Humanities Computing: Essential Problems, Experimental Practice[1]

Willard McCarty
King's College London, London, UK

For John Burrows, experimenter

Abstract

The application of computing to the disciplines of the humanities has two principal outcomes: useful results for the field of application and failures completely to demonstrate what is known. These failures, an inevitable feature of modelling, point to the key question for humanities computing, how we know what we know, and so to the beginning of its own scholarly enquiry. This, I argue, proceeds along three branches, the *algorithmic*, the *metatextual*, and the *representational*. Examining the first of these here I argue for research toward an open-ended, interoperable set of primitives based on previous work in the field and designed for the emerging digital library environment. To set the stage for their further development I argue that the field as a whole does not wait on a theoretical formulation of what humanists do, rather should look to the tradition of experimental knowledge-making as this has been illuminated in recent years by historians, philosophers, and sociologists of science.

Correspondence:
Willard McCarty,
Centre for Computing in the Humanities, King's College London, Strand, London WC2R 2LS, UK
E-mail:
willard.mccarty@kcl.ac.uk

1 This paper was first given at the humanities computing seminar of the Computer Science and English Initiative, University College Dublin (March 2000), then as a special lecture at Stanford University and the University of Georgia (April 2000). My thanks go to Susan Schreibman, Michael Keller,

1 Introduction

Although computing has been applied to research in the arts and letters for more than 50 years,[2] systematic effort to understand the consequences and implications from an interdisciplinary perspective is quite recent. Especially within the last decade, the growth of academic centres for humanities computing in the UK, Scandinavia, continental Europe, North America, and Australia has stimulated and supported international debate on the nature, function and institutional role of the new scholarly field.[3]

The following is a contribution to that debate. I begin with the central question, what happens when we apply computing to the humanities? I respond by exploring a provisional research agenda for the field and locating it within the tradition of experimental knowledge-making. In particular I argue that work in the history, philosophy, and sociology of science provides a robust basis on which to conceptualize the epistemology that frames this agenda. I advocate learning from the application of humanistic and sociological methods to the natural sciences in our continuing effort to understand the impact of a scientific device on the humanities.

2 What Happens

In operational terms, when humanities research is computerized the source materials become *data*—that is, computable information—and the research methods resolve into some combination of software and markup.[4] What happens intellectually is neither solely computational nor autonomously human but a combination or interaction of both—a thinking with, and against, the computer. The translation of source materials into data involves two fundamental steps: a choice of what to consider *to be the data* and a perceptual shift required to see these materials *as data*.

In making this choice one rules out not just what one regards as irrelevant—the illustration on the front cover? the text of the colophon?—but also those features of the source that are too difficult or wholly impossible now, or perhaps ever, to compute. Such exclusion raises interesting questions. The perceptual shift goes both ways: one can compare the source before conversion—with all its implicit contexts and the perceiver's tacit knowledge—with the denuded corpus of data. The difference illuminates what has not survived translation; it raises the question, why not? Selecting or designing software for the task in mind is likewise based on the never wholly successful attempt to map the aims of the study onto what software can do. Insertion of those unambiguous declarative statements (i.e. markup) normally involves the same creative conflict between computation and cognition: that which can be formulated in the rigorously consistent and explicit terms of the metalanguage is very different from thoughts about the object of study, whether these are inchoate in the mind or worked out in a 'natural language' argument. It is the difference, the gap between human knowledge and mechanical demonstration, that counts. Humanities computing lives and deals in that gap.

Another way of construing what happens is in terms of *modelling*, by which I mean 'the manipulation of symbol structures so as to bring them, more or less closely, into parallel with [a] pre-established nonsymbolic system, as when we grasp how dams work by developing a theory of hydraulics or constructing a flow chart'.[5] In terms of humanities computing, modelling is an iterative process of constructing and developing something like a computational 'knowledge representation' as this is defined in computer science.[6] In fact we might say that a model is a manipulable knowledge representation. I prefer to use the present participial *modelling* because it emphasizes the provisional, contingent nature of a continuing activity. Although representations, of knowledge or of anything else, can be equally provisional, the present participle usefully prevents us from mistaking what we do as a past participial result, some representation of what has been reached, understood, and so is *known*.

William Kretzschmar, and Nelson Hilton for the opportunities to work on it in public; to Charles Bigelow, John Bradley, John Burrows, Harry Collins, Vanessa Davies, Greg Dening, Susan Hockey, John Lavagnino, Tito Orlandi, Wilhelm Ott, Harold Short, Patricia Wright, and others I cannot (*mea culpa*) recall for help along the way.

2 See the bibliographies by Lancashire (1991) and Adamo (1994); for text analysis, see also the bibliography in McCarty (1993).

3 For an accounting and taxonomy of humanities computing centres, see McCarty and Kirschenbaum (2000); for recent public debate, see the essays listed on my homepage at http://ilex.cc.kcl.ac.uk/wlm/; Unsworth (2000a), the essays referenced there and others by the participants; those by Tito Orlandi, listed on his homepage at http://rmcisadu.let.uniroma1.it/~orlandi/; and the many discussions over the last few years on the *Humanist Discussion Group* (1997–).

4 For the prevalent sense of *data* as specifically computable, see the *OED* s.v. 1d. Here *markup* is taken to mean unambiguous declarative metalanguage inserted into or otherwise associated with primary linguistic data to identify items for processing that software cannot otherwise reliably treat.

5 Geertz, whose words I have quoted, usefully distinguishes this sense of the word, a model 'of' something, from a model 'for' it, 'the manipulation of the nonsymbolic systems in

Because machinery is so often directly involved, a 'model' in my sense is chiefly a physical device, something one runs, tinkers with, obtains results from, although the activity can of course take place in the head, as a thought-experiment. In either case, as it is a surrogate for or simulacrum of the thing studied a model cannot be true. Rather it is a pragmatic 'work of fiction' (Cartwright, 1983, p. 153) designed to reach a truth that is otherwise difficult or impossible to reach, or which one can intuit but not demonstrate. Often it will of necessity be quite crude—hence the term 'tinkertoy modelling' in physics[7]—but its verisimilitude is only a condition of its usefulness, not the point of the exercise. A completely successful model—one which performed exactly as expected—would only demonstrate that the knowledge it instantiated was trivial (in the mathematician's sense); it would be of no interest *as a model* to research. It might be of interest as a process that generates useful results for some other purpose, but properly speaking a model teaches us when it fails to correspond to what we expect of it. The model can fail in two ways, either by not working despite perfect implementation of what we know how to specify, or by working when we think it should not. Such failures drive research forward iteratively by suggesting improvements or pointing to the need for further investigation.[8]

Examples from the sciences will illustrate my point. In one biorobotics project researchers have modelled the walking and climbing behaviour of the *Blaberus discoidalis* cockroach, with benefits both to biology and robotics.[9] The model, Robot III, is well known because of its relative success, but as one of the researchers has said, its failures (e.g. to climb an obstacle of a kind the insect would easily scale) are much more interesting because they tell us what we do not know.[10] In another project mechanical engineers attempting to simulate how humans walk built a model literally from tinkertoys, not intending it to work but to demonstrate certain problems. 'Playing, with no hopes of success' they tried it, it worked and new ideas about walking have followed (Coleman and Ruina, 1998). Testable theories—for example, 'how the mind works' (to quote the title of Pinker (1997))—can also be considered models insofar as their primary value for research lies in what they do not comprehend. Thus Fodor, in his review of Pinker's book, quotes Ecclesiastes: 'the heart of the wise is in the house of mourning' (Fodor, 1998). As the linguist Edward Sapir famously said, 'All grammars leak' (1921, p. 38). Research follows these (well-constructed) leaks, though their often intermediary role means that they tend to be absorbed in the process of building better grammars.

