

Informatics, Architecture and Language

Matthew Chalmers

Dept. of Computing Science, University of Glasgow

matthew@dcs.gla.ac.uk

Abstract

Two complementary schools of thought exist with regard to the basic underlying assumptions and philosophies that guide our research in information navigation and access. As with all of HCI, and indeed most of informatics, we can place theories and design practices based in objectivity and mathematics at one end of a spectrum, and those emphasising subjectivity and language at the other. The first school of thought sees itself as part of traditional computer science, rooted in models that encompass the individual variations of users and that are often derived from experimentation and observation in controlled conditions. Mainstream information retrieval, cognitive psychology and task analysis exemplify such a philosophy. Complementary views are held by those who hold the sociological and the semiological as primary, and consider that objective categorical models are insufficient to model the complexity of human activity and ultimately of limited utility in guiding system design and development. Collaborative filtering, ecological psychology and ethnography are examples here. The techniques and systems presented in this book do not all lie towards one end of this spectrum, but instead show a variety of choices and emphases. This chapter, however, focuses on theory firmly towards the subjective and linguistic end of the spectrum: tools to let us place, compare and design techniques and systems. Such theory is noticeable by its absence in the majority of literature in this burgeoning research area. Here we try to redress the balance, aiming to build up a more abstract and general view of our work.

At the most applied level, this chapter deals with one approach to social information navigation systems, the path model [1], and describes its origin in an analogy with a theory of urban form: Hillier's Space Syntax [2]. More generally, we relate the use of and movement through information to use and movement in urban space. While architecture has already affected informatics in a number of areas, for example in the pattern languages of Alexander, here we use architecture as a stepping stone between linguistics and informatics. Through these links we wish to reinforce the view that all three are instances or subfields of semiology. In so doing, we aim to make more visible the range of assumptions and models that underlie all interactive information systems. We are often unaware of the models of knowledge and information that we build on, and the possible alternatives. Here we aim to make clearer some of those buried layers—the 'archaeology of knowledge' [3] that determines many of the strengths and weaknesses of any system for information navigation.

What underlies HCI?

The opposing schools of thought mentioned above were exemplified in a recent exchange in the pages of SIGCHI Bulletin, discussing appropriate metaphors for navigation and organisation of files on the desktop. The most recent contribution (at the time of writing) was [4]. At issue was the importance of location-based searching over logical retrieval. Fertig, Freeman and Gelernter put forward what might be considered a traditional computer science viewpoint, suggesting that continuing research, gradually adapting and extending the analysis approaches rooted in 1960s 'document retrieval', would eventually succeed. It would create better automatic tools for indexing the content of large databases and collections, made up of varied data types such as textual documents, images, sound files and so forth. It would, therefore, allow users to gain a full return on their investment in storing large amounts of information. This despite the fact that textual document retrieval has, even by its own measures of performance, not dramatically improved since its inception, and retrieval of image, video and audio content is still rudimentary. Nardi and Barreau suggested that such improvements would be valuable but would not address what they consider to be the paramount information management problem, the volume and heterogeneity of ephemeral information that comes and goes in everyday work. They see the traditional computer science approach as giving insufficient attention to the intricate details of how information on an individual's desktop is interwoven with the rest of that person's working environment of people, institutions and cultures. Such concerns are largely absent from the practice and theory of data retrieval. At the core of Nardi & Barreau's objections is a concern that "the alternatives offered by many developers of personal information management systems seem to view documents in the work space as a collection that can be easily characterized, ordered and retrieved based upon common characteristics, or based upon full text retrieval. These approaches ignore the complexity and variety of information in personal electronic environments. [...] Schemes that automatically characterize information may not provide enough flexibility to consider the richness of these environments, and schemes that allow characterization for visual retrieval may not easily accommodate all of the desired dimensions."

The same issues arise in the World Wide Web, where complexity and heterogeneity of information representation are increasingly problematic. The mix of data types restricts the coverage of traditional indexing techniques and limits the consequent power of search engines. As Tim Berners-Lee pointed out with regard to search engines for web data [5], "they are notorious in their ineffectiveness. [...] A web indexer has to read a page of hypertext and try to deduce the sorts of questions for which the page might provide the answer." Images, numerical data, audio, programs and applets: the variety of information is increasing along with the volume. Attempts are being made by various researchers to solve these problems, usually by adding some form of metadata. Metadata is data that represents the meaning or content of other data. It is often expressed in a formal vocabulary, and intended to allow programs to uniformly compare data originally from a variety of authors and sources.

Metadata may be handled automatically and objectively by indexers and search engines, but is often created manually. If metadata is manually created, it is open to be written, read and interpreted as each person sees fit. For example, web page authors or site managers may enter informally structured textual data such as captions or tags into a formally manipulated metadata scheme. For example, the PICS format proposed as a standard for web metadata [6] involves 'labels' that describe a web page or site e.g. suitability for children, whether it is commercial or not, and even, as Resnick and Miller suggest, 'coolness.' They also suggest that there might be many rating services, each of which could choose its own rating vocabulary. This opens the door to completely subjective metadata, however. For example, how objectively and uniformly with regard to its competitors will a large corporation describe its own web pages and products? How 'cool' will each home page be according to its author?

Also, even if the metadata describing a site is, to give an optimistic figure, one thousandth the size of its referent, the vast size of the web would lead to gigabytes of metadata overall. Ironically, we would then need 'metametadata' tools to find good rating services. We have not solved the problem, but deferred or even exacerbated it by adding to the indirection and complexity involved in acquiring useful information.

In other words, we only delay the problem of matching available information to users' interests and activities by stepping up to a metalevel. Since metadata open to individuals' use slips down to be just more data, we are back where we started.

Nardi and Barreau are amongst those involved in contemporary HCI theory, which is critical of modelling the mind (and hence the user) as an algorithmic processor—an approach that was until recently considered the firmest foundation upon which to build interactive information systems. Typical research involved short term controlled experiments in a laboratory-like setting, with experimental subjects introduced to new tools and techniques to be used in isolation from the other tools familiar from their everyday work, and away from their colleagues and workplace. Modern HCI theory criticises this as excessively reductionist, unrealistically examining use of tools that would normally be interleaved with other tools, as part of long term work within a community of use. Consequently they fail to take account of the complexity and situatedness of interaction, and this approach has not offered the hoped-for practical benefit to designers. Work such as Activity Theory [7] takes a more realistic view of the subjectivity, dynamics and social context of individual action. It broadens consideration from just the actor and the tool being designed to the other tools used, the intended outcome (at various levels of abstraction), and the community within which activity takes place. Activity theory has, however, been better for analysis and criticism than for driving system design, which we suggest is related to its greater concentration on activities than on artefacts i.e. on work's goals and actions rather than on the information and tools that represent and mediate work. Here we shift the balance the opposite way, focusing on information representation, categorisation and interpretation, and so becoming more directly linked to system design.

