Kirsh, D. Changing the Rules: Architecture in the new Millennium. Convergence. The Journal of Research into New Media Technologies. 2001 Vol 7(2) pp 113-125.

Changing the Rules

Architecture and the New Millennium

David Kirsh

Architecture is about to enter its first magical phase: a time when buildings actively co-operate with their inhabitants; when objects know what they are, where they are, what is near them; when social and physical space lose their tight coupling; when walls and partitions change with mood and task. As engineers and scientists explore how to digitise the world around us, the classical constraints of design, ruled so long by the physics of space, time, and material, are starting to crumble. Documents can be laid down in one place, automatically cloned, and a copy picked up in another. Meetings scheduled for 9am to 10am can be joined by latecomers at noon, who then participate in a captured form of the event and are 'edited into' the past. People on the West coast of the USA can participate, in a telepresent way, with their colleagues on the East coast, and hold a meeting against a virtual backdrop, such as a production line in their Taiwanese factory. Walls can seem to dematerialise, remote objects can be touched virtually, reshaped, passed through one another. Technology is moving inexorably so that being in one place at one time no longer need dominate how we work and play. Material boundaries are losing their meaning, and interface and information space are catch words that architects must master.

In this article I will discuss some of the theoretical ideas shaping our new conception of form, function and interactivity. My view is that of a cognitive scientist interested in how cognition is distributed throughout our environment. Since the ground rules defining the structure of environments are changing, our very idea of how we are embedded in the world is changing. Architecture is at a new frontier.

Redefining activity space

Much of architecture is about boundaries: defining space and movement by the thoughtful design of walls, enclosures and openings. The spaces I am personally interested in are the workspaces we spend most of our time in: offices, kitchens and meeting rooms. The boundaries here often blur the distinction between architecture and furniture, but the problems are constant. How do we design surfaces to facilitate the activities which our clients – the users of these spaces – want to perform?

The biggest question this raises is about the nature of activity. Do we really know what people do? There was a time when this sort of question did not seem so challenging. To discover the structure of a person's behaviour it was assumed sufficient to observe their movement, look at the goals they have, and the methods embodied in their action. We now know this is inadequate. Behaviour is far more complex, more densely interactive than this simple approach assumes.

Consider activity in a kitchen. According to cognitive theory a few years ago, the underlying structure of cooking activity ought to resemble the structure derived from the recipes specifying the steps to be taken. First you get the pots and pans you will need, then the ingredients, you heat the pan, and so on. But close observation shows that cooks are doing much more. Instead of proceeding step by step through a recipe we busy ourselves with additional actions. We clear our workspace, we wash dishes, we return items to the refrigerator as we finish using them, we lay out ingredients in a spatially interesting manner, we stop to read the recipe book, move it from place to place, we fill the salt shakers, nibble on snacks, pick up the phone, jot down notes or reminders, put on oven gloves, etc. None of these extras are written down as part of the recipe. But they matter to successful performance. Indeed, the best cooking books offer hints about these sort of things in the long discursive parts of the books (except for the nibbling, phone calling and annotation).

It is precisely these extras that good designers want to understand so they can place cutting boards where they will be most useful, so the phone can be conveniently positioned, so lighting is well sited, so open space will be available where movement and collaboration is most intense. Even when we think we understand the structure of an activity, say making a white sauce, cooks surprise us by the way they multitask, constantly thinking ahead and preparing the environment for future action. They rely on tricks for reducing memory load and speeding performance, they discover techniques for augmenting their perception, and improving accuracy. Evidently, everyday activity has more complexity than commonly appreciated.[1]

To represent this complexity cognitive scientists have been working on a conception of activity space that helps to make sense of behaviour. Formally, an activity space is an abstract blend of several components: the problem constraints or sub-goal structure implicit in a task, the physical space in which the task is to be performed, the sorts of actions an agent is capable of, as well as the concepts, plans and other intellectual or representational resources agents find in the environment or bring along in their head. An activity space is partly a mental projection on the part of the agent, partly a system of hard constraints imposed by the physical layout, and partly a set of logical dependencies derived from the sub-goal structure of the task itself.[2] With all this intellectual equipment you'd think we would have a good account of behaviour. Yet once again the richness of human activity makes this definition imperfect. Too often people perform seemingly irrelevant actions by habit, such as biting their nails when they think, or repeatedly washing the cutting board. We can't ignore these actions, as much as we might like to idealise from the facts, because when denied these habitual acts, agent performance often degrades.

