CHAPTER X

STRUCTURALIST INFORMATICS: CHALLENGING POSITIVISM IN INFORMATION SYSTEMS

Matthew Chalmers, Computing Science, Glasgow University

Abstract

It is often assumed that the information representations used in interactive information systems are grounded firmly and objectively in mathematics. This positivist view is at odds, however, with contemporary mathematics and philosophy of science, which sees subjective, social, and situated symbol systems — semiology — at the base of all human information use. We describe some of the reasons behind the paradigm shifts away from positivism in mathematics and linguistics, relate them to the information access approaches of information retrieval, collaborative filtering and the path model, and so argue that a similar paradigm shift away from 'folk informatics' is not only needed, but has begun.

1 INTRODUCTION

The 50th anniversary of the publication of Vannevar Bush's *As We May Think* (Bush, 1945) occurred recently. Perhaps the most seminal document in the history of information systems research, *As We May Think* is often held as foundational to hypermedia, the World Wide Web, and information retrieval (IR). Interestingly, Bush does not put forward that the collection of links used in connecting hypermedia objects is a solution in itself. A 'sea of links' merely replicates the problem of having a sea of records, tuples or documents. Bush proposed that human activity in the form of 'trails' or 'paths' was the solution i.e. particular human–selected sequences of links associating information objects, and not the vast set of possible associations that hypermedia, databases and libraries provide. Note also that such a path does not directly involve content. It relies on human interpretation of content, but the meaning or value of an object is primarily given by the chosen context of a path. A link between two objects thus only has meaning within the context of a particular path of a particular person. The hypermedia and retrieval systems that declare an origin in Bush's work generally do not work with the phenomena of context and person as Bush proposed: context and person are not part of their underlying model of knowledge. Re-reading *As We May Think* and seeing this disparity between underlying models suggests we should look more closely at information systems' models of knowledge and interpretation.

Two complementary schools of thought exist with regard to the basic underlying assumptions and philosophies that guide our research in information systems. One 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. In this paper we will focus on theory and practice of systems for accessing large bodies of information, and our examples from such schools will be information access approaches such as IR and collaborative filtering (CF). Within the context of this paper we will therefore treat 'information access' (and indeed 'informatics') effectively as representatives or synonyms for 'information systems.'

The first school of thought sees itself as part of traditional computer science, rooted in formal models and taxonomies that encompass the individual variations of users and that are often derived from experimentation and observation in controlled conditions. IR, based on content analysis of documents, fits this paradigm. Complementary views are held by those who hold the social and the semiological as primary, and consider that formal, mathematical models are insufficient to model the complexity of human activity and ultimately of limited utility in guiding system design and development. Collaborative filtering (Goldberg, 1992), sometimes known as recommender systems (Resnick, 1997), is an example here. In this paper we take a stance firmly towards the subjective and linguistic end of the spectrum, arguing that, even though theories and systems based on objective formalisations are extremely useful in particular situations, they are essentially inadequate to model and support the full complexity of human use of information. We suggest that this is due to traditional computer science having roots in reductionist, positivist models of knowledge and interpretation. As discussed in more detail later in this paper, we suggest that CF and its extensions such as the path model are examples of a nascent paradigm for information access systems that breaks away from these traditional roots.

In more general terms, we take the view that use of computer–based information is no different from any other human use of information. It is just another aspect of significant human action i.e. activity that both involves the use of symbols over time, and can be taken as symbolic. In other words, we treat informatics as an instance of semiology, as ultimately more akin to linguistics than our field's traditional ideal or exemplar, mathematics. We aim to make more visible some of the assumptions and models that underlie interactive information systems. We are often unaware of the models of knowledge and information that we build on, what they afford and what they inhibit, where they work well and where badly, how they work in isolation and in combination. Here we aim to make clearer some of those buried layers — the 'archaeology' (in the sense of Foucault (1972)) of information systems' foundations.

2 POSITIVISM AS INFORMATICS' TRADITIONAL PARADIGM

The opposing schools of thought mentioned above are exemplified by 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 (Nardi, 1997). At issue was the importance of location-based searching over logical retrieval. In response to Nardi and Barreau, Fertig, Freeman and Gelernter put forward what might be considered a traditional computer science viewpoint, suggesting that better indexing tools would allow users to gain the benefits of storing large amounts of information. 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.

At the core of Nardi and 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 the scale, complexity and heterogeneity of information representation is increasingly problematic. As Tim Berners-Lee pointed out with regard to search engines for World Wide Web data (Berners-Lee, 1997), "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. Due to the scale and distribution of ownership of information, most metadata has to be created manually by its owners, for example by page authors or site managers adding textual captions or tags. An example metadata scheme for Web data is PICS (Resnick, 1996).

