Pliny: A model for digital support of scholarship

John Bradley
Centre for Computing in the Humanities
King's College London
john.bradley@kcl.ac.uk

It could be argued that the digital humanities (DH) began with the work of Father Roberto Busa in the late 1940s and early 1950s when he conceived of his monumental *Index Thomisticus*. Early in his plans for the *Index* he recognised that to complete it would require the assistance of "some sort of machinery" (Busa 1980, p. 87), and shortly after a famous meeting with Thomas Watson of IBM (also described in Busa 1980) he came to believe that the machinery he needed was the computer. Perhaps because the computer offered the potential for radically new ways to work with texts compared to the approaches used in the past, the DH has become an evangelical field whose adherents believe that they should promote the role of computing in humanities scholarship to others. For more than 30 years DH practitioners have been anxious to convince the rest of the humanities that using the computer in ways that are beyond word processing (or, more recently, web browsing) should be an important component of their research methodology. From the very beginning they have claimed that it offers significant benefits -- see the early writings of Potter (1988) and Smith (1978) for a sense of some of the aspirations of these early digital humanists.

The development of software tools to support humanities scholarship has been one of the major strands of DH from its very outset. There has been, however, a general sense of disappointment within the DH community that their work has not yet had the widespread impact that had been expected by the early pioneers. Indeed, the recent Summit on Digital Tools in the Humanities (University of Virginia, September 2005) recognised that "only about six percent of humanist scholars go beyond general purpose information technology and use digital resources and more complex digital tools in their scholarship". (pg 4) and that "although humanists are on the verge of ... a revolutionary change in the scholarship... that such a revolutionary change has not yet occurred". (pg 5). There has been a widespread belief in the DH community that the mere passage of time, and the turnover of the generations, will eventually create a context in which digital approaches to scholarly research will be more accepted within the humanities mainstream. Unfortunately, my own experience at King's dealing with students has aligned me more closely to Richard Cunningham's recent post to the Humanist discussion list (Cunningham 2007) where he noted that his English students responded to his discussion about digital techniques with not only an unimpressed reaction, but perhaps even a shocked or angry one. As Cunningham then says:

"The anger came, I think, from having spent four or more years devoting themselves to a discipline they had now been disciplined into believing in, rather than questioning with the critical reasoning English scholars so often so loudly proclaim is their *raison d'etre*."

Indeed, later in the same posting he notes that "my students but also in others in my profession [...] seem to regard the computer and the digital revolution as a fancy that will pass."

I'm sure that a part of the resistance to the new techniques of the DH does indeed arise from issues such as the one raised by Cunningham. A negative response to the digital humanities may seem natural from a collection of students who, by their very choice of field to study, are perhaps suspicious of technology in the first place. However, it seems to me there are other problems inherent within the digital humanities, at least as it is thought of by DH practitioners, that limit uptake by the humanities research community at large. My own sense (described in Bradley 2005) of why the DH has had little impact so far is that the paradigm for computing that is held by the majority of non-digital humanists and that of the digital humanist community are significantly different and perhaps even incompatible. In Bradley 2005 I suggested that tool builders in the digital humanities would have better success persuading their non-digital colleagues that computers could have a significant positive benefit on their research if the tools they built fit better into how humanities scholarship is generally done, rather than if they developed new tools that were premised upon a radically different way to do things. As a colleague has remarked to me recently in a private correspondence:

Too much of IT rhetoric turns on dramatic as if the computer were the Apollo of Belvedere, who in the last line of Rilke's poem says "You must change your life" (Du musst dein Leben aendern). But people don't want to be told by anybody, even the Apollo of Belvedere and certainly not by computers, that they should change their life. (permission to quote here granted by the correspondent)

If a researcher came to a member of the DH community and asked them how s/he could use a computer to enhance his/her research work, what would the DH member say? Â The researcher probably already uses a word processor, and whether s/he recognises it or not, for this reason alone is doing some aspects of his/her research in ways that are different from those available to past generations. Furthermore, if the researcher is if not fully aware of the high quality scholarly resources available to his/her desktop through the browser, then of course s/he should be made aware of them. However, my contention is that the delivery of these resources via the web browser; which then limits what the user can do with them to reading them on the screen, printing them on paper, or perhaps saving the HTML page; underplays the potential of the digital materials behind these web pages. The browser user experiences the web primarily in terms of authors and users -- with s/he in the "user" camp. This is, of course, not entirely wrong. As schraefel *et al* (2004) note in their discussion of read-write Hypertext, scholars are indeed users of other people's materials. However, they also observe rightly that scholars are creators as well, and any system that most scholars experience as read-only will seem to them to support only a part of their scholarly life. There is, of course, a major strand within the digital humanities that focuses on creation -- the creation of digital scholarly editions. Furthermore, it has proven to be an extremely rich and fruitful way to think about scholarly editing, and the DH community has powerful models and tools to support this kind of work -- not least the TEI (TEI 2006). However, the preparation of scholarly textual editions represents an activity that is only a part (for many scholars a small part) of most scholar's work.