Moreover, and of greater consequence right now, whenever new technology is added to the mix, the job of understanding activity gets harder. Look at the effect which the television remote control unit has had on TV watching behaviour. Remote control has completely changed the *cost structure* of TV activity. In the days before remote control a person had to get out of the chair, walk to the TV, then change channels manually. If they wanted to check several channels they stayed in front of the TV, partially occluding the view of others, and rotated the channel dial to incrementally review availability. Reception areas with more than a hundred channels would have been unmanageable without random access (not to mention the trouble of having a dial with a hundred stops). Now, channel review can take place while the 'driver' remains seated. Channels can be randomly selected, even toggled back and forth. One result, apart from weight gain in the USA population, is channel surfing, an 'emergent' behaviour that is now intricately bound up with the way television programming is edited, delivered, and promoted.

To my knowledge no one successfully predicted the emergent behaviour of channel surfing. In retrospect, it is a plausible dynamic predictable from the change in the activity space associated with the 'task' of TV watching. But why did no one predict it? One reason, no doubt, is that very few theorists, if any, mapped out the activity space of TV watching, and analysed the cost structure of different TV watching environments. Another reason is that people are opportunistic. They co-opt the normal function of artefacts for uses the inventors of those artefacts never envisioned. So although remote control was not meant to facilitate channel surfing (as opposed to channel review), people soon found that the functionality of remotes supports a predilection they never knew they had.

If remote control can change our understanding of TV watching behaviour – a form of activity that is relatively simple – imagine how extreme the effects which digitally enhancing a desk or work surface will have on office behaviour. Imagine the effect of 'remote control' desks, filing cabinets, doors, lighting, and wall transparency. E-mail has already reshaped the average office day. But e-mail is just the start. How will office life be restructured when we have constant telepresent access to other offices? How will workflow change when our office helps us find documents and misplaced notes? What sort of partnerships will we develop with our smart desks?

And this is just the beginning. Computation is moving out of concentrated areas, such as computers, and into our walls, desks, ceilings and furniture in ways we have barely begun exploring. In a limited way in laboratories right now, rooms can sense activity, morph their appearance, and dynamically adapt to workflow or activity needs. We can collimate sound so that two people can speak across a room without others hearing them. We can project light onto walls or through translucent surfaces and create new types of dynamic trompe l'oeil. And, of course, we can create shared virtual environments inhabited by digital libraries, software agents, avatars, and telepresent versions of ourselves which can assist in the creation of interactive information spaces – *three dimensional intranets* – which can help us and our colleagues manage the impossible increase in information now confronting us. As William Mitchell put it, 'works of architecture function as both shelter and symbol, and the introduction of digital technology opens up new ways to perform the symbolising role. The resulting restructuring of a building's basic organisation compares to that which resulted when artificial light appeared as an alternative to natural light.'[3]

For a cognitive theorist this wild rush of new technology represents an opportunity to develop new theoretical analyses to help us avoid costly design mistakes and build more congenial structures. Just as thoughtful analyses might have led to predicting channel surfing, so it may be that with thoughtful analyses we can anticipate the way new technologies will change our workspace. If we do not start now, we can expect a string of trial and error designs that miss the mark.

In this new world, cognitive scientists and architects have a special opportunity to join forces. Whether or not we are far along with these new technologies, it is evident that one of the great architectural challenges of the twenty-first century is to unite the digital with the physical – to integrate new forms of telecommunications and computing seamlessly into everyday life. This same challenge to manage the digital and physical as complementary parts of a unified resource is an agenda for cognitive science too.

Digital objects are new kinds of objects

The first step in integrating different materials is to understand the basic nature of each. This is especially true for our efforts to wisely integrate the digital with the physical. At bottom, two things should be kept in mind:

(1) Digital and physical objects have different basic properties and so digital objects can be exploited in ways that are impossible with physical materials.

(2) Interactivity with digital objects can be defined by us, so the limiting factor in how we interact with digital objects has more to do with our imagination than with natural constraints.

Here is a little lesson worth reflecting on about the need to move beyond the customary rules of architectural composition and construction. Early efforts by computer interface designers to simplify the use of modems for faxing and making calls led to digitally recreating the appearance of telephones on the computer desktop. This seemed like a good idea, at the time, because people understand how to use telephones and don't understand how to use modems. Nothing facilitates comprehension better, it is said, than a good metaphor; and the telephone is certainly a good metaphor. Except in one respect: computer telephones support many more functions. We can enumerate a set of calls to autodial, each set for a specific time. We can save a call in digital form and then send it to colleagues. We can define response behaviours depending on the caller's identity – always take this call, let this other call ring eight times and then say we are not home. At some point the added functionality of computer telephony forced designers to give up on an image of a telephone which you answer by using your mouse to 'pick it up'. The new possibilities of merging computer and phones required new designs to support the new functionality and the new forms of interaction.