The claim that algorithms, themselves founded in mathematical logic and formal languages, form the ultimate foundation of informatics has also been denied by formal informatics itself. Wegner recently published a proof that interactive computing is an inherently more powerful computational paradigm than purely algorithmic computing [8]. He shows that the complexity and unpredictability of human interaction offers greater expressiveness and analytic power than formal algorithms. Interaction involves more than algorithms can express i.e. interaction is not reducible to algorithms, and interaction-based approaches to system design are more powerful than purely algorithmic ones. Computer science's traditional demand that one should be able to use formal languages to ground the behaviour of programs is therefore seen as inhibiting the expressiveness and power of the programs it builds. Wegner uses formal informatics to show its own limits, echoing the history of mathematics (and indeed physics) where the dream of reducing the world to a pure, clean and objective mathematical model has been shown to be an illusion.

We suggest that a belief or assumption that mathematical logic and objectivity form the ultimate basis of informatics is a naïve, reductionist stance. All too common even amongst computer science professionals, this is not based on the mathematicians' view of mathematics but on 'folk mathematics' i.e. a common-sense framework for understanding mathematics, widespread amongst the wider population. This is analogous to 'folk physics' that has a rough notion of inertia and Newtonian mechanics, and only a vague grasp of relativistic

physics and its successors, and 'folk psychology' the pre-scientific framework commonly used in comprehension and discussion of human behaviour [9].

In the next section, we attempt to look beyond 'folk informatics' and see what we can learn from mathematics' reconsideration of its foundations this century. In effect we ask the question: even if informatics is based on mathematics, mathematics is not firmly grounded in objective logic, so what is mathematics based on?

From Mathematics to Language

Some programs represent models and theories which are not intended to relate to the physical universe or our lives in it, for example in some branches of abstract mathematical theory. More often, however, information representations model aspects of our lives. We reduce and select from the infinite complexity of the world in order to gain the ability to store, index, manipulate and retrieve. This involves finite models, expressed in mathematically based formal languages, that represent properties and relations between real world objects. It has long been recognised in mathematics that such languages comprise of three types of structure: algebraic structures such as sets, structures of order such as lists, and topological structures such as graphs.

Since we use such mathematical schemes as foundations for building information representations, we should be aware of how solid and objective they are. This has of course been a concern within mathematics itself, for example in Cantor's demonstration of a paradox in set theory. The set of all sets is shown to be indeterminable because the 'metaset' of all sets is itself a set, and therefore should be a member of itself along with many other sets, but this is inconsistent with our definition of what a set is. The metaset slips down from metadata to data. Similarly we are familiar with the paradox in statements such as "This sentence is false."

To those determined to find objective foundations for mathematics, the persistent lack of a non-perceptual basis for such an apparently simple and familiar mathematical framework as Euclidean geometry was a particular concern [10]. This framework consists of five axioms, but no-one had been able to show that the axiom of parallels was founded in more than (human) claims of apparent self-evidence. This, to use Everitt and Fisher's phrasing of it, postulates that given a straight line and a point not on that line, it is possible to draw one and only one through the point parallel to the line. Also, Euclidean geometry is not in accord with modern physics' observations. On a cosmological scale, parallel lines do eventually meet. Even logic was under pressure, as quantum physics began to cast doubt on the law of the excluded middle whereby, given a logical statement p , only one of p and *not* p can be true. Hilbert led the formalist or foundationalist approach, trying to render mathematics genuinely consistent and independent of perception. As cited in [11], this meant that "the solid foundation of mathematics is in the consistency of its formal system: mathematics does not have to be 'true' as long as it is 'consistent,' and as long as this is the case, there is no need for further foundation."

Abandoning the claim to truth was not sufficient, however. The major blows to Hilbert's approach came from two directions. One was from within, in the form of Gödel's incompleteness theorem. This, in a self-referential manner related to that of Cantor, set up a paradox like "This sentence is false" wherein meta-mathematics, understood as a class, slips down into the formal system as a member of itself. Gödel proved that for any system of mathematics of significant expressive power, it was always possible to set up such a paradox or inconsistency. As later echoed by Wegner's proof, he therefore showed that the consistency of a formal language was only obtainable at the price of limited expressiveness. Gödel thus helped release mathematics from the illusion that it could consistently represent the world.

The other blow was from without, from Wittgenstein, who reinforced the view that mathematics was part of human history, and not an abstract ideal independent of it. He rejected the notion that mathematics' formal system can be solidly deduced from axioms. Proof is just another 'language game' [12], involving our invention of rules, systems and notations whose truth, as with all our natural language, is determined by our own social use rather than from axiomatic deduction.

Gödel demonstrated the problem of inconsistency in the mathematical structures we use for information representation, while Wittgenstein replaced the axiomatic basis of their truth and meaning with a social, linguistic basis. Together they were instrumental in discrediting positivism, the previously dominant paradigm whereby a symbol was a 'positive term' i.e. an objective absolute, based purely on a logical process of naming a thing in the real world. This concept of naming and reference, that connected mathematical symbols with things in the world, had been the core of positivists' concept of 'truth.' This concept was seen to be no longer justifiable. As a result, mathematics no longer claims to offer a means to consistently, truthfully and absolutely represent the world. We can choose to use it as a tool, pragmatically accepting its limitations and historical biases, but even then we should be aware that mathematics has a history of paradigm shifts and scientific revolutions, and we have no reason to believe that this historical process has ended.

Underneath mathematics, we find language: systems of symbols, with subjectivity tempered by socially constructed practices of proof and experiment. The contemporary view of language is that we cannot dig further: there is 'no exit from language' partly because all our consideration and modelling of the world is ultimately understood and used by means of perception and language. If natural language is formalised or abstractly represented, then when we interact with that representation, use it in the world, in our human interpretation and

activity, we necessarily involve subjective perception. We may try to step up to a metalevel by means of formalism and abstraction, as with the web metadata example earlier in this paper, but with human activity we slip back down to language. Within a controlled environment such as the computer we can manipulate such a 'meta-representation' as part of a more formal, finitely defined system. In interaction, however, the informal, subjective and infinite reassert themselves.

What, then, can we take from contemporary linguistics and semiology, that will help in our theory-building? Only a few decades before Gödel and Wittgenstein revolutionised mathematics' core, a similar paradigm shift had taken place within linguistics. Wittgenstein's language games have often been identified with linguistics' new paradigm: structuralism. Structuralism [13] combined linguistics and semiology, and displaced positivism in those fields.

Saussure's view was that, unlike positivism, naming is a relative or differential process, in that the elements of a language at any given time form a structure where any element only has meaning because of its relations and differences with other elements. Again we see a contrast with a positivist view of a one-to-one relationship of naming a unique, absolute, ideal thing in the world. If naming and language was based on such absolutes, then observed temporal variations of meaning could not happen e.g. 'cattle' refers to bovine animals in this century but meant 'all kinds of personal property' in the Middle Ages. Also, how could one language use only one word where another language uses two or more e.g. we use 'river' to cover both 'rivière' and 'fleuve' in French. (A 'fleuve' is a large river, which may flow into the sea but the word does not refer to just an estuary or firth.)