Pliny is a piece of software that explores some of the potential for computing support for activities involved in scholarly work that are not currently well addressed by the DH

community. It grows out of work I have been doing over a number of years (see Bradley 2003 for a relatively recent report on this, and Bradley and Vetch 2007 for some more recent thinking specifically about annotation). It aims to bring the benefits of computing to *existing scholarly practice*, rather than requiring the user to adopt new research strategies. Its design has been influenced by Vannevar Bush's visionary, but imaginary, machine, the *Memex*, which also aimed at supporting individual research centred on the reading of primary and secondary sources: "[a] device in which an individual stores all his books, records, and communications, and which is mechanized so that it may be consulted with exceeding speed and flexibility." (from *As We May Think*, Atlantic Monthly 1945). It also pays it dues to a long string of thinking about supporting research that finds it way through Bush, Ted Nelson, Douglas Englebart and others (see an interesting review of this in Trigg 1996). Although *Pliny* is actually publicly available now (see pliny.cch.kcl.ac.uk), it is not really a publication-quality piece of software. Instead, it should be viewed as a thought-piece that recognises issues that arise if we try to use the computer to provide support for a task that scholars often undertake -- that of developing an interpretation of materials that interest them.

So, what is humanities research, really?

Finding out what scholars actually do in any detail when they do their research has proven to be difficult, and the DH would benefit from an extensive and perhaps intimate survey of the actual research practices of those who are scholars but not currently in the DH community. From time to time a book has been be written that was clearly aimed at helping beginning researchers learn how to carry out research -- see Altick 1963 as an example -- and these can give us insight into what might be going on. A recent book like Griffin's *Research Methods for English Studies* (2005) acknowledges, however, that research methodology is not much discussed or described in any detail. She comments that even today "significant numbers of English studies academics in the UK" are "surprisingly in- or possibly non-articulate about what they do to achieve ... results" (p. 1). Carole L. Palmer -- a researcher in the Graduate School of Library and Information Science who has, as one of her research interests, an interest in studying humanities research -- remarks in her very useful overview of IS research directed at scholarly practice that within information science there has been much more focus on scientific research than humanities research, commenting that there "are considerably fewer works" that focused on this community. (Palmer and Cragin 2007, p. 4).

William Brockman *et al*'s report to the Council on Library and Information Resources (2001) reports on interviews with humanities researchers to get some insight in how they worked, and their report touched on what seems to me to be several important issues (described later in this article) that have influenced my work on *Pliny*. A further useful source for ideas comes from work that has been done in the Social Sciences -- see a brief introduction to this in Bradley 2003, including references to well established pieces of software that have been developed to support Social Science's qualitative research methodology whose design has influenced Pliny.

All these sources suggest that a central component of scholarly research is centered on extensive reading. Yet, in Palmer's more recent report (Palmer 2007) she observes that "few studies have examined the actual scholarly nature of the activity of reading, that is, how researchers interact directly with content and apply it to the research process." (p. 17) Even so, the limited amount of research available has relevant things to say to us here. Brockman *et al* claim "The process of reading and searching, developing context, and rereading and researching are at the heart of humanities scholarship", and then go on to point out that in the

humanities scholarly reading covers a broad range of both primary and secondary sources and includes intensive reading of key sources. Bates (1996) (quoted in Brockman *et al* p. 7) claims that "[t]he scholar who claims to be current and knowledgeable in a field must have read closely and be intimately familiar with a large number of particular works". What happens during reading has been described by few authors. John Lavagnino (Lavagnino 1997) observes that "reading as an involving process, not as interpretation or decoding".

Some researchers interviewed by Brockman *et al* mention annotation and notetaking as key activities that go on during reading for research. Brockman *et al* (p. 25) quotes a researcher that they interviewed who reported that s/he always wrote in the margins of materials s/he needed to absorb and that s/he "almost never underlined without writing in the margin" because otherwise s/he found him/herself "simply underlining, rather than absorbing".