If metadata, however, is not formally controlled to ensure consistent use, it becomes data: open to be written, read and interpreted as each person sees fit. The same problems of which metadata to choose (or trust) arise e.g. if a large commercial corporation sets up its own rating sites, how would you interpret its ratings of its own products' sites? When there is a large volume of metadata one will need to classify, index and search it, requiring 'metametadata' that may itself eventually become voluminous and problematic. In other words, with uncontrolled metadata we only defer the problem of matching available information to users' interests and activities — and when metadata open to individuals' use slips down to be just more data, the problem reappears.

Nardi and Barreau are amongst those in HCI who are 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 approach as excessively reductionist, as demanding so much control over experimental conditions that 'inconvenient' phenomena essential and central to everyday work are excluded: the interleaving of tools both in and out of the computer, application to a set of familiar work topics and goals, and the setting within a community of use. By not taking into account the complexity and situatedness of human activity, it is difficult to offer the hoped–for practical benefits to systems designers.

Work such as Activity Theory (Nardi, 1996) takes a more realistic view of the subjectivity, dynamics and social context of individuals' action, broadening 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 artifacts 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 toward the artifact, and towards system design, by emphasising how we represent, categorise and interpret information.

We see the need for a more radical break away from informatics' traditional belief in its ability to represent the world in logical and objective terms. We see this belief as naive, firstly because within formal informatics, the weakness of such a reductionist view has been demonstrated. Wegner recently published a formal proof that interactive computing is an inherently more powerful computational paradigm than purely algorithmic computing (Wegner, 1997). The main point for us here is Wegner's formal demonstration of the irreducibility of interaction to algorithms. We can gain greater expressiveness and analytic power with interaction-based approaches but only if we abandon our fields' traditional belief in algorithmic formalisability of information use.

Secondly, a reductionist stance does not reflect the contemporary state of mathematics. Wegner echoes this century's history of mathematics (and indeed physics) where the dream of reducing the world to a 'clean' and purely objective mathematical model was shown to be an illusion. In the next section, we see what we can learn from mathematics' reconsideration of its own foundations.

3 FROM MATHEMATICS TO LANGUAGE

Underlying interactive information systems are abstract data representations that formalise the content and behaviour of some aspect of our world. We use such abstractions in order to input, store, index, manipulate and present our information. We use mathematically–based formal representational systems — algebraic structures, structures of order and topological structures — to represent the properties and relations between real world objects.

Since we use such mathematical schemes as foundations for building information representations, we should understand mathematics' foundations. Prior to the 20th century, mathematics had hoped for convincing proof that it was founded firmly in formal, objective absolutes, based purely on a logical process of naming things in the real world. This positivist paradigm was common to many intellectual fields, but was weakened by enduring problems such as the inability to ground Euclidean geometry in more than perceived self–evidence (in particular, the axiom of parallels), by Cantor's demonstration that the set of all sets is indeterminable, i.e. the 'metaset' of all sets is a set, which keeps slipping down from metadata to data.

Hilbert led the formalist attempt to shake off such contradictions and inconsistencies by rendering mathematics genuinely independent of perception of the 'real' world, asserting (as quoted in (Karatani, 1995)) 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." Even having discarded the claim to truth, however, claims to consistency were attacked from two directions. One blow 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 wherein meta-mathematics, understood as a class, gets mixed into the formal system as a member of that class. The other blow was from without, from Wittgenstein, who overturned the notion that mathematics' formal system can be solidly deduced from axioms. Proof is a language game (Wittgenstein, 1958), 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 demonstrates the problem of consistency in the structures we use for information representation, while Wittgenstein replaces the axiomatic basis of their truth with semiological, socially constructed meanings.

A question one might then raise is: what is semiology based on? The contemporary view in this field is that we cannot dig further: there is 'no exit from language.' When we interact with a formalised or abstract representation of natural language or behaviour, we use it in our world of human interpretation and activity, and we necessarily involve natural language. We may try to step up to a metalevel by means of formalism and abstraction but inevitably we slip back down again. Within a controlled environment such as the computer, formal, self–consistent and finitely defined representational systems can operate, but in human–computer interaction the informal, subjective and infinite reassert themselves.

What can we take from contemporary semiology that will offer practical help in our informatics theory–building and information system design? We propose to draw from the paradigm that became dominant only a few decades before Gödel and Wittgenstein revolutionised mathematics' core, that has often been identified or equated with Wittgenstein's 'language game,' and that indirectly motivated mathematics' paradigm shift: structuralism (de Saussure, 1959). Here we aim to point out some of its key features that echo earlier discussion but that we later use in discussing information retrieval and collaborative filtering. A more extended discussion and comparison of these and other information access approaches is the subject of a forthcoming paper, (Chalmers, 1999).

Structuralism unified linguistics and semiology, and displaced positivism in linguistics. A word was held to be a symbol, and no longer seen as being part of a historically traceable process of absolutely and logically naming a thing in the real world.