Japanese has several words for the number 'one', used to suit the type of thing being enumerated. While we may use a word to signify a thing in the world, it does not refer to one absolute and abstract thing that each other language also has exactly one word for. Our meaning is derived from our use of the word's similarities and differences to the other words of our language. A word means what we use it to mean, or, to quote Wittgenstein, "the meaning of a word is its use in the language."

In Saussure's theory of natural language, the medium can be anything, including speech, written text and physical motion of the body e.g. sign language and dance. We can choose to use anything and any combination of media to communicate. It is this interpretative choice or reaction that creates significance, and so any action in any medium can be taken as significant, and hence as a symbol.

Representation in HCI and Information Access

Now we can see a fundamental limitation in the information navigation and access approaches that rely solely on the content of each information object in isolation, such the words inside a document, and ignore objects' use in human activity and objects' inter-relatedness. To assume that the words contained inside a document faithfully and fully describe the meaning of a document, irrespective of its use in language, is a naïvely positivist approach. Traditional content-based approaches can be seen as emphasising and operating on symbols and attributes which are contextually independent e.g. no matter who has a document and what activity they are involved in, the same set of words are contained inside the document. Of course, this specificity affords highly useful techniques such as indexing of contained words to allow quick searching, but we rely on the assumption that the context of use and the person involved are not significant. This is true when one wishes to find all documents that contain a given word, but false if one wishes to find all documents that one's colleagues find useful, or that conform to one's interpretation of a given topic. The assumption in itself is neither good nor bad; it is lack of awareness of it and its consequences that causes us problems. We should realise what assumptions our information representations are built on, and hence what they afford and what they inhibit.

That perception of a structure or representation is bound up with the perception of use became familiar to many within HCI via the ecological theory of perception of J.J. Gibson [14], later popularised by authors such as Don Norman. Gibson stresses the complementarity of perceiver and environment. The values and meanings of things in the environment arise from the perception of what those things provide or offer as potential actions or uses to the perceiver—in Gibson's terms, their affordances—and not by universally naming and categorising absolute or objective properties. He emphasises the way that a theory of meaning must avoid "the philosophical muddle of assuming fixed classes of objects, each defined by its common features and then given a name. As Ludwig Wittgenstein knew, you cannot specify the necessary and sufficient features of the class of things to which a name is given. [...] You do not have to classify and label things in order to perceive what they afford."

While many in HCI and information access have read Gibson or Norman, our field has not taken full account of the way that information representation schemes have affordances: they are also objects with characteristic strengths and weaknesses to choose in accordance with our uses, interests and abilities. Revealing the underlying models of knowledge and interpretation gives a common framework and vocabulary for comparing and analysing such schemes, as is the subject of a forthcoming paper [15].

Within HCI, Lucy Suchman [16] has been instrumental in establishing the importance of situated action: how particular concrete circumstances have a strong influence on behaviour, and how strict plans are often merely resources for more flexible, dynamic, contingent action i.e. more like maps than scripts [17]. Like Gibson, she generalises over objects in interfaces and objects in the physical world, treating them as elements of sign

systems, as linguistic expressions. “The significance of a linguistic expression on some actual occasion [...] lies in its relationship to circumstances that are presupposed or indicated by, but not actually captured in, the expression itself. Language takes its significance from the embedding world, even while it transforms the world into something that can be thought of and talked about.” [16, p. 58] Again we see the need to look beyond the content of the expression or object, towards the co-dependence and co-evolution of human behaviour and information structure, and the influence of context and situation of use not usually represented in our information systems.

Space Syntax

Having put forward the argument that information use is semiological, we can draw upon another fields or disciplines that have been accepted as semiological for some time. We focus on architecture here, amongst many semiological fields, as the path model of information access, discussed in a later section, was based by analogy with what we characterise as a structuralist theory of urban structure, Hillier’s ‘space syntax’ theory [2]. The path model treats information objects like spatial forms, histories of information use like individuals’ paths through the city, and language as the city. The notion of ‘language as city’ was also at the centre of Wittgenstein’s language games [12, p.8].

Around the start of the 20th century, around the same time as linguistics’ and mathematics’ revolutions, architecture took its most decisive steps on from the notion that buildings and cities are purely functional objects, exempt from significance to and influence from cultural and symbolic concerns [18]. The shaping and use of architectural form was then understood as being equally as semiological as the shaping and use of letter, word and document form in written texts. Word and written symbol usage corresponds to the motion and occupation of rooms and spaces in being a temporally ordered sequence of significant action.

In this section we outline space syntax, prior to a section using it to briefly characterise several traditional information access approaches and a section which lays out the path model in detail. Space syntax offers a view of city structure and development based on individuals’ movement. Moving through the city may be due to a variety of motives—plans, contingencies and so on—but Hillier and his colleagues have found that patterns of use and meaning for people are correlated with statistical consistencies in peoples’ paths. Such consistencies are correlated not so much with the nature of individual building forms and functions, but instead with the patterns of connectivity and visibility that make up the urban configuration.

Non-discursivity

As with other forms of language, and as activity theory and situated action emphasise, we often do not know how to talk about why we act as we do—why we read the city in the ways we do, or why we collectively tend to favour certain paths and routes and not others, or why the city configuration works or fails. This non-discursivity of configuration is an important issue in Hillier’s theory and in HCI theory, in Gibson and Suchman. Just as we are able to speak understandably and to understand what we hear without being able to express—to make discursive—why our language is as it is, movement in the configuration of the city is generally non-discursive. Space syntax is an attempt to find consistencies of use that we can talk about as theorists and designers i.e. to bring non-discursive aspects of architecture into the discursive.

Space provides the potential for paths of movement and view, infinite in the case of void, empty space but increasingly constrained as we introduce forms such as buildings. The city is at once a record of the functional processes that historically created it, and at the same time the strongest constraint on future development. Specific activities such as finding a particular building and assessing the most profitable location to place a specialist shop relate to the spatial form of the city through ‘general functionality’: the ways that we as individuals find a system of spaces intelligible, and the ways we move around in it.

Integration

The ‘integration’ of a configuration is a metric based on connectivity of spaces, and is a fundamental part of the vocabulary of the analytic side of the theory as a whole. We represent an urban configuration, such as the rooms of a building or the streets of a city, by a graph—an abstract representation of the configuration that simplifies but affords some useful analysis. Each node represents a separate subspace such as a room. The links between the graph nodes correspond to the connectivity of these spaces. Fig. 1 shows an example. The integration of a node is the maximum of the lengths of the most direct paths one could take to each other node i.e. the maximum shortest path length to any other node. We can combine the effects of the nodes that make up a graph by summing these integrations. There is an extra normalisation factor to take account of the size of the graph, but essentially we obtain a figure that describes the graph, and hence the urban configuration. This figure, the integration of the graph, shows (amongst other possible interpretations) how well connected the space is i.e. the degree to which moving around the space requires long routes that often connect you to only a small proportion of the nodes.