Should we expect less from the design of rooms and buildings? If we just use digital technologies as substitutes for existing physical ones, we will fail.

To make the best use of the additional functionality that the digital world can offer, we must think outside the physical box in two ways. First, we must think of how to use the digital in ways we cannot use the physical; and second, we must invent new forms of interactivity which liberate us from our conventional modes of relating to objects.

Nothing helps free thinking more than reflecting on ultimate principles, in this case, the irreducible differences between physical and digital objects.

Physical objects obey Newton's three laws. They have mass and hence inertia, so it takes force to move them and force to stop them. They have resistive capacity so any force applied to them encounters an equal and opposite force. And they accelerate or decelerate in conformity to $f=ma$. Physical solids have four other primary properties owing to their materiality: shape, size, position, and solidity; the properties of colour, and texture, etc. are more secondary.

None of these attributes are true of digital objects. Digital objects are computationally generated elements that can be displayed in wall projections, on screens, or be used in computer programs. They exercise no resistance, they have no mass, inertia, no intrinsic colour, shape, size, or solidity. They can be moved without concern for $f=ma$;

they exercise no opposite force when a user tries to move them; and because they have no mass it doesn't matter how large they are when a user tries to accelerate them. The physical dynamics they do seem to possess all have to do with input devices – the mouse or pointer being used to move them – not with the properties of the objects themselves.

Moreover, since the underlying notion of a digital object is that it is generated on demand, it lives in a kind of nether world of potentiality until called. There is no sense in which (as algorithm) it has a spatial location, although, if graphically realised, its display will always have a location, colour, shape and size. Even as a display object there is nothing privileged about its shape, size or colour at any moment, since on a larger screen or projected at a different angle it will have a dissimilar size, shape or corolla. All these changes have no more effect on its true identity than changing clothes or getting a suntan affects our own identity.

As for causation – and hence objecthood in the everyday sense – this is achieved only through mediated interaction. In a good interface, every effort is made to provide users with the feeling that they can directly manipulate the objects they see. The truth, though, is that users never actually interact with any digital object directly. Pointers or mice must be configured to pick up objects, to move them, to act on a screen window rather than the contents of the window. Users learn by experience how to manipulate input devices to achieve the changes they want in objects. It may feel to users as if they are interacting with those screen objects, but that is the result of their adaptation to the virtual. What they have really learned is how to make digital objects behave the way they want by mastering the use of input devices. An evil programmer, in one stroke, could subvert their actions and cause digital objects to behave in completely unpredictable ways. Because in the digital world there are no intrinsic physical constraints, causation is simulated.

This odd situation which digital technology creates is nicely portrayed by the way modern aircraft rely on simulations of the feel of flying to improve the control of pilots. Apparently, jets fly faster if their centre of mass is moved closer to the plane's nose, thereby changing the relative position of the centre of mass with respect to the centre of lift. The trouble is that in moving the centre of mass forward there is an increased chance that the plane will tip into a nose dive. To keep the plane flying on this knife edge the speed and sensitivity of adjustments is so great that pilots can no longer use mechanical means to control their planes. To assure fast enough response such jets now rely on digital networks to relay a simulated feel to the pilots. When a pilot pulls up on his steering wheel the computers inside the plane simulate the resistance of the ailerons delivering to the pilot the haptic information he or she needs to know what they are doing. Small computer adjustments augment and speed up these pilot reactions. To the pilot this force feedback is an integral part of the way he or she flies the plane. But, of course, there is no true resistance in the steering wheel. Pulling harder on the wheel is just a way of sending the number seven to the wing actuators instead of the number five. Computation is so irremediably built into planes that pilots could be in simulators.

Like it or not, we expect no less a marriage of computation with the materials, furniture and artefacts in buildings. What is real and what is augmented, what is material and what is virtual, will become less a question of causation and more a question of 'is this helping me work?' Design must cover both aspects. Given the pace of digital innovation there is little question, though, that it is the digital side that has become the motive for change.