Applied computing is no exception in its experimental application of formal constructs. Text-analysis systems, whether consciously or not, are based on roughly similar models of how meaning arises in a text—or more precisely, where it does. We hardly need reminding of how woefully inadequate, tinkertoy-like these models are, but they are none the less useful, for example, to give empirical focus to what we mean by 'context'. The difference that our machinery makes, which we are only beginning to realize, is that unlike the medieval exegetes who invented the con-

terms of the relationships expressed in the symbolic, as when we construct a dam according to the specifications implied in an hydraulic theory or the conclusions drawn from a flow chart' (Geertz, 1973, p. 93). As I have noted elsewhere, the literature on modelling is extensive but bewildering in its lack of consensus; the usage most companionable with humanities computing is found in experimental physics, where it is applied to the construction of equipment or software thus to model otherwise unreachable things and events (McCarty, 1994, pp. 278–80). On modelling in the sciences, see Groenewold (1961), Leatherdale (1974), Redhead (1980), and Cartwright (1983); in the humanities, see Frye (1991).

6 See Unsworth (2001) for a definition and detailed discussion of the term.

7 The term refers to an originally wooden, now plastic construction set for children, invented in the early twentieth century; see Cartwright (1983, p. 158).

8 Conventionally in the humanities we speak rather of raising questions, or more recently of problematizing. Given the promotional rhetoric of computing, however, and the pervasive misunderstanding of what research is for, the shock-value of 'failure' is needed as a salutary reminder that, as Fowler has said about the commentary, good scholarship 'does not solve problems but makes them worse' (Fowler, 1999: 442).

cordance (McCarty, 1993), we can manipulate as well as use a model of textual meaning. One of the clearest examples in humanities computing is the work of Burrows, whose application of numerical techniques to questions of stylistics and attribution typically proceeds by careful, judicious probing of results.[11]

Positive results have pragmatic value: they attest to a successful technique, which in turn adds to the stock of methodological knowledge. Potentially this technique constitutes an exportable object, of interest in other circumstances and disciplines, and so a valuable commodity by which the interdiscipline makes itself useful to others. But, again, for humanities computing the completely successful technique is in itself trivial—until, perhaps in a new context, it can be made to fail. Then it becomes interesting once again as scholarship.

What do I mean by 'scholarship' in humanities computing? This is the primary question I wish to raise here. The answer is not at all obvious: we cannot simply assume we know what this is. We cannot just import a definition from elsewhere, *because each field defines itself in the nature of its scholarship*; indeed, divergent schools within a single field may look with like incomprehension or disapproval on the work of each other. Philosophers do philosophy, which does not look at all like classical philology; in turn, the classical philologist may well not be able to see what the sociologist does *as* scholarship. Computing humanists must grow their own definition. I have begun already by adapting what I think is a commonly held view in the humanities, if an old-fashioned one, that the point of our work is to raise better questions rather than come up with solutions. Through the discussion of modelling I have laid particular stress on what I think is the commanding role of experiment in raising these questions. After offering a provisional typology for this work and developing a broad example of research, I will devote the remainder of this paper to the ethos of experiment and where we may look for help with it.

3 Kinds of Humanities Computing

Let us divide the work of humanities computing into three fundamental branches: *algorithmic*, *metalinguistic*, and *representational*.[12] The first emphasizes the application of analytical algorithms to source materials; for example, to generate a table of collocation frequencies for a given text or to analyse inflected words morphologically. The second is concerned chiefly with addition of metalinguistic tagging to identify entities software cannot reliably treat, such as chapter divisions, references to persons, or figures of speech. The third, though it also of course uses algorithms, focuses on arranging, formatting, or otherwise transforming the appearance of data. This may be done analytically to reveal patterns, as the KWIC concordance does for language[13] or McGann's deformational filtering for images,[14] or synthetically in the design and construction of electronic scholarly reference works, such as editions, commentaries, lexicons, and the like (McCarty, 2002).

9 This research was featured in the January 2000 issue of the journal *BioScience*; see Reizmann *et al.* (2000). See also the homepage of the Biologically Inspired Robotics Laboratory at http://biorobots.cwru.edu/.

10 James Watson, quoted in Menzel and D'Aluisio (2000, p. 105).

11 On unexpectedly good results, see Burrows and Craig (2001); on results that are surprising because they are better than they should be, see Burrows (1992, pp. 173–4). The clearest statements and demonstration of modelling are in Burrows (2002).

12 Compare McCarty (1999a). Here I concentrate almost exclusively on text. For visual data, see note 14; for music, see the International Symposium on Music Information Retrieval (2000), Online Music Recognition and Searching (2000) (with bibliography), and Roland *et al.* (2001).

13 That the 'keywords [of a KWIC concordance] form a column which guides the eye' was noticed by its inventor Luhn in 1959 (Luhn, 1966); for its consequent uses in research, see Fischer (1966), Preston and Coleman (1978, p. 9), and McCarty (1993, p. 57).

14 See McGann (1997). The broader topic of what is popularly called called 'visual thinking' or 'visualization' is both broad and ill-defined; it includes art history, cognitive psychology, semiology, and design theory (specifically for text, theory of typographic form), among several other

To date the most successful and influential of these branches has been the metalinguistic, principally through the efforts of the Text Encoding Initiative, which has done fundamental research into formal textual structures; research continues in the application of the *TEI Guidelines* and in the development of the Extensible Markup Language (XML).[15] The presentational branch, active for the last several years in hypertext research, has more recently expanded under the name of 'new media studies'.[16] Here, however, I will focus on the algorithmic branch as a means of approaching the question of our activity as a whole. Most of what I say about that is, however, applicable to all three.

The algorithmic branch of humanities computing has, to put the matter bluntly, not done as well from the perspective of the individual scholar.[17] In practice, humanists usually cobble together what they can out of the ragbag of commercial and freeware applications, with academic software added in when anything at all is available for a given area of research. Rarely do humanists have something written for them, even more rarely write it themselves. In recent years, despite scattered proposals, no comprehensive initiative in software development for humanities research has progressed much beyond the discussion stage.[18] There are several obvious practical reasons, but here I would like to concentrate on the scholarly problems that I think need attention before the more worldly matters can sensibly be addressed.

Some years ago, in one such initiative for a new text-analysis system, Sperberg-McQueen (1996; see Simons, 1998) argued that 'we need an open, extensible system.... [whose] architecture, if we insist on calling it that, will be an emergent property of its development, not an *a priori* specification. We are not building a building; blueprints will get us nowhere. We are trying to cultivate a coral reef; all we can do is try to give the polyps something to attach themselves to, and watch over their growth'. Since then the wisdom of his argument for independent but coordinated development has become particularly obvious, especially in the now unavoidable context of the digital library, which has not so much redefined our goals as clarified them. I am not referring here to the electronic analogue of local collections nor to the kind of thing exemplified by the Perseus Digital Library,[19] whose work is centrally important to our realization of it. Rather I am referring to the idea of a singular, world-wide entity—a heterogeneous, geographically unconstrained working environment of mutually compatible data and software to which independent, otherwise uncoordinated efforts contribute.[20]

In important respects the digital library is not at all a radical departure from the dual origins that the term identifies. Its model for aggregation of knowledge is in essence the ancient idea of the research library: a large number of modular resources in more or less standard format that a variety of readers with unforeseen purposes may combine and recombine *ad lib*. As a digital entity it also builds on a broader trend in computing hardware and software toward this recombinatorial liberty. The development of hardware has, as we all know, put computers directly into the hands of individuals. At the same time, the graphical user-interface

fields. For art history (with aspects of psychology), see, for example, Arnheim (1969); for psychology, Wright (2000); for graphical semiology, Bertin (1967); for design theory, see Tufte (1983, 1990, 1997) and the *Genre and Multimodality* project (2000) (with bibliography); for typographic theory, Baudin (1994); for software design, Petre *et al.* (1998).

15 For the TEI, see *Text Encoding Initiative* (2000); for a popular account of XML, see Bosak and Bray (1999) with further readings; see also Connolly (1997).

16 For hypertext research, see McCarty (2000); for new media studies, Manovich (2001).

17 The problem for the presentation of encoded text is noted by Ott (2000a, p. 93).

18 Note the *Elta Software Initiative* (Horton *et al.*, 1998), the Text Analysis Software Planning Meeting at Princeton (Sperberg-McQueen, 1996), and the Text Software Initiative (Ide and Veronis, 1993); note also the closely related *Building Blocks Project* (NINCH, 2000), which has ambitions to go further than earlier discussions, such as Bakewell *et al.* (1988) and Bearman (1996), to create an actual research agenda.