Figure 1. A floor plan with its connectivity graph. Each node in the graph represents one 'room' and each arc between nodes represents a door connecting two rooms. Each node is labelled with its integration, showing the maximum path length to other nodes. The node of lowest integration may not be as good a movement node as the well-connected central node that bridges two rings. The top left 'dead end' node is a good occupation node

Visibility and Intelligibility

Visibility also plays its part. Although connectivity may make paths feasible, if the consequences of movement are not apparent—one cannot see where or to what a path may lead—then we diminish the confidence with which we make choices of movement. As we look from any point, we are aware of how our lines of visibility extend outwards and connect to spaces we could move into. An 'intelligible' configuration is one in which what you see from the nodes of the system is a good guide to the global pattern of depth. More formally, Hillier treated intelligibility as the correlation between integration and the connectivity of lines of visibility.

He suggests that minimising graph depth and integration support intelligibility, and are themselves best supported by combinations of occupation and movement spaces. Occupation means the use of space for activities that are at least partly and often largely static, such as conversing, reading, sleeping, cooking, and working at a laboratory bench. 'Dead end' spaces with a single link are essentially occupation-only spaces, links that collectively form rings tend to help the integration of spaces, and spaces with more than two links and connect at least two rings are the best movement spaces. Different types of space afford different patterns of occupation and movement according to their patterns of branches and rings. Occupation spaces afford detailed, localised activity, while movement spaces maximise the flexibility and efficiency of motion.

Movement through the city means that we encounter buildings and people along the way. No matter how planned such activity is, we gain a by-product in the encounters along our way. Configuration affects how well we can find our goals as well as come across the situations where unforeseen actions may take place e.g. in going to one shop in the city centre, we might pass another that prompts us to take a detour. Hillier suggests that maximised integration is the best way to gain the usefulness of this by-product, and leads to positive feedback between movement and development—the "multiplier effects which are the root source of the life of cities." To use Gibson's term, maximising affordances for encounter is what makes cities good.

Architectural space limits the patterns of co-presence amongst the individuals living in and passing through an area, and therefore co-awareness and affordances for interaction. In this way the spatial becomes involved with the social. As a social fact and a social resource, spatial configuration defines the 'virtual community' of an area: the pattern of natural co-presence brought about through the influence of spatial design on movement and other related aspects of space use.

Configuration does not always incrementally co-evolve with the varied movements of everyday life and work in a 'natural' or 'passive' way. Cities also show structures that are the result of active planning and shaping on a large scale. This may be for economic reasons, e.g. a city constructing a train line to encourage public access to its centre, or it may be for political or ritualistic reasons: 'to make a statement'. Examples are the dominant central axis of Brasilia, and the streets and spaces that focus attention on the centres of US government in Washington. Grand axes are extreme cases of how, at all scales, we can use paths and configurations as symbols and instruments.

Throughout space syntax, the function, categorisation, or content of the buildings involved are deliberately de-emphasised, and indeed it is Hillier's thesis that intelligibility and movement dominate them because all other aspects of function pass through them, and influence the urban form through them. These are strong claims, and yet Hillier presents many predictive analyses, backed up by a good match with observed behaviour in existing urban forms, without recourse to content analysis. He does not need it for his theory to work, and even apparently intangible phenomena such as house prices, burglary rates and locations of teenage and drunkard 'hangout' are quite accurately predicted by this very powerful theory which is now beginning to have wider application [19].

Information structures seen as architectural plans

Part of the strength of Hillier's presentation of his theory was the analysis of example configurations of streets in cities such as London, and his involvement in urban regeneration projects such as King's Cross. We now use

space syntax to look at simple prototypical cases of traditional information access techniques, as represented by simple connectivity graphs.

We employ several criteria in discussing each technique. Integration is usually first along with perceptibility. We must use the words ‘perceptible’ and ‘visible’ more loosely here, as we have no analogue of the objectively measurable lines of sight. Intelligibility, describing their informal correlation, is then mentioned. Another issue related to intelligibility is scale. This is not considered at length by Hillier, perhaps due to the size and flatness of cities generally restricting the variation in the number of perceptible buildings to within a few orders of magnitude. We discuss scale here, because in information environments the number of objects varies from a handful to billions. We also look for good occupation and movement objects, and a tendency towards the minimisation of graph depth.

Hypermedia

Let us reuse the architectural floor plan from above, and consider the graph of connectivity to represent bi-directional links in a simple hypermedia structure, as in Fig. 2. We do not include the object that provides access to the hypermedia structure, for example the desktop or command line, in this or the following figures.




Figure 2. A simple hypermedia graph. Textual nodes are marked in black, while nodes of other types are in white. Arcs are bi-directional links. Integration values are unchanged from the architectural floor plan of Fig. 1.

In simple or ‘pure’ hypertext, connectivity is the sole associative medium and the only perceptible nodes on a page are those which are linked to it. The scale of a hypermedia object is usually relatively small, being kept under the control of a small number of authors—often one. Integration is dependent on the author of the links, and will worsen in larger graphs unless connectivity is kept high. This not generally being the case, there is a tendency towards deep tangled networks of connections. Perceptibility tends to be poor and intelligibility tends to be low—it is easy to become “lost in hyperspace”. This also contributes to the poor general functionality of hypermedia; the authoring process does not tend to produce shallow graphs, or a mixture of objects serving as good general purpose circulation spaces and others as local ‘occupation’ spaces.

One of the reasons connectivity stays low is to take advantage of one of hypermedia’s strengths, namely that the explicit association expressed by links is directly under the subjective control of the author. The author need not rely on indirect association expressed by, for example, using rules for categorisation or metrics of similarity to classify or measure objects’ association with others. He or she doesn’t have to say *why* two objects are associated, just that they are associated.

The Web — Indexed Hypermedia

If the web was purely a large hypermedia structure then it would suffer badly from the scale-associated problems described above. Unlike ‘normal’ hypermedia its authors are huge in number and there is no sense of collaborative authoring of one large design or configuration. These problems of scale have been attacked primarily by indexing, which boosts connectivity. Considering our hypermedia example above as a web site, we can see the effects of adding a search engine in Fig. 3. Indexing has helped the integration of the non-textual objects and some of the textual objects, but the central non-textual nodes are still distant from the textual nodes to the left. In an architectural analogy, the search engine is like a narrow corridor that connects all objects but does not help with perceptibility.




Figure 3. A text-based search engine, accessible from any object, has been added to the structure of Fig. 2. Bi-directional links connect the engine to textual objects while unidirectional links (light grey lines) show the linkage from non-textual objects. Integration values confirm that indexed objects are still distant from non-textual objects. On the right is a ‘floor plan’ with the search engine as a narrow corridor accessible from all rooms but with ‘one-way’ doors (in grey) blocking access into non-textual rooms.