It is at this point that we must free our thinking. Digital objects can be delivered at lightning speed almost anywhere. They can be duplicated essentially for no cost and they need next to no room to be stored. Indeed, there is no longer any real sense to the question of where they are stored since 'the network is everywhere' and systems are becoming increasingly distributed. One file, a thousand copies anywhere, just in time, and just the way you want it (big font, small font, annotated, read only, modifiable). Moreover, the very rooms in which these digital elements are embedded soon will contain thousands of sensors and small actuators. How shall we exploit this new-found interactivity? This collapse of space, paper and information? I shall briefly mention three ways: telepresence, three dimensional intranets, intelligent furniture.

Telepresence

Telepresence refers to being in one place and having the full experience of being present in a different place. Someone experiencing telepresence is able to perceive and act as though at the remote site. Although full telepresence remains several years in the future it is easy to project broadband video-teleconferencing on walls, panels and work surfaces.

With telepresence the notion of *place* weakens dramatically. Walls become telepresent portals. Desk work surfaces can unite with other desk work surfaces to support remote collaboration. We can keep a robust telepresence in two or more offices at once to allow for opportunistic encounters. When telepresence is coupled with rich media capture of activity, time, as well as space begins to lose its structuring power on how we must *prepare* to do what we do where we must.

Three dimensional intranets

A second major change digital everything makes possible is three dimensional intranets and anytime, everyplace personal information networks. At present, corporate and personal intranets are accessed through computers. As these move out of laptop and desktop boxes to walls, windows, furniture, and ceilings, and as our means of interacting moves beyond the confines of mice, pointers, and tablets, we will interact using speech, gesture, and everyday objects. Sensors will be everywhere and processors will be devoted to making sense of this new level of context awareness. Perhaps this will take another five years in laboratories, perhaps it will take fifteen.

Even now, though, projection of information onto walls and work surfaces allows users to interact standing up and by touching and pointing, often on a writable surface such as a whiteboard. The extra space this provides for organising information means that we now can massively increase the number of entry points to information we have in a room. Intranets and digital libraries can take to the walls. But with that increase in power comes the need to design the interface to all this information with the presentation parameters of walls and ceilings in mind. Information architecture will have to take account of physical architecture.

Moreover, with such extra space available we can begin to take advantage of the changing properties of documents and exhibitions. For example, already we cannot simply 'print out' meteorological data, which may occupy gigabytes on disc. Data and programs have become integrated in complex applications, which let us interactively visualise the data. Sometimes we need to stand up and move around to deal with these large visualisations. The same holds true when we want to interact more completely with resources in learning environments, games and exhibitions. We need the extra space of three dimensions. The issues involved in providing services in such distributed, multi-layered environments are poorly understood. But they are changing the way we work.

Intelligent furniture

The third change concerns our furniture and the effects on office activity it will have once it too is linked into the architectural network. Already there is a trend toward increased reconfigurability in office landscapes, with furniture playing an important role. But so far, excepting a few laboratories, furniture has not been enhanced digitally to track our actions, adapt to requirements by digital extensions, and support information presentation. Furniture makers have concerned themselves with physical ergonomics not cognitive ergonomics and cognitive amplification.

Again, though, as sensors, projectors, and actuators begin to saturate our furniture our manner of relating will change. If a telepresent partner calls a colleague to join us to discuss diagrams, our desk might extend visually onto the wall or actuators might extend a further projectable surface to let us spread things out. Information at hand can be moved by wand, mouse or gesture to any part of the three dimensional information space we will inhabit. Once our orientation focuses squarely on supporting information and interaction with information in all its forms, the way we think about the specification of a room or office will change enormously. Furniture and walls are part of an interactive system.

Workspace Enclosure

The glimpse of the near future just gestured at may seem more like science fiction than hard-headed empiricism. Buildings still need to support and enclose. Yet if I am right, the requirements of support and enclosure no longer need dominate our experience of buildings. The digital can take us beyond the tyranny of materials.

As an example of how differences between digital and physical objects matter to architecture it is illuminating to examine the nine variables that Christopher Alexander et al presented as influencing a person's sense of enclosure in their workspace.[4]

Based on an empirical evaluation of office dwellers Alexander *et al* drew a system of implications for enclosure design. These implications were meant to be universals of design – invariant consequences of the psychology of users. Almost every consequence is made either otiose in a digital office or must be reconceptualised.

According to Alexander et al, 'People cannot work effectively if their workspace is too enclosed or too exposed. A good workspace strikes the balance'.[5] This basic idea plays out in nine recommendations borne out by an analysis of what 17 men and women thought were their best and worst workspaces. Let us look at these now and explore how they are changed by the possibilities digital enhancement creates.

