions in cognitive neuro-science would lead us to expect. Let us then examine some of these opinions.

Other views from cognitive neuroscience

Although most cognitive theories of space influential in the built environment have been decisively ego-centred, for example 'personal space' and 'human territoriality, and even the ways in which cognitive maps have been brought into our field has placed undue emphasis on the position of the subject at the centre of the map, not least by requiring the map to be loaded with value (and other) information (O'Keefe and Nadel, 1978: 75), in the cognitive neuro-science the distinction between allocentric and ego-centred models of cognition is a principle theme. As Petersen et al argue in the final chapter of a recent neuroscience reader on spatial cognition and language: 'It is now well established that the vertebrate hippocampus subserves a spatial mapping function that is both multimodal and allocentric; that is, external space is represented independent of the momentary position of the organism, in terms of the relations between objects and the places they occupy in what appears to be an objective, absolute framework'. (Petersen et al., 1996: 556)

Advocacy of an allocentric view goes back at least to one of the seminal books in cognitive science in the last quarter of the of the twentieth century – O'Keefe' and Nadel's 'The Hippocampus as a cognitive' map which focused on issues which are very close to our theme here. O'Keefe and Nadel contrasts 'routes' to 'maps' as cognitive entities and as metaphors for different cognitive bases for spatial behaviour. Routes are associated with the taxon system, and can be seen as connected series of specific behaviours implying landmarks and specific responses to them. Maps are associated with the 'locale' system, and imply the availability of 'an aggregate of interrelated information with no necessary specification of guides'. The 'cognitive maps', which he argues make up much of the hippocampus, are of the latter kind, and generate that he calls 'place hypotheses and exploration'. Although both maps and routes are clearly used in different mixes in different circumstances, routes having the advantage of being very specific in particular circumstances, and the disadvantage of being inflexible – for example, they cannot easily be reversed or used to solve other direction-finding problems with – and of being highly vulnerable to partial loss of information, while maps have the disadvantage of being much less specific, but also the advantages of being much more flexible – they can be used to solve the whole range of local direction-finding problems – and much more robust under partial loss of information.

What then are these allocentric spatial representations like ? In contrasting the respective roles of 'conceptual structure' and 'spatial representation' at the interface between spatial cognition and language, Jackendoff argues that spatial representations are not images, but 'geometric (or even quasi-topological) in character, rather than algebraic', and must be 'independent of spatial modality' and 'suitable for encoding the full spatial layout of a scene and for mediating among alternative perspectives ("What would this scene look like from over there"), so it can be used to support reading, navigating and giving instructions'. (Jackendoff, 1996: 9). It seems that although these spatial representations have the capacity to organise and inter-relate items from the physical world, they are not in themselves anything like representations of physical entities, but have a more abstract character.

Johnson-Laird is more explicit. He argues that we use spatial 'models' to mediate our relation with circumstances, but he emphases the formal aspects of the models. Talking of how we solve problems of a series of propositions describing the relative location of objects, he argues: 'The most likely way in which an inference is made involves setting up an internal representation of the scene depicted by the premises. The representation may be a vivid image or a fleeting abstract delineation – its substance is of no concern. The crucial thing is that its *formal* properties mirror the spatial relations of the scene, so that the conclusion can be read off in almost as direct a fashion as from an actual array of objects' (Johnson-Laird, 1996: 437) – my emphasis). He goes on: 'The key feature of spatial models is not that they represent spatial relations – propositional representations also do that – but rather they are functionally organised on spatial axes and, in particular, that the information in them can be accessed by way of these axes' (Johnson-Laird, 1996: 445-6) . He even takes the argument one step farther: 'Human reasoners use functionally spatial models to think about space, but they also appear to use such models in order to think in general' (Johnson-Laird, 1996: 460).

Bowman extends this theme arguing a fundamental role for spatial cognition in structuring thought in general: 'If any domain has plausible claim to strong language-independent perceptual and cognitive organisation, it is space…..Our mental representations of space are constrained not only by our biology but also by their fit to the world 'out there. ….Little wonder it has seemed likely to many investigators that the language of space closely mirrors to contours of non-linguistic spatial understanding. Several kinds of empirical evidence support the assumption that children know a great deal about space before they can talk about it, and that they draw on this knowledge in acquiring spatial words.' (Bowman, 1996: 387).