Improving connectivity with indexing tends to improve integration, especially for structures of large scale. This is dependent on the data types of objects, however: if many objects cannot be indexed then integration becomes skewed. Site authors can allow nodes to be indexed by adding captions or other metadata but then face the difficult task of guessing what indices to offer (as in the Berners-Lee quote above). Perceptibility is not improved and, since highly connected areas still tend to have low perceptibility, intelligibility is also not improved. Search engines are good circulation spaces, as we tend to pass through them for varied reasons and in varied directions. They tend to reduce overall graph depth and make other nodes more like occupation spaces. In this way they aid the general functionality of the overall configuration, and would be good places to introduce features that afforded awareness of other people if the web was to directly support the social aspects of use. Another asymmetry arises with search engines: the most used indices are in search engines and catalogues which cover many sites, and yet site authors and administrators have direct control only over their own site. They have a tendency to ignore these 'external' links and work primarily with their own 'internal' links. From the user's point of view this distinction is not so important, and they need not give greater authority to connections defined by site authors when they can use the links from external search engines.

Indexed/Categorised Databases

If we now do not consider inter-object links, and focus on indexing alone, we rely on search engines completely (Fig. 4). Generally we categorise some pre-existing corpus of data and use category-specific indices as our associative medium. We thus partition information in ways that may not be in accord with use.




Figure 4. On the left we see two sets of four objects. Each set contains homogeneous objects. Shown below each set is an index for that type. On the right is a corresponding floor plan, with two groups of rooms inaccessible from each other. Connectivity amongst each set is afforded by its index, represented as a narrow corridor that offers little perceptibility.

An index and search engine generally aid integration on a large scale but do little to support perceptibility between objects of a type. They tend to offer neither between objects of different types. More formally, in a heterogeneous set of objects integration values are very high as one must move via different engines. Intelligibility is therefore extremely poor. As with the web, indexing aids general functionality. Direct association such as authored, tailored hypermedia linkage is de-emphasised in order to take advantage of indexing's handling of scale. Access to manageably sized subsets and more specific association depends on subclassification schemes (classifications of objects' components) and metrics such as Boolean matching, treatment as a high-dimensional vector space, and probabilistic models. These are indirect associative media, and problems of handling non-discursivity and overly objective formalisms tend to occur.

Visualisation

In visualisation, the primary medium for association is the low-dimensional space that objects are positioned in. Configurationally, this is similar to an index (Fig. 5). Visualisation offers a means to connect heterogeneous objects and allows us to make perceptible complex interrelationships.




Figure 5. Objects of arbitrary types can be connected by placement in a visualised space. This space plays a configurational role similar to that of an index and offers the potential for maximal perceptibility. It is represented here as a space above the rooms, cf. the narrow corridor of Fig. 3, as it offers access to and visibility of all rooms.

Integration is generally very good, and the potential for perceptibility is high due to the great deal of information that can be handled when presentation and interaction are in accord with our perceptual skills. Realising the potential for perceptibility is not a trivial matter, however. Position, colour and other graphical dimensions must convey association, and their perceptual subtleties and interrelationships—bound up with our culture and language—make this a difficult design problem. Scale is another significant problem. Features such as screen

resolution, legibility of symbols, occlusion and clutter all inhibit perceptibility of object sets of the scale handled routinely by large database systems. Many visualisation systems add in a search engine in order to help with such problems. While intelligibility can be very high for smaller sets of objects, this tends to decline for larger sets.

Just as search engines aided general functionality, so with the embedding space of visualisation. The space is even more clearly a 'circulation space', and positioned objects are dead-end occupation spaces in a low-depth configuration that supports a variety of paths of movement.

The path model: an information access approach based on activity

Given the common semiological features of architecture and informatics, and the independence of categorisation of spaces and predictive power of Hillier's architectural theory, it would appear useful to work towards an analogous informational theory. This would be based on movement through and visibility of computer-based information. Association and similarity of information would be interwoven via consistent patterns of use. The configuration would record the functional processes that created it, and in this way constrain and guide future development of structure. The path model is a step in this direction.

What should we take as the fundamental basis, or model, of information access? Traditionally, informatics has focused on the form of the datum, with functions involving use and activity being built around the categorisation and formalisation of data representations. As Gödel showed, this approach leads to formal inconsistency—the metaclass slips into the class. Wittgenstein showed that separation between the structure and use of a language of representation ignores the social origins of meaning. We need to shift away from a view centred on formal categorisations that are separated from use, with activity peripheral.

The analytic and predictive power of Hillier's theory comes from activity being at the centre of the model. Movement and visibility in space, rather than the individual building, is at the centre of his theory. Presence at a spatial position and visual perception were quite easy to define as separate phenomena in the case of the city. With information we are forced to conflate perception of an object and presence in an object, since we have no strict equivalent of the reader being at a precise position within one and only one object. So, while movement through information space is an analogous basis for information representation, movement should be construed as the change in perception of information as one reads, writes, hears, scrolls, changes 3D viewpoint, and so forth. We move amongst information in many ways and, whether explicitly or implicitly, in this way we build a path of associations.

In the case of a simple web browser we might construct paths by logging which pages are successively loaded, approximating the order in which they are presented to the user. Use of web tools is usually interwoven with that of others such as mail tools and editors. Also, there is little, if any, fundamental difference between objects on and off the web. Therefore to understand web use we should also log use of documents and tools on the desktop. Ideally we would look at all media that form the environment of the user e.g. where that person is, the papers and books on the desk, what music they can hear, who else is in the room, what the weather is. The list appears endless and, somewhat inconveniently, it is. We should of course concentrate our efforts first on the most obvious and amenable targets for logging. Nevertheless, we as system builders (or information authors) should be aware of our assumptions about which categories of information and tools—computer-based or not—will be important to our systems' users (readers).

As with the city, our reasons for movement may be myriad. The association between items of information may express pre-planned action or our behaviour may be more contingent upon the situation. Rather than attempt to model the meanings in users' minds, Hillier was able to take statistical consistencies in users' paths as the social expression of meaning, and in this way bring the non-discursive into the discursive. Similarly we do not attempt to directly represent or formalise the consciousness of each user, but instead look for meaning in statistical consistencies within the collection of users' paths through information objects.

In the architectural case, uniqueness or identity of spaces was relatively straightforward and familiar, because of the way that individual spaces were defined as partitioning the total space under consideration. Here we begin by assuming that each information object has a unique name or identifier. An object might be a document, a program, an image, a person, or even a path, but its identifier is not dependent on the type or form of the object. Internal changes do not change the identifier of an object, and a copy of an object has a different identifier to the original. The metric of object similarity should not be based on identifiers or naming but on association, or linkage, in paths.