Saussure's view was that naming was a relative or differential process. 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, and not by a one–to–one relationship of naming a unique, absolute, ideal thing in the world. While we may use a word to signify a thing in the world, we do 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: stone and steel in sculpture and architecture, paint and ink in printing and graphic design, and so forth. Symbol associations span the media used for expression: the picture of a pipe, the word 'pipe,' and the wooden pipe are interlinked. It is interpretive choice or reaction that creates significance, and so any action in any medium can be taken as significant, and hence as a symbol. Language, therefore, is the dynamic pattern of individuals' time–ordered symbolic activity in all media: subjective, situated and social.

4 REPRESENTATION IN INFORMATION SYSTEMS

Now we can see a fundamental limitation in the relatively common view that ignores information objects' use in human activity and objects' interrelatedness, treating each object in isolation and, as is generally the case in IR, restricting consideration to the content of each object (e.g. the words and images inside a document). To assume that the symbols contained inside an object faithfully and fully describe its meaning, irrespective of perception and use, is a positivist approach.

That perception of a structure or representation is bound up with the perception of use became familiar to many within information systems, especially HCI, via the ecological theory of perception popularised by authors such as Don Norman, but originally developed by J.J. Gibson (1979). 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 on the part of 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." We start to see how contemporary philosophy has already influenced modern HCI theory.

Within HCI, Lucy Suchman (1987) 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 (Schmidt, 1997). 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, on the other hand, 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" (Suchman, 1997:58). We 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 usually unrepresented in our information systems. In Suchman we again see the influence of contemporary discourse on language and representation.

The most direct influence of contemporary philosophy on information access may be that of Heidegger's philosophical hermeneutics and Searle's speech act theory on the workflow model. In *Understanding Computers and Cognition*, Winograd and Flores (1986) took account of the situated use of information in work but, as (Chalmers, 1999) discusses in more detail, they then imposed a reductionist and objectified categorisation of user activity that made workflow an information access method that is famously inflexible and disempowering of the individual. Contrarily, workflow may be well suited to organisations which need strict control for safety, efficiency and legal reasons e.g. in performing complex medical trials and processing insurance claims.

In contrast with Gibson and Suchman, and indeed Winograd and Flores, traditional content-based approaches operate only on the subset of symbols and attributes that are conveniently accessible and contextually independent e.g. no matter who has a textual document and what activity they are involved in, the same set of words are accessible inside the document. Of course, this specificity affords a great strength: techniques such as indexing of contained words allow quick and automatic searching of large volumes of data, but we rely on the assumption that the context of use and interpretation of the person involved is not significant.

This assumption is true when one wishes to find all documents that contain the word ‘pipe,’ but false if one wishes to identify the documents that one individual person would consider to be about pipes, or would be useful to that person in learning how to make a pipe. Tasks such as the latter two seem irredeemably fuzzy and difficult to traditional IR but perfectly natural to us in our daily lives. Specification of information need generally requires the user to define what they want in terms of content and in a formal query language. For the former objective task IR works well, but for tasks involving subjectivity the information representation excludes the requisite phenomena of context and person, and interaction — querying — is notoriously difficult.

It is in this area that collaborative filtering is strong. CF deals with similarities in profiles of choice and use of information, and handles tasks such as “name a movie that I’ll enjoy” with ease. Recommendations of useful or interesting information objects are specific to the person involved, but independent of their particular current activity or work context. Formalising why one liked or disliked an object is avoided or reduced, for example by selecting from a five point rating scale from 'very good' to 'very bad,' and so one can react according to one's informal understanding. Even simpler is Amazon.com's book recommender (www.amazon.com) where each book purchase serves as a positive rating and is added to the purchaser's profile.

Similarity of two profiles is considered high if they have similar ratings for the same objects. Given a person’s profile, a small set of 'neighbours' is determined, being the most similar profiles. The system presents to a user a small number of objects that were rated highly in neighbours but that he or she has not yet rated. Sometimes one can explicitly choose the subset of people from whom profiles should be drawn, allowing use of one's knowledge of colleagues' and friends' expertise.

CF also cuts across boundaries between apparently heterogeneous data. Books, movies, music, people: CF operates on the names or identifiers of objects and not (necessarily) the content. It circumvents IR’s intransigent difficulty with comparing symbols in different media e.g. in measuring the similarity of a word and an image.

CF systems need ‘bootstrapping.’ Objects that have not been rated can not be recommended, so CF relies on external sources, such as other information access tools, to introduce new objects to at least one user, so that their use or choice can be added to their profile. This is not a great weakness, however, as we of course use many tools (including IR tools) in our everyday work, each of which can offer use and selection information. CF is therefore explicitly dependent on other sources of information, in direct contrast to the independent (or isolated) operation assumed by IR.