19 See the *Perseus Digital Library* (especially under 'Publications') and Smith *et al.* (2000).

20 On digital library research, see McCarty (2000, III.D); as the entries there make clear, the idea of the digital library tends to reflect a strong

mediating their application has provided a visual metaphor of recombinatorial space and so provided impetus to conceptualizing individual programs as components. Software developers now argue for and and work with 'open' architectures of reusable components (although this openness is, of course, severely constrained by commercial interests). The Web, which Smith (1997, p. 240) has identified as key to the infrastructure for computing and communications in the future, provides an indefinitely expanding and various source—indefinitely extensible library shelves—from which to populate the individual 'desktop'.

In other words, that which we might call in computing terms *the end-user realization of modelling* is helping to realize a deeply familiar working environment. With that familiarity comes great potential for the humanities that hardly a moment's reflection is required to grasp.

Admittedly the problems facing the realization of such a library are formidable, perhaps least of which are the technical ones.[21] Among these, the ability of independently developed components successfully to interact with each other—the unsolved problem of 'interoperability'—is the central one.[22] Although we must not overlook these problems or treat them lightly, we also cannot afford to be stopped by them. Hence in the interest of framing our agenda—those things now imaginable to be done—let us put the world-wide digital library into the future-perfect tense, as McGann (1997) recommends. That having been done, it becomes obvious that the digital library provides the natural if not mandatory context in which humanities computing research (and much else) is to be done. This means primarily thinking henceforth in terms of interoperable components that implement broadly applicable functions. Again, this is not a new idea—it has been around at least since the Unix Toolbox[23]—but its scope and urgency of application are.

In humanities computing perhaps the most notable effort in the direction of interoperable components is the Tübingen System of Text-Processing Programs (TUSTEP), 'a professional toolbox for those academic fields where texts are the object of research'.[24] The problem—better, the research opportunity—presented by TUSTEP as by other self-contained systems lies in the distinction currently made for the digital library between two kinds of online component-based systems: *distributed* on the one hand and *heterogeneous* or *federated* on the other. Although TUSTEP is not an online system, like the former kind it is a finite collection of components 'carefully designed to work with each other', in contrast to the emerging notion of 'cooperating systems in which individual components are designed or operated autonomously' (Paepcke *et al.*, 1998, p. 33). Autonomous operation is exactly what interoperability requires.

In the language of systems theory, such components are known as 'primitives', that is, the lowest-level building blocks of the system;[25] by definition, interoperability specifies an open-ended system. Our task, then, is not to define all the relevant primitives, rather a sufficient number to allow actual research and a means for defining others. Depending on the kind of scholarship, this will involve to varying degrees the widely

interest in and the potential for collaborative work and user-centred design, which in turn leads to user-built configurations of resources.

21 See the extensive discussion of these problems in McCarty (2002) and the references.

22 See Fox and Marchionini (1998) and especially Paepcke *et al.* (1998).

23 For the application of Unix tools to humanities research, see Orlandi (1990); more generally, Kernighan and Plauger (1976).

24 Ott (2000a, p. 93); see also Ott (2000b).

25 Dawson and Medler (1997), s.v. 'primitive'; see also 'functional analysis'; discussion in *Humanist*, vol. 14, in several numbers from 258 to 408, s.v. 'primitives'. Unsworth (2000b) is discussed below.

dispersed scholarly literature illustrating or discussing methods of analysis, centrally including but not limited to the literature of humanities computing,[26] and two crucial sources of tacit methodological knowledge: the analytical genres, such as the lexicon or commentary, and computing systems designed for research purposes.[27] Extracting this knowledge is a complex subject for which I have no space here.[28] Suffice it to say that the raw material is abundant, but gathering and forming it into a cogent set of primitives is a major project.

4 Defining the Primitives

An example from everyday life will help illustrate the nature of primitives and the kind of analysis that identifies them. Imagine the world of a builder, which we might describe as a loosely organized system of objects and actions. For the sake of simplicity, consider only the objects: houses, garages, walls, ceilings, floors, timber, paint, glue, glass, bricks, and so on. Within this system we cannot call 'house' a primitive, as it is composed of other objects within the builder's world, for example, walls, ceilings, floors, bricks, timber, glass, paint, etc. 'Brick', however, is a primitive, as the builder acquires bricks ready-made; similarly the timber, glass, and paint.

The implicit analogy between builder and scholar quickly breaks down if pressed—the builder's world is defined by physical, social, and economic constraints not applicable to the scholar's. The point of my analogy is only to raise the analytic question of the scholar's primitives, or more precisely what these primitives might be in the computational environment of a digital library.

Two approaches to this question have been identified. We can begin, as Unsworth (2000b) has proposed, from the top down, identifying as the scholar's 'irreducible currency' a set of fundamental abstract operations that describe what humanists intentionally do and that apply across the genres, periods, and theoretical approaches. His provisional set includes 'discovering', 'annotating', 'comparing', 'referring', 'sampling', 'illustrating', and 'representing'. From these, he argues, we should in principle be able to develop interoperable software tools exchangeable 'across all manner of boundaries of type or token'. At the other end of the spectrum, we can work from the bottom up, as in TUSTEP if not at an even lower level, implementing the discrete mechanical actions that scholars are actually observed to take in the performance of scholarly work or that the computational environment itself requires. This broadly sociological approach is conventional within computer science and software engineering circles, where ethnographic methods of observation *in situ* are now commonly used (e.g. Wixon and Ramey, 1996).

The appeal of the former is that the high-level primitives are the recognizable as well as intentional actions of scholarship, independent (so that argument runs) of accidentals. The appeal of the latter is based on the fact that individuals are seldom if ever reliable informants about their own actions. Observation of what they actually do allows us to

26 The bibliographic problem is immense; see McCarty (1999c, n. 5).

27 For textual studies this includes, for example, TUSTEP (see note 24), TACT (Wooldridge, 1991; Lancashire et al., 1996), and the Oxford Concordance Program (Hockey and Martin, 1987; Hockey, 1988).

28 See my discussion of artefactual analysis for the commentary form in McCarty (2002).

obtain tacit knowledge of the habitual and (so this argument runs) below the level at which theoretical commitments operate. But we are not forced to choose: both may be used in what Bateson describes as 'a sort of pincers maneuver' on reality. 'In scientific research', he remarks, 'you start from *two* beginnings, each of which has its own kind of authority: the observations cannot be denied, and the fundamentals must be fitted' (Bateson, 2000, p. xxviii).

The central question, we might say, is at what level of detail or 'granularity' to locate the primitives—just how primitive to make them. If universal formalisms for 'comparing', etc., can be written, then we have our answer. Otherwise—more or less as now, or in something yet to be devised—our primitives will need to be statements in one or more high-level programming languages. This in turn implies either that scholars will need to become programmers as a matter of course or that the practice of scholarship will need to be much more collaborative than it now is.

A cogent argument establishing the possibility of such universal formalisms has yet to be made. In fact we have serious reasons to doubt that it can, even at a level considerably lower than 'comparing', etc. Consider three examples of quite low-level yet common scholarly functions: alphanumeric sorting; compiling a frequency list of word-forms; lemmatizing the word-forms of an inflected language. As Ott (2000c) has argued, experience to date suggests no way of lemmatizing in software without allowing for arbitrary human intervention, and sorting can only be considered a primitive if we accept a routine that treats its input as character-strings rather than word-forms. Such a routine could not, for example, distinguish between letters and numbers, treat double-character combinations (such as 'LL' in Spanish or 'NG' in Welsh) as single letters, nor, as Ott points out, conform to the conventions for various kinds of indexes. Furthermore, compiling a frequency list is dubious for the same reason as sorting when lists ordered alphabetically are needed.