Let us consider that at any given time we model our perception of objects as a set of object identifiers. A path has an identifier, an author identifier, and a time-ordered sequence of object sets. For example:

```
path12, Matthew, ((doc77, image3), (doc77, image3, Bill), (image3, Bill))
```

Periodically, we record the set of objects most obviously perceptible to a user, and add this set to their path. Each such record might be made either after a chosen time interval (which might lead to contiguous object sets which are equal) or when some discrete move was considered to have been made (i.e. a change in the set of

perceptible objects). Other schemes are also clearly feasible, for example timestamping each set. (Current implementations rely on timestamping.) More difficult to justify and to implement would be the association with each object of a probability describing relative strength of perceptibility. Perhaps some tools involved in information access will not be able to co-ordinate with others in order to define the set for a given moment in time. Instead, either concurrently or as part of a later analysis, another program combines individual logs. By means of timestamps, for example, it weaves logs together to form the sets of contemporaneously accessed objects.

A null path has no objects in it. Two paths are equivalent if they have the same author and contain the same object sets in the same order. Timestamps would offer even greater specificity. One can imagine a variety of metrics of path similarity, with each giving a maximum value for equivalent paths, and decreasing as authors, object sets and orders differ. Contiguous path segments, and paths defined by selectively filtering out particular objects from a larger path or inserting selected objects into a shorter one, could be compared similarly. This allows us to compare our recent behaviour with our own past (i.e. comparing a short segment of a path, with end 'now', with the full path so as to find similar episodes in the past) and of course with the paths of others.

Two objects are similar if they consistently tend to be in similar paths. This is irrespective of the objects' heterogeneity or homogeneity of content or form, as it is their use that we focus on here. A path's start and end might be defined by the earliest and latest of all logged activity by the person involved. Alternatively we might delimit paths upon noticing prolonged periods of inactivity, or abrupt changes in the membership of the set of objects worked with. One danger here is that inactivity in the logged behaviour may be due to an inability to log significant activity, rather than a lack of activity that might be logged. We should also be aware that such partitioning based on inactivity, object set consistency, etc. is pre-categorisation or pre-judgement that may not suit later use.

We emphasise that we do not assume that following paths in this way is the only way to move through information, or that a tool based on path extension would be the only one in active use on a computer at any one time. Instead here we look for a way to model the movement made via whatever ensemble of tools the user cares to use. This ensemble may or may not include a tool based on paths, but nevertheless we can consider the tool that initially offers information access, e.g. the desktop upon login, to be an object. Like most objects, it offers access to a number of subsequent objects, in this case mail tools, web browsers, path-based tools and so forth. In this way, our analysis can always involve a single configuration of objects.

Perceptibility and Intelligibility

When perception of an object A makes another object B perceptible then we say that A is connected to B. This does not mean that the ability to perceive the connection is necessarily easy or natural. As mentioned earlier, we diverge from Hillier here, whose 'line of sight' visibility was simpler to measure, treated as if it was an objective, formal property. Here, it is more obvious that perceptibility may involve complex actions and interpretations understood only by the community of use (and not by us outsiders, the analysts), but nevertheless the space syntax analogy suggests we represent this by a practical, binary property: connectivity. If one could use perceptibility to define (unidirectional) connectivity, one could form graphs of connections between objects, and apply Hillier's integration to express how well connected individual objects and configurations are, and intelligibility as the correlation between visibility and integration. The degree of perceptibility for the community of users will hinder or help their everyday use, and will be affected by the degree to which their informal language of use is in accord with the formal language of analysis. As yet, we cannot go so far. Here we must rely on more intuitive notions of perceptibility. Note another aspect of movement and view that Hillier assumes but that we cannot: people can see each other within the same space, and often in connected spaces. This is a simpler assumption to accept for cities than for most information systems.

Perceptibility is a vital feature of movement, determining affordances for movement and our confidence in moving. An intelligible configuration is therefore one in which what is perceptible of the objects is a good guide to the global structure. One perceives where to go, as well as how objects fit together at the larger scale. When at a given object, it is often the case that some connected objects are more perceptible than others. An HTML page makes linked pages visible by colour, font, location on the page, a textual description and so forth. The changing background colour of the page, a position far down a very long page, or having an unhelpful description may make a link less perceptible. The page design may also reflect, or fail to reflect, where we might move in the larger-scale configuration e.g. a clear step up to the parent directory in a graphical file browser, or another confusing step in a maze of 'See also' cross-references in manual pages.

An extreme case, showing maximal connectivity and minimal perceptibility, would be a tool akin to the 'Go To' field in a web browser, but isolated and used only once. One could type in an arbitrary URL, and then jump to the page. All objects are connected to the tool, but none are perceptible. This is somewhat similar to the case of web search engines, which are connected to a very large number of pages but do not make those pages easily visible or perceptible.

When we look at objects and how we move around them, we can again focus on occupation and movement. Occupation means the use of objects for activities that are at least partly and often largely 'static': reading a long piece of text on page, interacting with an embedded Java application. Making a strong analogy with Hillier would suggest that 'dead end' spaces may be best for occupation and that such 'focused' texts or applications should be kept distinct from highly-connected objects. Connecting objects, such as well-connected ring-bridging objects, complement occupation objects. They offer flexibility and efficiency of movement and support more dynamic activities. Together with designing so as to minimise depth, implementing such distinctions should enhance the general functionality of the configuration, increasing the affordances of encounter with objects and, assuming some kind of mechanism for perceiving people in nearby information spaces, of encounter with people.

Adaptation and Social Effects

Some information access techniques require a start-up or training period, whereas others operate immediately. In either case, the pace of creation and evolution can be much faster than that of urban environments. Cities generally present the same visible configuration to all and, like the basic access techniques discussed above, do not adapt differently and specifically to each individual user. Given the dynamics of computers, however, one must ask how quickly information access can be tailored to individual habits and needs.

In the path model, we define a new path as we move through information, ostensibly expressing our interests or activities. The path will be continually logged, adding to the configuration of paths shared by our community of users and adapting the relative path similarities. Given our current position at the end of our path we may find corresponding periods in other similar paths, and objects that consistently arise slightly further along from those periods. These can be presented to the user i.e. made perceptible, affording future movement and triggering off adaptation in the form of further logging, reconfiguration, and presentation. Now we see what 'retrieval' can mean here. Having created or accessed an information object in the course of our work, we can be presented with other 'relevant' objects without dependence on their homogeneity. Furthermore, relevance is derived from the context of similar human activity, and not raw data content or 'expert opinion'.

It is feasible that mere presentation of possible paths for information movement might be logged, trigger reconfiguration, and spark new presentation. Setting such a 'hair trigger' which did not wait for (or give weight to) user choice among possible paths of movement might offer an interesting 'guided tour' if the pace were calm. If too fast, however, the tool might bolt off wildly into the informational horizon. This sprint toward infinity, a surfeit of adaptation, would be at least as bad as having no adaptation at all. In between these two extremes, controlled adaptation would move at the pace of the user, taking account of evolving activities and choices, and continually offering appropriate awareness of objects and affordances for action.