Another characteristic of CF is that it does not offer contextually–specific recommendations. It does not take account of the temporal order of activity. A profile generally records what a person liked or used, but nothing of the time of use and the other symbols used around that time i.e. the context of use. Recommendations are therefore based on all of a person's history, and not their current task or activity. The path model extended CF by adding time to the information representation, in order to gain this contextual specificity (Chalmers, 1998). Paths, echoing *As We May Think*, are timestamped logs of symbol use over time e.g. logs of URL accesses in a web browser. The path model defines the context of a symbol's use by an individual as the other symbols used by that person within a selected window of time: user–selected, in current implementations. Information is presented to the user by taking the recently–used symbols in their path, finding the most similar windows of activity in the past of their chosen set of individuals, and finding symbols that were frequently used thereafter in the windows but were not used recently by the user.

In summary, CF and the path model break away from positivist approaches to information access such as IR because symbols are defined relatively, by similarities and differences in individuals' use, and not by appeal to objective classification. Heterogeneity of content types is unproblematic. They are based on very few *a priori* categories, instead relying on ephemeral, subjective and adaptive categorisation. Formalisation of information need by the user is avoided or minimised. These features form our basis for the claim that they are practical examples of a complementary paradigm to IR, and early examples of a structuralist or even poststructuralist informatics.

5 CONCLUSION

The implementation work in CF and paths has progressed so far with relatively little theoretical discourse to underpin it. In this paper we aim to help redress this imbalance, and to aid in future development, by making more explicit some of the representational issues and assumptions that informatics shares with other semiological fields. Our field has not taken full account of the way that information representation schemes have affordances: they are tools or objects built on characteristic strengths and weaknesses relative to our uses, interests and abilities. Such characteristics in themselves are neither good nor bad. It is lack of awareness of them and of their consequences that causes us problems. We should understand what assumptions our information representations are built on, and hence what they afford and what they inhibit, and so move beyond naive 'folk informatics'.

To achieve this understanding, and to broaden and strengthen our field's foundations, we must abandon reductionist, positivist theories of informatics and rebuild on a base of contemporary philosophy and semiology. This requires a paradigm shift, a Kuhnian revolution within our science. Informatics is younger than fields such as mathematics and linguistics, but sadly seems not to have learnt from them. It is less mature and relatively isolated in having a positivist theoretical infrastructure underlying its everyday practice — despite the contributions of a number of writers and researchers. As with any other symbol and tool, our field is most usefully understood as just one of many, interwoven with others involved in human language and meaning.

References

Berners-Lee, T. (1997) "World-Wide Computer", Comm. ACM, 40(2):57-58.

Bush, V. (1996) 'As We May Think', Atlantic Monthly, July 1945. Reprinted in ACM Interactions 3(2):37-46.

Chalmers, M., Rodden, K. & Brodbeck, D. (1998) 'The Order of Things: Activity-Centred Information Access', Proc. WWW7, 359–367.

Chalmers, M. (1999) 'Comparing Information Access Approaches', to appear in J.ASIS.

Foucault, M. (1972) The Archaeology of Knowledge, trans. A.M. Sheridan Smith, Routledge.

Gibson, J.J. (1979) The Ecological Approach to Visual Perception, Lawrence Erlbaum.

Goldberg, D. et al. (1992) 'Using collaborative filtering to weave an information tapestry', Comm. ACM 35(12):61-70, December 1992.

Karatani, K. (1995) Architecture as Metaphor: Language, Number, Money, MIT Press.

Nardi B. (ed.) (1996) Context & Consciousness: Activity Theory and Human-Computer Interaction, MIT Press.

Nardi, B. & Barreau, D. (1997) ''Finding and Reminding' Revisited: Appropriate Metaphors for File Organization at the Desktop', SIGCHI Bulletin, 29(1):76-78.

Resnick, P. & Miller, J. (1996) 'PICS: Internet Access Controls without Censorship', Comm. ACM, 39(10):87-93.

Resnick, P. & Varian, H. (eds.) (1997) Special Issue of Comm. ACM on Recommender Systems, 40(3).

de Saussure, F. (1959) Course in General Linguistics, trans. Wade Baskin, McGraw-Hill.

Schmidt, K. (1997) 'Of Maps and Scripts: The Status of Formal Constructs in Cooperative Work', Proc. ACM Group 97, 138–147.

Suchman, L. (1987) Plans and Situated Actions: The Problem of Human Machine Communication, Cambridge University Press.

Wegner, P. (1997) 'Why Interaction is More Powerful than Algorithms', Comm. ACM, 40(5):80-91.

Winograd, T. & Flores, F. (1986) Understanding Computers and Cognition, Addison Wesley.

Wittgenstein, L. (1958) Philosophical Investigations (3rd ed.), trans. G.E.M. Anscombe, Oxford University Press.