(1) Wall immediately behind. 'You feel more comfortable in a workspace if there is a wall behind you. (If your back is exposed you feel vulnerable – you can never tell if someone is looking at you or, or if someone is coming at you from behind.)'.[6]

In a digitally enhanced office the sense of where a wall is can be strongly influenced by trompe l'oeil, graphical pattern, visual movement and transparency. Too little is known empirically about these influences. Nonetheless, a temporary feature of walls is that their decorations and appointments do not move, nor does their transparency vary. In his survey, Alexander relied on this assumption that walls are flat surfaces, uniformly coloured, serving as material shields, protecting a person from approach and prying eyes. No data was offered to show whether colour, texture, acoustic properties, trompe l'oeil, or other treatments alter the distance that counts as 'immediately behind'. And no effort was made to distinguish the psychological distance a person feels from an opaque wall from the psychological distance they feel from a transparent wall. Digital projection will change all this.

(2) Wall immediately beside. 'You feel more comfortable in a workspace if there is a wall [immediately] to one side. (If your workspace is open in front and on both sides you feel too exposed.)'.

Another question Alexander did not address was which type of wall or shield protects us most from a feeling of exposure: sound walls, visually opaque walls, or physical approach walls? In normal physical environments there are few ways to dissociate shielding of sound, sight and physical approach. Aside from playing with the visual and acoustic properties of glass, which can dissociate light from sound and approach, there are few ways of separating the remaining factors – approach and sound. In a digital world where such dissociation is easier the import of each factor needs to be explained. Does knowledge of video access, whether or not we see who can see us, vitiate the secure feeling a wall provides? What effect does the ability to turn on and off transparency to sound, sight and movement have on our sense of exposure? What is the effect of the ability to make a wall one way transparent, letting us see, hear or step out while denying that service to others? Digital walls offer greater precision in selectively blocking out discomforting factors.

(3) Amount of open space in front of you. 'There should be no blank wall closer than 8 ft in front of you. (As you work you want to occasionally look up and rest your eyes by focusing them on something farther away than a desk.)'

Digitality alters the meaning of open space. For Alexander, the space that matters is above all physical space. Despite the rhetoric about the physiology and psychology of gazing it all comes down to physical distance rather than perceived distance. Yet if we focus on psycho-physiological reasons, then what matters should be a person's feeling of openness regardless of whether that feeling is created by wall projections or physical reality. Since, any estimation of workspace comfort must surely depend on the regular tasks performed in that space, changing a person's activity space should effect their space needs. If one's desktop is no longer one's primary work surface because one works with digital paper on the walls, whiteboards, corkboards, calendars and the three dimensional intranet, a new set of activities become relevant in determining how much space is felt to be necessary. Comfort and activity cannot be divorced.

(4) Area of the workspace. 'Workspaces where you spend most of the day should be at least 60 square feet in area. (If your workspace

is any smaller than 60 square feet you feel cramped and claustrophobic.)'

The same concern about perceived versus real applies when we explore the amount of workspace area a person needs to feel comfortable. Are we concerned with the true area or the psychological feeling of area? A person living inside an office, whether with wall projections or not, certainly knows the rough distance to the walls. Wall projections change perceived distance when the projection system is on, but they neither change the space available for movement nor the psychological space when they are not active. Assuming full scale wall illusions – rich three-dimensional life-size images of the world outside right now – it is an as yet unanswered empirical what physical area is necessary to eliminate claustrophobia. This is especially true if a person is not surrounded by four walls.

(5) Enclosure around the immediate workspace. 'Each workspace should be 50 to 75 percent enclosed by walls or windows. (We guess that enclosure by windows creates about half the feeling of enclosure that solid walls have, so that a workspace which is surrounded half by a wall and half by a window is considered to have 75 percent enclosure ...)'

The motive for enclosure, presumably, is privacy and sense of security. As such, any digital walls that alter the sense of privacy or security may have an effect on the design of enclosures. Selective blocking of sound will increase a sense of privacy and/or security. Moreover, such rules of thumb as Alexander suggests must surely vary with the frequency and type of traffic to be expected. What are the effects of working in an office with remoter collaborators rather than present ones. If video, telepresence, projection and portable furniture are at hand the psychological dynamics of enclosure will be different.

(6) View to the outside. 'Every workspace should have a view to the outside. (If you do not have a view to the outside, you feel too enclosed and oppressed by the building, even if you are working in a large open office. ...)'