Finally we come back to Petersen et al: 'Some, but not all, of the spatial maps identified by neurobiological and behavioural research impose a structure that goes beyond, and in consequence alters, our interpretation of the information available in the input alone. For example, the hippocampus appears to impose a Euclidean framework onto non-Eucllidean inputs (O'Keefe and Nadel, 1978) who see in this process the instantiation of a Kantian a priori notion of absolute space…. We propose that in ‘distorting’ the sensory inputs these spatial maps may impose an order and a structure that our spatial conceptual representations require.’ And finally: ‘We point out the importance of a careful analysis of the intrinsic ‘organising factors’ that interact with spatial information to structure our knowledge of the spatial world. These organising factors act like a kind of ‘syntax in accord with which inputs to spatial systems are ordered, and in doing so they contribute meaning to the spatial representations themselves. This is perhaps clearest in the allocentric map observed in the hippocampus, but is also observable in other cases’ (Petersen et al., 1996: 569)

So, is there a syntax of spatial cognition?

It is of course something like a syntax that our model proposes. In arguing that people use acquired knowledge of metric and visual integration as a means of both structuring and understanding the artificial spatial environments constructed by architects and planners, and of the allocentric grid as a limiting form, we are in effect proposing something like a syntax of spatial cognition. Is this possibility worth exploring? At least one important body of opinion would oppose this idea in principle. In making the case for the ‘embodied’ mind (in contrast to the Cartesian disembodied mind), Lakoff and Johnson argue that elementary schemes of spatial relations, such as those found in the English propositions, are both perceptual and conceptual, and as such offer powerful structuring devices for thought in general. But this works, they argue, not by creating a syntax, but by the metaphorisation of the elementary spatial schemes. Their opposition to the idea of a more complex formal development is summed up in their critique of Chomsky, who they regard as only the latest inheritor of the Western a priori philosophies of the disembodied mind: ‘Syntax’ they argue, ‘is real enough, but it is neither autonomous nor constituted by meaningless, unintepreted symbols. Rather it is the study of symbolisation – the pairing of meaning with linguistic expressions, that is, with phonological forms or categories of phonological forms….from a neural perspective, symbolisation is just a way of discussing neural connectivity’ (Lakoff and Johnson, 1999: 498).

If we were to proceed by strict analogy between a syntax of space and that of language – and I have strongly argument elsewhere that no such analogy can be made (Hillier and Hanson, 1984: Chapter 1) – this would seem to imply that Lakoff and Johnson would be at least sceptical of the idea that even architectural and urban space was accessible through a cognitive syntax. However, on closer examination, it is not obvious that Lakoff and Johnson’s argument is fully consistent in these respects. For example, they argue that at the level of elementary schemes of spatial relations, such as that involved in the word ‘contain’, the traditional distinction between the perceptual and the conceptual is obliterated. They do not say so explicitly, but this is clearly because what we are seeing *is* a relational scheme, and relational schemes are in their nature 'conceptual' – indeed the difficulty philosophers always found in granting reality to relations was their conceptual rather than visual nature.

The problem with Lakoff and Johnson's argument is why should the obliteration of the distinction between the perceptual and the conceptual be confined to elementary schemes of relations and not re-appear to some degree at least in the more complex spatial situations that typify everyday life, whether in forests or cities. It seems a priori (this is not intended as a provocation) far more likely that what is the case with elementary schemes of relations should also be found to some degree in more complex situations. The obvious question then arises: might not the elementary schemes themselves form the basis for some kind of syntax (as was suggested in Hillier and Hanson, 1984: Chapter 2), and on what grounds would be expect that they would not? Why should their extension from elementary situations be confined to metaphorisation?

The second point at which Lakoff and Johnson's arguments might be turned around concerns learning. 'Why', they ask, 'is it possible for our concepts to fit so well with the way we function in the world. They fit so well because they have evolved from our sensorinotor systems, which have in turn evolved to allow us to function well in our physical environment.' (Lakoff and Johnson, 1999: 44). If this is the case, then surely we would expect spatial learning to go beyond elementary spatial schemes, and begin at least to engage with some of the configurational complexity that is necessarily involved both in living in the material world and in living with others in that world. No one who has watched lions hunt can doubt that complex configurational calculations involving several lions and at least one prey, are made throughout the hunt. To propose that this reflects some kind of relation between the lawful behaviour of spatial configuration and the brains of discrete beings inhabiting them is no more surprising than we internalise enough of the laws of physics to throw an object to that its parabola leads it to fall or strike at a particular point. If our spatial knowledge were not in this sense lawful, then we would surely not be here.

By far the most likely reason for the acquisition of lawful knowledge of spatial configuration would seem to be of the kind that Lakoff and Johnson describe: what we learn is the *invariance* of the spatial behaviour of the complex environments in which we live. To learn to throw a projectile so that its parabola leads it to strike a certain point in three dimensional space surely depends on having learnt the spatial invariance of the material world. It would surely not be unlikely if something similar turned out to be the case with space. This suggests a fundamental link to Gibson

(Gibson, 1986), whose core concept for perception is that what we learn to see is the invariants of objects. What we learn about space, perhaps, are the invariant behaviours that space adopts as we experiment with it and interact with it by moving about in it and placing objects in it. These are of necessity configurational properties which affect how we see and how we go.