The project also runs into philosophical difficulties, which are revealed if we ask what claim we might make for such universal formalisms as well as for the ways of combining them. The most ambitious claim would imply one-to-one mapping of thought-forms onto algorithms, in other words a version of the 'strong AI' claim to assimilate human patterns of action to formal mathematical structures. This is what Hofstadter (1995) calls the 'Boolean Dream', which for us implies 'a law-bound vision' of a finite mechanical means for autonomous construction of scholarship. Truly labour-saving! Apart from the bone-chilling dread this vision should inspire, the trouble with it is that it is not happening, and we can see no way for anyone to make it happen. It is flawed in principle—forgetful, as Agre (1995) argues, of the philosophical tradition from which it comes: it ignorantly makes 'an effort to domesticate the Cartesian soul into a technical order in which it does not belong'. Thus, he points out, 'a pervasive and ever more clear pattern of technical frustrations' results. Agre concludes, as I do about the scholarly value of humanities computing, that these frustrations *are* the point; they point to the value of failure in a recursive 'cycle of formalization', in which diagnosis of the technical impasses would reveal difficulties with the underlying conception and so initiate 'new and more informed rounds of formal modelling. The privilege in this cycle', he points out, 'does not lie with the formalization process, nor does it lie with the critical diagnosis of technical impasses. Rather, it lies with the cycle itself, in the researcher's "reflective conversation with the materials" of technical and critical work.'

Dennett (1984) reaches much the same conclusion about artificial intelligence research. Discussing the 'frame problem' in AI (roughly, as I understand it, the profound dependence of real-world knowledge on context), Dennett argues that it is far more than a technical embarrassment or a curious puzzle. Rather it is 'a new, deep epistemological problem—accessible in principle but unnoticed by generations of philosophers—brought to light by the novel methods of AI, and still far from being solved' (Dennett, 1984, p. 148). In his view, then, AI is philosophy continued by other means. For my purposes it does not matter at all whether he is right about the 'frame problem' or about philosophers; my point is made by his heuristic use of AI to drive particular sets of formally articulated statements and the assumptions they embody until the model they constitute breaks. In other words, to quote Agre, AI offers 'a powerful way to force an idea's internal tensions to the surface through prolonged technical frustrations'. Here I quibble: the manifestation of an idea in software is *not* the same idea but another, closely related one, more accurately a model of it, from which with care we learn much about the original idea—but let that pass. I also leave to the philosophers and their colleagues in AI the question of whether they need each other. All I wish to borrow from Dennett and Agre is the shared point that AI can be viewed as mechanized philosophy, and that from the perspective of the humanities the *scholarly* value of doing mechanized philosophy results from the problems it turns up, not from the successes. Furthermore, I think that roughly the same is implied by Newell and Simon (1976, p. 105–6) when they describe computer science as an experimental practice, with deeper understanding of nature as the point and improved techniques the byproduct. More about that later.

The opposition to be understood, then, is not computer science or AI versus the humanities—they seem, actually, to have some important interests in common—rather it is the instrumentalist view of computing, directed toward making things work, versus the scholarly one, which prospers when they do not. In isolation, the instrumentalist view results either in the humanist's dismissive sneer that the computer is 'just a tool' or in the technologist's falsely optimistic confusion of the realizable goal for the quest—hence the extravagant claims that full realization of the quest is imminent, as Simon and Newell infamously proclaimed in 1958 (p. 6), or that it is at least soon, as Minsky declared in 1967. The problematic habit of mind I identify—the overpowering lust for solutions—betrays itself not just when an emperor of our realm imagines unreal clothing but also when he or she confuses a model with the reality it approximates. Pinker's engagingly written book, *Words and Rules* (1999), provides a good example. What he depicts is more like a cartoon than a photograph, and though it may be only a matter of time before our cartoons become indistinguishably photographic, still they will not be true, and that will still be the point.

Earlier I suggested observation of what scholars actually do, as if this were a straightforward task. But as the sociologist Harry Collins demonstrates (1995), observation of human actions is also not nearly as simple as might seem. The basic problem he identifies is that although many human actions are *regular* (walking is an example), most of these are too dependent on an ever-changing social context to map one-to-one onto observable behaviour; an infinite regress of rules would be required to cover all situations. *Behaviour-specific* actions (military marching is an example) are perfectly mechanizable, but comparatively few actions, he notes, are like these. Unfortunately, we cannot come up with a universal toolkit just by separating out behaviour-specific from regular scholarly actions because, Collins observes, many actions can be executed as either kind depending on intention and desired outcome. Hence it is not possible consistently to tell by observation or introspection which is which.

Ryle illustrates the point by asking how we might tell solely by observation the difference between an involuntary twitch and a deliberate wink, which itself can be a fake wink, a parody, a parody of a fake, and so on.[29] The same applies, for example, to mechanisms of scholarly reference, especially the infamous *cf.* (L. *conferre*, 'compare), which seems purely behaviour-specific but whose intention may be quite different from a reasoned imperative to consult the listed examples; indeed, it may not be possible to figure out the purpose behind the reference. Hence a straight hypertextual mechanization of a set of references could easily pervert the author's intention.[30]

According to Collins's argument, then, we are left with two projects and a question. The projects are first to do the sociological analysis to identify the behaviour-specific actions in humanities research, then to write the software modules this analysis defines. Let us put these projects on our agenda, but for the sake of argument I will assume here that they are done. The question that remains—a somewhat clearer form of the one with which I began—is how we characterize the actual practice of humanities computing implied by Collins's argument: an unpredictable, always varying combination of behaviour-specific and regular scholarly actions. How do we define a coherent scholarship in such a higgledy-piggledy practice and so come up with a shareable set of primitives?

6 The Analogy of the Experimental Sciences

The beginning of an answer is: we look around for analogous patterns in older, better established research activity and ask how scholarship is defined there. Not surprisingly, we find an intriguing analogy with the

29 Ryle quoted and amplified in Geertz (1973, p. 6f).

30 I discuss this example extensively in McCarty (2002).

experimental sciences, in which, like us, researchers work in laboratories with equipment, collaborate across areas of specialization, use empirical methods on data, and employ numerical analysis. The immediate question is, of course, how far this analogy can be pushed. Does it show us anything of substance about humanities computing, and if so, what?

This is certainly not the first time someone has suggested an analogy with the sciences. Yet curiously we seldom if ever take the analogy with science further, except in those disciplines (such as linguistics and history) with overlap into the social sciences. In the humanities it lurks mostly unexamined in the background, occasionally invoked to distinguish humanities computing from its help-desk origins or to give an insecure colleague the illusion of certainty. Here, however, I propose a rather more sustained look at the idea, and specifically the idea of *experimental* rather than theoretical science.

I recommend the sciences to your attention with no particular reverence for the scientific as a privileged mode of knowing. The word 'science', Searle (1984, p. 11) points out, 'has become something of an honorific term'; he would do without it if he could. I agree; I think we should shun it in naming what computing humanists do, certainly in the Anglophone world.[31] At the same time, we cannot afford to ignore our natural kinships wherever they may be. Close studies of laboratory practice, the grubby bits of science—which the Göttingen number theoretician Edmund Landau once contemptuously called the *Schmieröl* (engine-grease) of knowledge (Galison, 1997, p. xvii)—in fact help greatly to counteract this anti-intellectual reverence by illuminating the uncertainties and makeshift arrangements, the fudging, the historical contingencies, the role of hunch and play. They help 'bring science to the same epistemological level as other knowledge-making activities', as Collins remarks (1992, p. 185). '[W]hat we are all aiming at in intellectual disciplines', Searle (1984, p. 11) declares, 'is knowledge and understanding. There is only knowledge and understanding, whether we have it in mathematics, literary criticism, history, physics, or philosophy'. Such a levelling of the playing field is very important for proponents of humanities computing, who in their professional insecurity may be tempted to take to some theoretical high ground, wash off the *Schmieröl*, and so, I would argue, lose their reason for being.

My focus, then, is on experimentation, not the sciences as such.

In the quarter-century since Newell and Simon (1976, p. 105) cautiously suggested that the core activity of computer science is experimentation, we have come to know a great deal more about what this means.[32] The philosopher Ian Hacking has noted that until the 1980s theory dominated historical and philosophical thinking about the sciences (Buchwald, 1995, p. 1). As a matter of course, experimentation was assumed to be entirely subservient to theory, and replication of experimental results to be a simple matter of honestly following instructions under essentially identical circumstances. Close inspection of actual practice has utterly changed our view of all that, however.