Since paths are associated with people we make manifest the community's use of information. Association and meaning are thus socially determined. The patterns of co-perceptibility in the city are strongly linked to co-awareness and the social, and similarities in the way people act in information potentially can make those people perceptible to others. Paths offer a means to find out about people's past activities and, as paths grow and adapt, people's ongoing activity. This may be useful when, for example, this enhances interpersonal communication and community. The act of writing a proposal for a research topic could trigger presentation of relevant references and authors, and also local people with experience in the topic. Note, however, that we do not have inherent symmetry as is normal in physical visibility on a street, or audibility in conversation. Like all information that identifies people and may be accessible to others, either at the time of action or at some later date, paths raise issues of privacy and invasiveness [20].

Implementations

So far, application of the path model has been in two areas: a URL and file recommender tool and 2D visualisation of URLs and files. This work was carried out while the author was at UBS Ubilab, in Zürich, Switzerland. An earlier system, for URLs only, was described in [1], and so we offer only a brief description of the more technical details. Web usage was a convenient area for early experimentation because the URL naming system offered unambiguous references to heterogeneous objects, we use the web regularly, and browser activity is easy to log. We extended the Muffin web proxy (muffin.doit.org) to log URLs in a relational database. We now also log the use of local files inside the *xemacs* editor, with logging triggered by switches between editor buffers. Since we only recorded loading of URLs and switching of buffers, and not visibility on screen, a path involves not object sets (as discussed above) but a sequence of individual timestamped URLs. Each user can turn path logging on and off at will. By default, each path is potentially visible to all those who contribute paths i.e. the set of paths is treated as a shared resource. We are also experimenting with adding the content of web pages and files to the path.

http://www.arosabergbahnen.ch/	18.0
http://www.arosabergbahnen.ch/Grafiken/collage.jpg	15.0
http://www.klosters.ch/images/tn_gotschna.jpg	9.0
http://www.klosters.ch/gotschna.html	9.0
http://194.158.230.224:9090/telenet/CH/180/2.g.html	9.0
http://www.arosa.ch/main.html	9.0
http://www.arosa.ch/ski/auswahl.html	9.0
http://ad.adsmart.net/src/goski/mountains^1?adtype=ac&bgcolor=F...	9.0
http://www.arosabergbahnen.ch/home.html	9.0
http://www.arosa.ch/	9.0

Figure 6. A snapshot of a URL recommendation list from our system. Starting to choose a day for a ski trip, the author accessed web pages with detailed weather reports for the mountains of Switzerland, including the *telenet* service of a local university. Recommendations were drawn from six sequences of past activity in three people's paths, and were mostly for web sites of ski resorts near Zürich, such as Arosa and Klosters. The numbers in the right hand column are the tally values used in ranking relevant URLs.

We treat the person's most recent path entries as an implicit request for recommendations. Every few minutes, the system takes the most recent path entries i.e. the end of the path, and searches for past occurrences of each URL (or, more generally speaking, each symbol). Currently, this search can either cover one's own path or all paths within the shared path set. In the interface, a slider is used to set how long this 'recent' period is, and one can select which paths (or owners) to draw from, thus allowing people to use knowledge of their colleagues to steer the recommendation process.

The system then collects the context of each past occurrence of each of the most recently used symbols—the path entries following soon after each past occurrence. Another slider sets this period's length. The system then collates these symbols, removes any which are amongst the set of recently used symbols (since we don't need to recommend symbols the user just used), and then presents the top ten from this ranking as a recommendation list. The system thus recommends to the user symbols that were frequently used in similar contexts but that it has not observed recently in the user's path. The people whose paths contributed to recommendations are not identified in this list. An example recommendation list is shown in Fig. 1, above, and four snapshots of a visualisation of the same example are shown in Fig. 2.

This example is intended to demonstrate how recommendations, such as ski information given weather pages, might not be the most obvious thing for a system to do until one considers the context of use and history behind it. The recommendations suited the author extremely well, as they were useful in his particular situation. If the winter weather was good, then a ski trip could go ahead. In this case, he got the Klosters information from his own path, and the Arosa information from another person's path. Never having been to Arosa, or to the Arosa web site, the recommendation was therefore both novel and relevant. The example also offers a contrast with tools based on content. The likely recommendation would simply have been for yet more weather pages. Lastly, the recommendations include heterogeneous data: JPEG images as well as pages of HTML. This demonstrates the ability to mix media that are usually indexed and searched by disjoint systems. Since paths involve identifiers, and similarity involves patterns of use of identifiers, it doesn't matter what the content 'inside' an identifier is.

As with [1], co-occurrence statistics were used to define a similarity metric in the layout algorithm of [21], and the resultant layouts were displayed in the visualisation tool of [22]. Sections of paths cross at multiply accessed symbols. Symbols adjacent to these recurrences are brought near to each other, building associations between different periods of time. Layouts now involve a relatively small number of symbols, which tend to form 'strands', each of which is a sequence of symbols collected prior to symbol tallying i.e. one of the similar periods of past activity used in the recommendation process. Our previous visualisations involved layouts of complete paths, involving thousands of URLs. In both recommendations and in visualisations, we then used statistics of co-occurrence over entire paths. This diluted the specificity of recommendations and also involved a co-occurrence matrix of size $O(N^2)$, where N is the number of symbols in paths. We now avoid this quadratic data structure, taking advantage of database indexing facilities to generate required statistics on the fly. Recommendations are made periodically—another slider in the interface sets how often—and take a few seconds to make. Most of the work takes place on a server machine devoted to the database. We can make a matching visualisation in under one minute on a Silicon Graphics O2.

Statistics for recommendation and visualisation thus form categorisations or abstractions over path entries that are not fixed *a priori*. Instead, recurrence statistics are made anew for each person at the time of, and using the context of, each recommendation operation. Each new entry in a path changes the pattern of symbol co-occurrences throughout the shared configuration of paths. Even if you have not accessed new information recently, your recommendations may still vary as other people's activities change their paths.




Figure 7. Four snapshots of an interactive ‘map’ of URLs based on the path model, scaled down from actual screen size, and with the frame and controls of the surrounding visualisation tool (described in [22]) cropped. Following access of web pages with detailed weather (*wetter* in German) reports for the mountains of Switzerland, recommendations involved URLs for web sites of ski resorts near Zürich, such as Arosa, Davos, and Klosters, as well as URLs for other weather pages. Each snapshot shows the same positions, based on co-occurrence of URLs, but shows in a lighter colour those URLs that match the given search string e.g. *arosa* in *www.arosa.ch*. The *Klosters* snapshot also shows three URLs which have been ‘popped up’ by the tool to show textual detail.

Ongoing work focuses on adding the symbols contained in web documents (such as words and URLs) to paths, in extending the interface to allow more active steering of the recommendation process, and in more specialised tools for logging and recommending Java software components. In adding content, the goal is to treat words, filenames and URLs equally within the same model. In active steering, the user could choose to add symbols to (and remove others from) the ‘recent’ symbol set. At one end of a spectrum would be the removal of all automatically collected symbols, and the explicit specification of a query-like set of terms. At the other end would be the current implicit or passive specification.