We do not there need necessarily to agree with O'Keefe that the spatial cognition that we impose on our surroundings is a pure Kantian system, since its most likely origins would seem to lie in the invariance that is found to underlie our spatial transactions with our ambient world However since such cognitive models seems to have been crucially involved in the construction as well as the understanding of the spatial order of the city, this may prove to be a route to save the Kantian hypothesis.

Notes

[1] This analysis of visual integration also seems to offer a possible redefinition of the question of whether or not least line maps might be algorithmically defined. Visual integration clearly offers a way to define the fewest and longest lines between any pair of points in a complex space. The general question then becomes: is there a set of lines which would serve all pairs*?*

[2] In terms of the constraints on real space patterns of course we must also take into account the contrary geometries of moving in and occupying spaces, as set out in Chapter 8 of Hillier 1996a, since this will always be a factor promoting the greater width of at least some spaces in the system, as will of course also movement capacity issues. Here we are dealing with theoretical limits.

[3] If we do seek to assign a cognitive role to the orthogonal grid as a reference point for dealing with complex urban spaces, does this then challenge the syntactic idea of intelligibility, which is expected to be weaker in a pure grid than a deformed grid. It does not. The grid we are proposing here is a synchronic conceptual model, giving an all at once – and probably simplified - picture of the space of a system. Syntactic intelligibility is about the ease which with a synchronic picture of the grid can be built up step by step by moving about in it and seeing it from different points, and this, it would seem, might actually involve this abstract conceptual model as a reference point. So if anything, the theory of the grid as a conceptual model seems to clarify the concept of syntactic intelligibility.

References

Bowman, M., 1996, "Learning how to structure space for language", in P. Bloom, M. Petersen, L. Nadel and M. Garrett (eds.), *Language and Space*, Cambridge, MA, MIT Press

Carvalho, R., Iida, S., and Penn, A., 2003, "Scaling and universality in the micro-structure of urban space", *Proceedings of the Fourth Space Syntax Symposium 2003*, London

Conroy, R., 2000, *Spatial Navigation in Immersive Virtual Environments,* PhD thesis University of London (UCL)

Gibson, J., 1986, *The Ecological Approach to Visual Perception*, Lawrence Erlbaum Associates

Hillier, B., et al., 1986, *Spatial configuration and use density at the urban level: appendix on distance factors*, Unit for Architectural Studies, UCL

Hillier, B. et al., 1987, "Creating life; or, does architecture determine anything", *Architecture & Behaviour*, Editions de la Tour

Hillier, B., 1996a, *Space is the Machine*, Cambridge, Cambridge University Press

Hillier, B., 1999, "Centrality as a process: accounting for attraction inequalities in deformed grids", *Urban Design International*, Volume 3/4

Hillier, B., 2001, 2003a, “A theory of the city as object”, *Urban Design International*, 2003, also *Proceedings of the Third International Space Syntax Symposium 2001*, Atlanta, http://undertow.arch.gatech.edu/homepages/3sss/

Hillier, B., 2003b, “The knowledge that shapes the city”*, Proceedings of the Fourth Space Syntax Symposium 2003*, London

Jackendoff, 1996, “The architecture of the linguistic-spatial interface”, in P. Bloom, M. Petersen, L. Nadel and M. Garrett (eds.), *Language and Space*, Cambridge, MA, MIT Press

Johnson-Laird P, 1996, “Space to think”, in P. Bloom, M. Petersen, L. Nadel and M. Garrett (eds.), *Language and Space*, Cambridge, MA, MIT Press

Lakoff, G. and Johnson, M., 1999, *Philosophy in the Flesh*, Basic Books

O’Keefe, J. and Nadel, L., 1978, *The Hippocampus as a Cognitive Map*, Oxford, Clarendon Press

Penn, A., 2003, “The shape of habitable space, *Proceedings of the Fourth Space Syntax Symposium 2003*, London

Petersen et al., 1996, “Space and language”, in P. Bloom, M. Petersen, L. Nadel, M. Garrett (eds.), *Language and Space*, Cambridge, MA, MIT Press

Shannon, C. and Weaver, W., 1948, *The Mathematical Theory of Communication*, Chicago, IL, University of Illinois Press.

Tversky, B., Morrison, J.,Franklin, N. and Bryant D.,1999, “Three spaces of spatial cognition”, *Professional Geographer,* 51 (4)

Tversky, B. 2002, *in conversation*

Turner, A, 2002, “Depthmap, v2.11 (computer program) UCL”, London, introduced in Turner, A, 2001, “Depthmap: a program to perform visibility graph analysis”, *Proceedings of the Third International Symposium on Space Syntax 2001*, Atlanta, GA, pp. 31.1-31.9