We now know that, as Hacking (1983, p. 150) says, '[e]xperimentation

31 I might not avoid the cognate terms if I were writing in another European language and so in a different cultural context, however. Similarly, in English, it seems to me, the term 'informatics' has all the wrong associations, suggesting a highly specialized, technical–theoretical study such as information science, identification with which would pull us away from the life-blood of interdisciplinary collaboration; see McCarty (1999b).

32 In addition to the works I discuss here, see Cartwright (1983), Achinstein and Hannaway (1985), Latour and Woolgar (1986), Gooding *et al.* (1989), Buchwald (1995), and Pickering (1995).

has a life of its own', that experimentation is genuinely epistemological on its own grounds. Thus the intellectual excitement of recent work in the history, philosophy, and sociology of science: for example, in historian Peter Galison's unifying focus on 'all that grubby, unplatonic equipment' at the centre of the disparate bundle of activities involved in nuclear physics (Galison, 1997, p. xvii); in Hacking's philosophy of experimental knowledge, springing from 'the remnants of an historical process of becoming and discovering' brought to light when about 40 years ago philosophers unwrapped 'the mummy of science' and precipitated a crisis of rationality (Hacking, 1983, p. 1); and in Collins's wide-ranging demonstration of the social artisanship by which science is constructed within networks of researchers (1992).

Facts, we have learned, do not speak for themselves; they are *facta*, 'things made' by people within overlapping historical, philosophical, and social contexts.

Here at my peril I skirt hotly contested issues in the culture wars fought over whether knowledge is socially constructed; I can perhaps do no better than defer to Hacking's very recent book *The Social Construction of What?* (1999), which he calls a 'primer for noncombatants' in these wars. I run considerable risk of grossly oversimplifying quite subtle arguments and the complex historical and social circumstances they involve. I compound the dangers by confining myself to the work of only three people out several labouring in the confluence of their three vital disciplines. Yet—this is my only excuse—my purpose here is limited to attracting attention to work that, I am convinced, computing humanists must be aware of in order to see the integrity of what they do and its place within the broader intellectual landscape. Much help is at hand if you know where to look.

In the following section I will illustrate three major sources of such help by sketching briefly how the three scholars I have mentioned depict experimental practice as it appears in their respective disciplines: the history, philosophy, and sociology of science.

7 History

In his magisterial study, *Image and Logic: A Material Culture of Microphysics* (1997), Galison shows us in detail how individuals engaged in a bundle of rather disparate activities came together and successfully collaborated to solve the most difficult experimental problems of their day. His historiography is centred on the machinery these activities had in common (e.g. the devices invented to study subatomic interactions) and the users' intellectual and pragmatic negotiations of knowledge across disciplinary boundaries. He proposes and develops an anthropological–linguistic analogy for these negotiations, which he calls the *Trading Zone*.[33] In a typical trading zone people from mutually incomprehensible cultures come together to trade objects of interest. To do that, as in fact happens, they develop a highly restricted proto-language or *pidgin* for their negotiations. This pidgin allows them to reach

33 See Galison (1997): Chapter 9 especially, but also Chapter 1 and Sections 3.11, 6.7, 6.11, 8.1, 8.5, and 8.7.

agreement among themselves about objects of trade even though outside of the zone, within their own cultures, their understandings and uses of these objects differ radically.

Galison's analysis of collaborative physics is easily generalizable to our situation, likewise equipment-centred, likewise involving researchers with very different points of view coming together over the equipment. His trading zone metaphor brings to our notice and encourages us to value the unifying methodological perspective on the humanities that common machinery affords. By extension, it defines the role of the person in whose professional care this perspective resides—the computing humanist—as dealer in and adaptor of methodological goods. Teaching and various forms of collegial service have shown me that the methodological goods in each of the humanities are in fact a portable wealth, although like the techniques Galison discusses they tend to be understood in very different ways in the sometimes radically different contexts in which they are used. These goods migrate easily across the disciplinary boundaries because they depend essentially on data structures and types (discursive or tabular text, image, sound) not on interpretative context or meaning. Thus, for example, the historian, literary critic, and linguist interested in the language of their respective texts (or the same text) may all use the same methods of text analysis in the same ways, though their questions and the conclusions they draw are radically different.

The skills required of the methodological trader are both social and intellectual. The former raise essentially administrative questions about how to create the circumstances in which an academic culture will foster such interchanges across disciplines and among the departments where they live. Not a simple matter, especially given the socio-political realities of intellectual 'turf', the lack of money for disciplinary experimentation, and (where it still exists) the implications of tenure. Among other things, a considerable literature on interdisciplinarity awaits our reading and inward digestion (especially Newell (1998)). I will say very little about such matters here because so much depends on national and local conditions. The best we can do internationally is to make the intellectual case for humanities computing; the administrative imagination and political will required are local matters.

The intellectual skills needed by the methodological trader point to curriculum—what we want the next generation of professional computing humanists to know. One of course assumes a high level of technological competence; what is at issue, however, is the wide-ranging methodological knowledge and deep familiarity with research that a scholarly approach to applied computing requires. Space prevents me from saying much more about curriculum here, but by implication I am arguing for inclusion of the history, philosophy, and sociology of science at the root level of a (post)graduate programme. The metaphor of the trader proves insufficient to cover these intellectual skills—unless, that is, we understand his or her wily cunning as an analogue for the computing humanist's analytical insight and cogent argumentation. Thales the philosopher would have understood this (Aristotle, *Politics*, 1259a).

8 Philosophy

Let me begin by repeating an important point Hacking makes in *Representing and Intervening* (1983). He notes that 'Philosophers of science constantly discuss theories and representations of reality, but say almost nothing about experiment, technology, or the use of knowledge to alter the world … Philosophy of science has so much become philosophy of theory that the very existence of pre-theoretical observations or experiments has been denied' (Hacking, 1983, p. 149f). Only within the last 20 years or so, that is, has experimentation become philosophically visible and its history emerged out of the shadow of theory. For an applied field such as humanities computing, in which widespread practice comes long before articulation of anything we could call a theory of the field, Hacking's (1999, p. 199) 'ecumenical descriptive epistemology' of experiment offers a welcome beginning. His argument suggests that we are not necessarily moving toward such a theory. It is even to be doubted whether, in general, fields of enquiry need a unified or unifying theory at all, to develop and prosper.[34]

Let me summarize very rapidly what Hacking does. He sets the scene for a philosophy of experiment by showing the insufficiency of representation as a basis for scientific reality. Instead, '[w]e shall count as real', he declares, 'what we can use to intervene in the world to affect something else, or what the world can use to affect us' (Hacking, 1983, p. 146). Citing a number of historical cases, he demonstrates that 'no statements of any generality are to be made' about which comes first, theory, experiment, technology, or whatever (p. 166), and he documents how history has been rewritten in certain cases to put theory first when it was not there at the time (pp. 160–1). He shows how through experiment the reality of once hypothetical entities, such as the electron, has been established by our ability to manipulate them. Experiment, in other words, means knowing by doing. Using the example of the microscope he shows how experimental observation involves active, interventionist skill: thus we are in part 'convinced about the structures we seem to see because we can interfere with them in quite physical ways' (p. 209). Similarly, he takes apart the notion of theory, distinguishing the separate and integral role played by the adjustments we commonly make to bring theory into accord with the world (p. 214). This gives formal recognition to the otherwise illegitimate-seeming compromise of pure ideas with laboratory realities.

Hacking's view gives a prominent epistemological role to the unregulated or unexpected. Thus the importance of play. Modelling (with its strongly representational element) involves it, but he has in mind an additional trajectory, away from the testing of a pre-existing idea to discovery of the unexpected. Thus observation counts for very little; what matters is being observant and so able to catch 'the instructive quirks or unexpected outcomes' (Hacking, 1983, p. 167). *Sitzfleisch* is as necessary as ever, but insights come unexpectedly when rigorous procedures bump up against anomalous cases—when, that is, we are alert enough to see

34 This is what Culler calls the 'myth of foundationalism' in English studies, for which see Culler (1988, Chapter 2).

instances *as* anomalous. An experiment may falsify a pre-existing theory, as Popper said it should, but that is not its sole function. We need not aim narrowly at instantiating, confirming or refuting theories, however comfortable that would seem, although an experiment may sometimes do just that.