Will a changing view of a street or a forest satisfy the desire for a window overlooking life? Trompe l'oeil that are static may have limited value in alleviating a sense of enclosure. But little is known about the psychological sense of openness delivered by a photo-realistic three-dimensional image of an active scene (particularly if it is a real time image of a real scene, such as our backyard, or the street outside, or whatever it is that is seen in the picture window across the hall.

7. No other person should work closer than 8 feet to your workspace.

> '(You should be able to hold a conversation either on the phone or in person with someone without feeling as though someone else can hear every word you are saying ...)'

Once sound walls are commercially available what grounds will support this requirement? Given rule eight, that it is desirable to have other people working nearby the question becomes, once more, one of privacy.

(8) Other people nearby. 'It is uncomfortable if you are not aware of at least two other persons while you work. On the other hand, you do not want to be aware of more than eight people. (If you are aware of more than eight people, you lose a sense of where you are in the whole organisation. .)'

With telepresence it is easy to have remote colleagues work nearby. It is also possible to digitally control how loud they sound, how near they seem to be, and so on. It may still be true that comfort in an enclosure or office requires the presence of two or more colleagues, but how this constrains physical design is suddenly less clear.

(9) Hear only similar workplace noise to your own. 'Your workplace should be sufficiently enclosed to cut out noises which are of a different kind from the ones you make. There is some evidence that one can concentrate on a task better if people around him are doing the same thing, not something else.'

Digitisation of the workplace means that computational devices may pre-process all ambient sensory elements such as sounds, sights and smells. As long as there is a way of shielding and capturing an original stimulus, it can in principle be transformed and then regenerated in a new form, much as now happens in cockpits, so someone on the receiving end cannot tell what the original stimulus is. In this case, sound generated by nearby workers might be blocked near them and modified and redirected to the appropriate person. Different neighbours could all hear different sounds coming from the same enclosure.

Given the universality which Alexander et al believed their pattern language exercised over architectural design it is of some significance, I believe, that not one of their variables carries the physical consequence it was once thought to imply. The impact of the digital on architectural design is evident in the way the normal rules for deriving formal constraint are overturned.

Conclusions

How architects resolve the material with the virtual will shape our experience of buildings and enclosed space for the foreseeable future. The problem is not theirs alone. Fundamental concepts are in transition and it is the job of cognitive scientists, new media analysts, and computational scientists to understand the implications for human interactive behaviour. I have argued that the concepts of psychological space and place are changing; that the concept of enclosure is changing; that our idea of what is furniture and what is architecture are changing. Architects, until now, have been able to concern themselves almost exclusively with the experience of spatial structure, functional effectiveness and aesthetic feel. But henceforth they will have to be concerned with information architecture too, for how people navigate through information space will also be linked to how they move around their offices and use the shapes therein. Buildings will forever serve to support and house, but as the digital permeates our physical world the way we experience architected structures will have less to do with the material nature of those structures. The wired world is not the same old world.

Acknowledgements I thank Ted Krueger for thoughtful discussions and his encouragement to write this article. The author also gratefully acknowledges support for research on this and related topics by the National Science Foundation (KDI grant IIS 9873156) and the Office of Naval Research.

Notes

1 For an account of some of these tricks and techniques see D. Kirsh, 'The Intelligent Use of Space', *Artificial Intelligence*, 73 (1995), pp. 31-68; D. Kirsh, 'Complementary Strategies: Why we use our hands when we think', *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society* (Hillsdale, NJ: Lawrence Erlbaum, 1995); D. Kirsh, 'Adapting the environment instead of oneself', *Adaptive Behavior*, 4, no. 3/4 (1996), pp. 415-452.

2 See D. Kirsh, 'A Few Thoughts on Cognitive Overload', *Intellectica*, 30 (2000), pp. 19-51. Also available at http://icl-server.ucsd.edu/~kirsh/Articles/Overload/published.html (18 April 2001).

3 William J. Mitchell, 'Architectonics: the poetics of virtuality' in *The Virtual Dimension: Architecture, Representation and Crash Culture*, ed. John Beckman (New York: Princeton Architectural Press, 1998), p. 213.

4 Christopher Alexander et al, *A Pattern Language: Towns, Buildings, Construction* (New York: Oxford University Press, 1977), pp. 847-852.

5 Alexander et al, p. 847.

6 For the rules quoted here see Alexander *et al*, pp. 848-850.