Conclusion

In the path model and related systems, information representation has human activity at its centre. Categorisation, in the form of statistical consistencies in similar past activity, is built around this shared, adaptive, relative system of symbols. In our recommendations and visualisations, we aim to show and support user activity, and to take account of the interdependence and interweaving of information tools such as web browsers, editors and so forth. We have described how the path model was built by analogy with a theory of architecture, and more generally tried to make clearer the similarities between the fields that can profitably be used in both system design and theory. These similarities exist even while there are substantial differences. For example, mutability: bits are easier to move around than bricks. This difference is diminishing, however, as some information structures, such as home pages of large corporations and the search engines we so rely on, are made persistent, reliably available facilities. In architecture, banks, galleries and museums increasingly have their supporting structure on the outside, so that internal partitions and passageways can be easily reorganised. Buildings are made more adaptable to their current users while URLs become more corporate and fixed. We suggest that it is the common semiological basis of informatics and architecture that makes for these similarities. Human activity involving information in computers not so special. It is little different from activity involving other media of communication and representation. We thus gain an understanding of how the choices for future significant action of each individual, in and through symbols of all types, are influenced by the past and current activity of others. This gives rise to the ‘hermeneutic circle’ of representation, interpretation and action. We have tried to make clearer how these foundations lie deep—almost hidden—beneath the theory and practice of informatics, and to show how we can build upon them in new ways when we better comprehend their origins, strengths and weaknesses.

Antoni Gaudi, architect, said that “originality is going back to origins.” In this paper we have looked back through informatics to mathematics, in order to make more ‘ready to hand’ our materials and tools for building new approaches to information representation. We consider that contemporary mathematics and philosophy show us that activity, language and the social are at the centre and origin of informatics, rather than on the periphery of a core of objectivity and positivism. This view does not make us discard positivist approaches for those based on structuralist, or even post-structuralist ones. As was pointed out by Ricoeur, positivist approaches allow access to open semiological systems while structuralist approaches, such as the path model, involve a closed, relative system of symbols [23]. Structuralism needs positivist approaches to connect each new symbol into its representational systems, and so begin another turn of the hermeneutic circle: taking a new symbolic representation, interpreting it in terms of how we use it with others, and creating new patterns of

action and signification. Thus, while path-based systems and related collaborative filtering systems may be able to represent and support more of the complexity and subjectivity of information than positivist information access approaches, the former are dependent on the latter.

We believe that the understanding gained from re-examining the origins of informatics opens up new space for originality in informatics theory and systems design. This shift towards what Wegner called the interactive paradigm will let us handle information that seemed intractable, combine tools that seemed unrelated and interweave theories that seemed incommensurable. This can be done if we face up to informatics' limitations and historical biases, and see what Paul Churchland calls the 'deeper unity' of modern intellectual discourse all the way from neurophysiology to philosophy of science. Informatics can then step forward to take its place with contemporary mathematics, philosophy and semiology.

Acknowledgements

The author wishes to thank the editors for discussion and comment related to this paper, and also Dominique Brodbeck, Paul Dourish, Luc Girardin, Christian Heath, Aran Lunzer, Marissa Mayer, Kerry Rodden, Lisa Tweedie, and those who directed the laboratories where the bulk of this work was done: Hans-Peter Frei of UBS Ubilab, and Yuzuru Tanaka of the Meme Media Laboratory, Hokkaido University. And Jeff Mills (Liquid Room), and Mogwai (Young Team).

References

1. Chalmers M, Rodden K, Brodbeck D. The Order of Things: Activity-Centred Information Access. Proc. 7th Intl. World Wide Web Conf. (WWW7), Brisbane, 1998, 359–367.
2. Hillier B. Space is the Machine. Cambridge University Press, 1996.
3. Foucault M. The Archaeology of Knowledge. Trans. A.M. Sheridan Smith, Routledge, 1972.
4. Nardi B, Barreau D. "Finding and Reminding" Revisited: Appropriate Metaphors for File Organization at the Desktop. ACM SIGCHI Bulletin, 29(1):76-78, January 1997.
5. Berners-Lee T. World-Wide Computer. Comm. ACM, 40(2):57-58, February 1997.
6. Resnick P, Miller J. PICS: Internet Access Controls without Censorship. Comm. ACM, 39(10):87-93, October 1996.
7. Nardi B. (ed.), Context & Consciousness: Activity Theory and Human-Computer Interaction. MIT Press, 1996.
8. Wegner P. Why Interaction is More Powerful than Algorithms. Comm. ACM, 40(5):80–91, May 1997.
9. Churchland P. Folk Psychology. In: On The Contrary: Critical Essays 1987-1997. P.M. Churchland & P.S. Churchland, MIT Press, 1998, 3–15.
10. Everitt N, Fisher A. Modern Epistemology: A New Introduction. McGraw-Hill, 1995.
11. Karatani K. Architecture as Metaphor: Language, Number, Money. MIT Press, 1995.
12. Wittgenstein L. Philosophical Investigations (3rd ed.). Trans. G.E.M. Anscombe, Oxford University Press, 1958.
13. de Saussure F. Course in General Linguistics. Trans. Wade Baskin, McGraw-Hill, 1959.
14. Gibson J. The Ecological Approach to Visual Perception. Lawrence Erlbaum, 1979.
15. Chalmers M. Comparing Information Access Approaches. Submitted for publication.
16. Suchman L. Plans and Situated Actions: The Problem of Human Machine Communication. Cambridge University Press, 1987.
17. Schmidt K. Of Maps and Scripts: The Status of Formal Constructs in Cooperative Work. Proc. ACM Group 97, Phoenix, Nov. 97, 138–147.

18. Colomina B. *Privacy and Publicity: Modern Architecture as Mass Media*. MIT Press, 1996.
19. Major M. (ed.), *Proc. 1st Intl. Symp. on Space Syntax*. University College London, April 1997.
20. Bellotti V, Sellen A. Design for Privacy in Ubiquitous Computing Environments. In G. de Michelis, C. Simone and K. Schmidt (Eds.) *Proc. Third European Conf. on Computer Supported Cooperative Work, (ECSCW'93)*, pp. 77-92. Kluwer, 1993.
21. Chalmers M. A Linear Iteration Time Layout Algorithm for Visualising High-Dimensional Data. *Proc. IEEE Visualization 96*, Oct.–Nov. 1996, 127–132.
22. Brodbeck D, Chalmers M, Lunzer A, Cotture P. Domesticating Bead: Adapting an Information Visualization System to a Financial Institution. *Proc. IEEE Information Visualization 97*, Phoenix, Oct. 1997, 73–80.
23. Ricoeur P. *Hermeneutics and Human Sciences: Essays on Language, Action, and Interpretation*. Trans. J. Thompson, Cambridge U. Press, 1981.