Hacking's analysis of what it means to see with a microscope gives us reason to explore the same question: can we be realists about the knowledge we gain because we manipulate entities to probe further into the unknown? We know, for example, that a concordancer shows us textual detail by transforming the text;[35] that collocation is an idea resulting from that transformation; and that by learning to manipulate collocations and so taking collocation on as a textual reality, we probe further into language. Yes, it does seem we are realists as experimenters are realists.

9 Sociology

Finally, what do we learn about experiment from sociology?

In his book *Changing Order: Replication and Induction in Scientific Practice* (1985, 1992), Collins notes that the '[n]ew and timorous sciences, such as the whole range of social sciences'—and with them I include activities comparable with the sciences, such as humanities computing—'have tried to develop by apeing what they believe to be the method of the natural sciences'. This is of course a false picture: we know now, as Hacking (1999, p. 198) says, that '[t]here is no one scientific method; the sciences are as disunified in their methods as in their topics'. The idea of a unitary method arises from what Collins (1992, p. 159) calls the 'the "algorithmical model" ', which 'encourages the view that formal communication can carry a complete recipe for experiment with all that follows'. According to this model, experimenters are essentially skilled rats in a behaviourist's maze, or as Collins imagines, functions of a computer-like research machine called 'science'. In the first part of his book (pp. 29–49), Collins adopts Douglas Adams's metaphor of the Earth as a very large computer from *The Hitch-Hiker's Guide to the Galaxy* to show that science does not and cannot work like that.

Using a number of case studies ranging from construction of a laser through detection of gravity waves to experiments in the paranormal, Collins focuses on how during a scientific controversy facts are made within the social group of concerned allies and enemies he calls the 'core set' (1992, pp. 142–5). He argues that in practice, during such a controversy, the ideas of scientific *inference* (e.g. that tomorrow will be like today) and *replicability* (that we can demonstrate this) turn out to be not at all straightforward. The negotiations within a core set may include many forms 'not normally viewed as belonging to science' when on purely scientific grounds 'experimenters' regress' cannot be broken—that is, the common situation in which establishing absolute validity of experimental results requires valid replication of the experiment, which in turn is established by achieving the results to be tested, and so forth *ad*

35 Although it is true that concording antedates humanities computing by roughly 750 years, concording did not become an experimental activity, with focus on transformations of the text, until interactive concordancers became available; see McCarty (1993).

nauseam. Exactly how a case is resolved, Collins argues, depends on the degree to which the putative discovery threatens to disturb 'the conceptual web' of the field or fields involved. Parapsychology, for example, 'threatens too much to too many to be easily acceptable. That is why its more uncompromising proponents are forced to live in a world of their own' (Collins, 1992, p. 139). Thus '[t]he core set,' as he describes it, ' "funnels in" social interests, turns them into "non-scientific" negotiating tactics and uses them to manufacture certified knowledge' (p. 144). In other words, '*the objects of science are made by hiding their social origins*' (p. 188).

Prying into these origins certainly relativizes science, but this is visible only within the core set, or as he says, '*[t]he degree of certainty which is ascribed to knowledge increases catastrophically as it crosses the core set boundary in both space and time*' (Collins, 1992, p. 144). Interestingly for our purposes, the idea of the core set implies a much shorter distance than we might have imagined between knowledge of the natural world and knowledge in the humanities, which are certainly no stranger to social construction and have no 'mantle of infallibility' to be divested of (Collins, 1992, p. 159). Few in the humanities have dreamt that theories are to be tested through observation of neutral facts by unbiased observers; seeing that this is not true of the sciences either brings us closer, recommends more strongly the analogy of experimental science. It also, of course, raises the question of where the differences lie between the sciences and the humanities, which for my purposes is thankfully the simpler question of what these differences are at the level of data and software.

10 **Conclusion**

The issues and arguments I raise here undoubtedly raise hackles, and I am acutely conscious of being an untrained neophyte, clumsily fingering the subtle goods of three well-explored realms. I hope, as Hacking (1983, p. xv) says, that the issues and arguments are nevertheless 'clear enough and serious enough to engage a mind to whom they are new, and also abrasive enough to strike sparks off those who have been thinking about these things for years'. It does seem to me that much more effort is required to clarify the arguments in the history, philosophy, and sociology of science *with respect to each other*. That is perhaps being fought out now, but we who would take up all three and apply their analytical tools to humanities computing need especially to understand how they interrelate.

I will conclude by summarizing why I think the effort is worth the risk. What, in other words, does the analogy of experimental science do for us?

First of all, through Galison's metaphor of the trading zone, it gives depth of meaning to the often *ad hoc* and ill-understood practice of collaboration in the computer-assisted humanities: it helps us to understand how people from very different fields can do laboratory-based research together well—do research better than alone, or at all—while remaining undiluted intellectually. Experimental science suggests the importance of our material culture as computing humanists, in particular its unifying perspective, which brings together a rather different assortment of disciplines and skill-sets than books and libraries did and continue to do. Again thanks to Galison's metaphor, experimental science throws light on the humanities computing specialist as methodological trader and demonstrates the portability of the goods in trade. By analogy, it links the principal disciplines that make up science studies to the curriculum he or she needs to follow in order to understand and guide this new laboratory practice. It suggests that the disparate, theoretically disunified activities that happen under his or her aegis make up a genuine epistemological practice and so, as happened to experiment itself, give us a language and a set of analytical tools with which to discuss humanities computing in practice. Experimental science suggests a complex relationship between computing and theory in the humanities. It raises the question, do we in fact also hypothesize entities, then establish their reality by manipulating them, as I have suggested for concording? The analytical tools we borrow reveal philosophical complexity in the data we have, suggest that we act on our artefacts to know them. The socio-political dimensions of knowledge-making in the humanities hardly come as a surprise, but it is most salutary to be shown that they are as crucial to a scientific approach as to any other.

How, then, does knowledge-making in the humanities differ from knowledge-making in the sciences? Are we, as Rorty recently said in the *London Review of Books,* within sight of 'the end of the epistemic wars' between C.P. Snow's 'two cultures' (2000)? Galison's (1999) report on his historical research into representations of objectivity from the late eighteenth to the early twentieth century demonstrates well enough that scientific reality is not simply 'out there', and our own work that reality for the humanities is not simply 'in here'. A convergence from both sides into the gap between the demonstrable and the beautiful[36] suggests that the analogy I recommend will work. Perhaps what humanities computing will turn out to be is the beginning of an epistemic love-making, with much joy and many children.

But what about the primitives? My argument drives toward the conclusion that they are to be discovered pragmatically, gradually, by experimentation—that we do not necessarily have to wait for or look to an overarching theory or satisfactory formalization. We have enough experience with computing now to make a start, in the (future-perfect) context of digital library research, to imagine a basic set of interoperable primitives. Once our colleagues can play with these like Lego pieces, others will emerge out of tacit knowledge into the reality of the manipulable.

36 I refer to the intangible but incorporated quality in both humanistic and scientific work; for the latter, see Friedman's brief discussion of 'the implicit assumption that the laws of nature are elegant and beautiful' (1999, pp. 12–14).

References

Achinstein, P. and Hannaway, O. (eds) (1985). *Observation, Experiment and Hypothesis in Modern Physical Science*. Cambridge, MA: MIT Press.

Adamo, G. (1994). *Informatica e discipline umanistiche. Vol. 5. Bibliografia di informatica umanistica*. Rome: Bulzoni Editore.

Agre, P. E. (1995). The soul gained and lost: artificial intelligence as a philosophical project. In Güzeldere, G. and Franchi, S. (eds), *Constructions of the Mind: Artificial Intelligence and the Humanities. Stanford Humanities Review*, **4**(2): 1–19. Available: http://www.stanford.edu/group/SHR/4-2/text/agre.html

Arnheim, R. (1969). *Visual Thinking*. London: Faber and Faber.

Bakewell, E., Beeman, W. O., Reese, C. McM., and Schmitt, M. (eds) (1988). *Object, Image, Inquiry: The Art Historian at Work*. Los Angeles, CA: The Getty Information Institute.

Bateson, G. (2000). *Steps to an Ecology of Mind*. Chicago, IL: University of Chicago Press.

Baudin, F. (1994). *L'Effet Gutenberg*. Paris: Ed. du Cercle de la Librairie.

Bearman, D. (ed.) (1996). *Research Agenda for Networked Cultural Heritage*. Santa Monica, CA: Getty Art History Information Program.

Bertin, J. (1967). *Semiologie Graphique*. Paris: Mouton/Gauthier-Villars. Trans. *Semiology of Graphics: Diagrams, Networks, Maps*. Madison, WI: University of Wisconsin Press, 1983.

Boden, M. A. (ed.) (1990). *The Philosophy of Artificial Intelligence*. Oxford Readings in Philosophy. Oxford: Oxford University Press.

Bosak, J. and Bray, T. (1999). XML and the second-generation Web. *Scientific American* (May): 79–83. Available: http://www.sciam. com/1999/0599issue/0599bosak.html.

Buchwald, J. Z. (1995). *Scientific Practice: Theories and Stories of Doing Physics*. Chicago, IL: University of Chicago Press.

Burrows, J. (1992). Computers and the study of literature. In Butler, C. (ed.), *Computers and Written Texts*. Oxford: Basil Blackwell, pp. 167–204.

Burrows, J. (2002). Questions of authorship: attribution and beyond. *Computers and the Humanities* (in press).

Burrows, J. and Craig, H. (2001). Lucy Hutchinson and the authorship of two seventeenth-century poems: a computational approach. *Seventeenth Century*, **16**: 259–82.

Cartwright, N. (1983). *How the Laws of Physics Lie*. Oxford: Clarendon Press.

Coleman, M. J. and Ruina, A. (1998). An uncontrolled walking toy that cannot stand still. *Physical Review Letters*, **80**(16): 3658–61. Available: http://prl.aps.org/

Collins, H. A. (1992). *Changing Order: Replication and Induction in Scientific Practice*, 2nd edn. Chicago, IL: University of Chicago Press.

Collins, H. A. (1995). Humans, machines, and the structure of knowledge. In Güzeldere, G. and Franchi, S. (eds), *Constructions of the Mind: Artificial Intelligence and the Humanities. Stanford Humanities Review*, **4**(2): 67–83. Available: http://www.stanford.edu/group/SHR/4-2/text/collins.html

Connolly, D. (1997). *Extensible Markup Language (XML)*. W3C. Available: http://www.w3.org/XML/

Culler, J. (1988). *Framing the Sign: Criticism and its Institutions*. Oxford: Basil Blackwell.

Dawson, M. R. W. and Medler, D. A. (eds) (1997). *Cognitive Science Dictionary*. Edmonton: University of Alberta. Available: http://web.psych.ualberta.ca/~mike/Pearl_Street/Dictionary/dictionary.html

Dennett, D. C. (1984; rpt. 1990). Cognitive wheels: the frame problem of AI. In Boden, M. A. (ed.), *The Philosophy of Artificial Intelligence*. Oxford Readings in Philosophy. Oxford: Oxford University Press, pp. 147–70.

Fischer, M. (1966). The KWIC index concept: a retrospective view. *American Documentation*, (April), pp. 57–70. Reprinted in Elias, A. W. (ed.), *Key Papers in Information Science*. Washington, DC: American Society for Information Science, 1971, pp. 121–34.

Fodor, J. (1998). The trouble with psychological Darwinism. Review of Steven Pinker, *How the Mind Works*, and Henry Plotkin, *Evolution in Mind*. *London Review of Books*, **20**(2). Available: http://www.lrb.co.uk/v20n02/fodo2002.html

Fowler, D. (1999). Criticism as commentary and commentary as criticism in the age of electronic media. In Most, G. W. (ed.), *Commentaries—Kommentare. Aporemata: Kritische Studien zur Philologiegeschichte, Band 4*. Göttingen: Vandenhoeck & Ruprecht, pp. 426–42.

Fox, E. A. and Marchionini, G. (1998). Toward a worldwide digital library. *Communications of the ACM*, **41**(4): 29–32.

Friedman, J. (1999). Creativity in science. In *The Humanities and the Sciences*. ACLS Occasional Paper, No. 47. New York: American Council of Learned Societies, pp. 9–14. See under Publications at http://www.acls.org/

Frye, N. (1991). Literary and mechanical models. In Lancashire, I. (ed.), *Research in Humanities Computing I*. Oxford: Clarendon Press, pp. 3–12.

Galison, P. (1997). *Image and Logic: A Material Culture of Microphysics.* Chicago, IL: University of Chicago Press.

Galison, P. (1999). Objectivity is Romantic. In *The Humanities and the Sciences*. ACLS Occasional Paper, No. 47. New York: American Council of Learned Societies, pp. 15–43. See under Publications at http://www.acls.org/

Geertz, C. (1973). *The Interpretation of Cultures*. Reprinted 1992. London: Fontana Press.

Genre and Multimodality (2000). University of Sterling. Available: http://www.gem.stir.ac.uk/

Gooding, D., Pinch, T., and Schaffer, S. (eds) (1989). *The Uses of Experiment*. Cambridge: Cambridge University Press.

Groenewold, H. J. (1961). The model in physics. In Freudenthal, H. (ed.), *The Concept and the Role of the Model in Mathematics and Natural and Social Sciences*. Dordrecht: D. Reidel.

Güzeldere, G. and Franchi, S. (eds) (1995). *Constructions of the Mind: Artificial Intelligence and the Humanities. Stanford Humanities Review*, **4**(2). Available: http://www.stanford.edu/group/SHR/4-2/text/toc.html

Hacking, I. (1983). *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science.* Cambridge: Cambridge University Press.

Hacking, I. (1999). *The Social Construction of What?* Cambridge, MA: Harvard University Press.

Hockey, S. (1988). *Micro-OCP User Manual*. Reprinted Oxford: Oxford University Press, 1993.

Hockey, S. and Martin, J. (1987). The Oxford Concordance Program Version 2. *Literary and Linguistic Computing*, **2**: 125–31.

Hofstadter, D. R. (1995). On seeing A's and seeing As. In Güzeldere, G. and Franchi, S. (eds), *Constructions of the Mind: Artificial Intelligence and the*

Humanities. Stanford Humanities Review, **4**(2): 109–21. Available: http://www.stanford.edu/group/SHR/4-2/text/hofstadter.html

Horton, T., *et al.* (1998). *The Elta Software Initiative*. Available: http://www.kcl.ac.uk/humanities/cch/elta/

Humanist Discussion Group (1987–). http://www.kcl.ac.uk/humanities/cch/humanist/

The Humanities and the Sciences (1999). ACLS Occasional Paper, No. 47. New York: American Council of Learned Societies. See under Publications at http://www.acls.org/

Ide, N. and Veronis, J. (1993). What next, after the Text Encoding Initiative? The need for text software. *ACH Newsletter*, **15**(1). Available: http://www.ach.org/ACH_Newsletter/15.1_Newsletter.html#afterTEI

International Symposium on Music Information Retrieval (2000). http://ciir.cs.umass.edu/music2000/

Kernighan, B. W. and Plauger, P. J. (1976). *Software Tools*. Reading, MA: Addison Wesley.

Lancashire, I. (1991). *Humanities Computing Yearbook 1989–90*. Oxford: Oxford University Press.

Lancashire, I., Bradley, J., Wooldridge, R., McCarty, W., and Stairs, M. (1996). *Using TACT with Electronic Texts*. New York: Modern Language Association.

Latour, B. and Woolgar, S. (1986). *Laboratory Life: The Construction of Scientific Facts*. Princeton, NJ: Princeton University Press.

Leatherdale, W. H. (1974). *The Role of Analogy, Model, and Metaphor in Science*. New York: American Elsevier.

Luhn, H. P. (1966). Keyword-in-Context Index for technical literature (KWIC Index). In Hays, D. G. (ed.), *Readings in Automatic Language Processing*. New York: American Elsevier.

Manovich, L. (2001). *The Language of New Media*. Cambridge, MA: MIT Press.

McCarty, W. (1993). Handmade, computer-assisted, and electronic concordances of Chaucer. In Lancashire, I. (ed.), *Computer-based Chaucer Studies*. *CCH Working Papers*, **3**: 49–65.

McCarty, W. (1994). Encoding persons and places in *The Metamorphoses* of Ovid. Part 2: the metatextual translation. *Texte: Revue de Critique et de Théorie Littéraire*, **15/16**: 261–306. Available: http://www.chass.utoronto.ca/french/litera/Revue_Texte/

McCarty, W. (1999a). Humanities computing as interdiscipline. *Is Humanities Computing an Academic Discipline?* IATH, University of Virginia, 5 November 1999. Available: http://ilex.cc.kcl.ac.uk/wlm/essays/inter/

McCarty, W. (1999b). A note on terminology: *humanities computing* or *humanistic informatics*? Available: http://ilex.cc.kcl.ac.uk/wlm/essays/inter/terminology.html

McCarty, W. (1999c). We would know how we know what we know: responding to the computational transformation of the humanities. *The Transformation of Science: Research between Printed Information and the Challenges of Electronic Networks*. Max Planck Gesellschaft, Schloss Elmau, 31 May–2 June 1999. Available: http://ilex.cc.kcl.ac.uk/wlm/essays/know/

McCarty, W. (2000). A serious beginners guide to hypertext research. Available: http://ilex.cc.kcl.ac.uk/wlm/essays/achallc2000/hyperbib.html

McCarty, W. (2002). A network with a thousand entrances: commentary in an electronic age? In Gibson, R. K. and Kraus, C. S. (eds), *The Classical Commentary*. Leiden: Brill (forthcoming).

McCarty, W. and Kirschenbaum, M. (2000). Humanities computing units and institutional resources. Available: http://www.kcl.ac.uk/humanities/cch/wlm/hcu

McGann, J. (1997). Imagining what you don't know: the theoretical goals of the Rosetti Archive. Available: http://jefferson.village.virginia.edu/~jjm2f/chum.html

Menzel, P. and D'Aluisio, F. (2000). *Robo Sapiens: Evolution of a New Species.* Cambridge, MA: MIT Press.

Minsky, M. (1967). *Computation: Finite and Infinite Machines*. Englewood Cliffs, NJ: Prentice–Hall.

Newell, A. and Simon, H. (1976). Computer science as empirical enquiry: symbols and search. In Boden, M. A. (ed.) *The Philosophy of Artificial Intelligence*. Oxford Readings in Philosophy. Oxford: Oxford University Press, pp. 105–32.

Newell, W. H. (ed.) (1998). *Interdisciplinarity: Essays from the Literature*. New York: College Entrance Examination Board.

NINCH (National Initiative for a Networked Cultural Heritage) (2000). Building Blocks Project. Available: http://www-ninch.cni.org/bb/project/project.html

Online Music Recognition and Searching (2000). Available: http://www.kcl.ac.uk/kis/schools/hums/music/ttc/IR_projects/OMRAS/

Orlandi, T. (1990). L'ambiente Unix e le applicazioni umanistiche. *Archeologia e Calcolatori*, **1**: 237–51.

Ott, W. (2000a). Strategies and tools for textual scholarship: the Tübingen System of Text Processing Programs (TUSTEP). *Literary and Linguistic Computing*, **15**(1): 93–108.

Ott, W. (2000b). *TUSTEP*. Available: http://www.uni-tuebingen.de/zdv/zrlinfo/tustep_eng.html

Ott, W. (2000c). *Humanist*, **14**(0262) [2], 22 September. Available: http://lists.village.virginia.edu/lists_archive/Humanist/v14/

Paepcke, A., Chang, C.-C. K., García-Molina, H., and Winograd, T. (1998). Interoperability for digital libraries worldwide. *Communications of the ACM*, **1**(4): 33-43. ACM Digital Library.

Perseus Digital Library. http://www.perseus.tufts.edu/

Petre, M., Blackwell, A., and Green, T. (1998). Cognitive questions in software visualisation. In Stasko, J. T., Domingue, J. B., Brown, M. H. and Price, B. A. (eds), *Software Visualisation*. Cambridge, MA: MIT Press, pp. 453–80.

Pickering, A. (1995). *The Mangle of Practice: Time, Agency, and Science*. Chicago, IL: University of Chicago Press.

Pinker, S. (1997). *How the Mind Works*. Harmondsworth: Allen Lane, The Penguin Press.

Pinker, S. (1999). *Words and Rules: The Ingredients of Language*. London: Weidenfield and Nicolson.

Preston, M. and Coleman, S. S. (1978). Some considerations concerning encoding and concording texts. *Computers and the Humanities*, **12**: 3–12.

Redhead, M. (1980). Models in physics. *British Journal of the Philosophical Society*, **31**: 146–7.

Reizmann, R. E., Quinn, R. D., Watson, J. T., and Zill, S. N. (2000). Insect walking and biorobotics: a relationship with mutual benefits. *Bioscience*, **50**: 23–33.

Roland, P.-Y., Cambouropoulos, E., and Wiggins, G. A. (eds) (2001). *Pattern Processing in Music Analysis and Creation*. Special issue of *Computers and the Humanities*, **35**.

Rorty, R. (2000). Being that can be understood is language. *London Review of Books*, **22**(6): 23–5.

Sapir, E. (1921). *Language: An Introduction to the Study of Speech.* New York: Harcourt Brace.

Searle, J. (1984). *Minds, Brains and Science: The 1984 Reith Lectures.* Harmondsworth: Penguin.

Simon, H. and Newell, A. (1958). Heuristic problem solving: the next advance in operations research. *Operations Research*, **6** (Jan–Feb): 1–10.

Simons, G. (1998). In search of task-centered software: building single-purpose tools from multipurpose components. SIL Electronic Working Papers 1998-004. Available: http://www.sil.org/silewp/1998/004/SILEWP1998-004.html

Smith, D. A., Rydberg-Cox, J. A., and Crane, G. R. (2000). The Perseus Project: a digital library for the humanities. *Literary and Linguistic Computing*, **15**(1): 15–25.

Smith, J. B. (1997). The king is dead; long live the king. Keynote speech for Hypertext '97, Southampton, UK, 6–11 April. *Proceedings of the Eighth ACM Conference on Hypertext*, p. 240.

Sperberg-McQueen, C. M. (1996). Text Analysis Software Planning Meeting, Princeton, 17–19 May: Trip Report. Available: http://www-tei.uic.edu/orgs/tei/misc/trips/ceth9605.html

Text Encoding Initiative (2000). Available: http://www.tei-c.org/

Tufte, E. R. (1983). *The Visual Display of Quantitative Information*. Cheshire CN Graphics Press.

Tufte, E. R. (1990). *Envisioning Information*. Cheshire CN Graphics Press.

Tufte, E. R. (1997). *Visual Explanations: Images and Quantities, Evidence and Narrative*. Cheshire CN Graphics Press.

Unsworth, J. (comp.) (2000a). Is humanities computing an academic discipline? An interdisciplinary seminar, Institute for Advanced Technology in the Humanities, University of Virginia. Available: http://www.iath.virginia.edu/hcs/

Unsworth, J. (2000b). Scholarly primitives: what methods do humanities researchers have in common, and how might our tools reflect this? Humanities computing: formal methods, experimental practice. In the Seminar on Humanities Computing, King's College London. Available: http://jefferson.village.virginia.edu/~jmu2m/Kings.5-00/primitives.html

Unsworth, J. (2001). Knowledge representation in humanities computing. eHumanities Lecture, National Endowment for the Humanities, 3 April. Available: http://www.iath.virginia.edu/~jmu2m/KR/KRinHC.html

Wixon, D. and Ramey, J. (eds) (1996). *Field Methods Casebook for Software Design*. New York: John Wiley.

Wooldridge, T. R. (ed.) (1991). *A TACT Exemplar*. CCH Working Papers 1. Toronto, Ont.: Centre for Computing in the Humanities. Reprinted as individual papers in *Computing in the Humanities Working Papers* (s.v. *CCHWP 1*). Available: http://www.chass.utoronto.ca/epc/chwp/titles.html#Postprints

Wright, P. (2000). The psychology of layout: consequences of the visual structure of documents. In Power, R. and Scott, D. (eds), *Using Layout for the Generation Understanding or Retrieval of Documents*. American Association for Artificial Intelligence Technical Report FS-99-04, pp. 1–9.