Annotation/notetaking by itself is not the end of the process of research, of course. We see a hint that by itself it is only the beginning of a larger research process in Lavagnino (1997) where he continues his remark noted above to say that reading "can lead to interpretation, but only by way of generating reactions that we *subsequently* seek to describe or explain" (my highlight). Palmer (2007), in her section on "Writing", quotes writings by H.J. Rheinberger where he remarks that in the sciences the "scribbling" and "jotting" of ideas are "nearer to the materialities of scientific work than are research communications", and therefore (to now quote Palmer herself) "play[s] a key role in mediating information". In the fourteenth chapter of his 1963 book Altick makes this transformation from notetaking to interpretation building more explicit by proposing a rather formal non-digital technique. He describes how to record notes on note cards while reading. Then, as ideas begin to emerge from the reading they are recorded on separate "concept" note cards and linked to the reading note cards by recording lists of references to relevant note cards on the concept cards. The clear naming of the concept on the card is a part of the process of having the ideas emerge -- Altick insists on each concept card having "a readily intelligible caption".

At some point an interpretation emerges and matures. The scholar is then ready to report his/her finds to others. In traditional scholarship this is the article or monograph. The task, then, is to take the elements of the interpretation and order and present them in a fashion so that they are intelligible to a reader. Here, we find Brockman *et al* claiming that the researchers they interviewed reported that they use Word as a note organiser, and it appeared that they did not find it very satisfactory.

From this we can see three possible phases to traditional scholarly research:

- first, reading and annotation and/or notetaking;
- then, organising notes to find relationships. This is usually a personal, internal activity. The researcher expects to develop organising principles for these notes from intensive and extending thinking about the material and his/her reactions to it. One of the steps in developing an interpretation is the naming of a concept that one has discovered in the texts, and to relate it to other concepts. Note that although searching through materials is a part of the task of finding relevant materials, it is only a part of the work -- what one finds must then be organised, collated, identified and related together into some sort of unified structure.
- finally publishing (in an article), in which one turns what one has found as a personal interpretation into a public object.

(in reality, these "phases" are intermingled, of course -- reading proceeds in parallel with thinking about how what one observes while reading fits together with pre-existing ideas).

One can find pieces of software to assist with notetaking and note storage, and several of them are very good indeed. See, for example *Zotero* (Zotero 2007), which is integrated into the FireFox browser and (quoting the website) allows one to "collect, manage, and cite your research sources"; or *Tinderbox* (Tindexbox 2007) which (to quote its website) "lets you record ideas quickly and keep them where you'll find them again when you need them". As striking as these tools are, I like to think that Pliny is different from them in that it is trying to explore how an integrated tool could be developed to assist the scholar with all *three* aspects of their research rather than focusing primarily on the activity of note taking and subsequent note searching. Let us now move to look at how these three phases are represented in Pliny itself. In the folllowing sections we will take each one in turn.

Reading and Annotation in Pliny

Pliny provides personal annotation mechanisms that can be added to established public and private digital resources (currently web pages, images, text as PDF), and it is designed in such a way that annotation tools for other digital media (such as, say, XML/TEI documents, or video or audio) could be added as further extensions. It can even be used to record notes that arise from the reading of non-digital (print) material. This recognises the widely held view that scholarly research involves reading (and presumably notetaking) across a broad range of sources, and potentially across both digital and non-digital objects.

Some part of the DH community has been driven by the technologies of the WWW for some years now. Almost all funded research at my institution that involves digital techniques is aimed in the end at producing a digital resource that is accessible over the WWW and has the user using a web browser to access it. In spite of this, *Pliny* is not a web application and a user does not open his/her browser to use it -- although as we shall show shortly *Pliny* does integrate a web browser into its operation and supports the managing of notes taken while reading web pages. Notes are not stored by *Pliny* on some kind of remote annotation server, they are stored on the user's personal computer -- as befits material that should belong to the researcher who created them.

You can see *Pliny* being used to store notes associated with a web page (from the *Proceedings of the Old Bailey* site: http://www.oldbaileyonline.org) in figure 1. *Pliny* uses its host machine's standard web browser to display the web page (here, it is Internet Explorer, but on other platforms it could be, say, Firefox) in the larger left hand area, and provides an area (called a *reference area* in Pliny terminology) to the right where the user can record notes. Here, the user has recorded two notes to describe two observations. The dividing line between the reference area for the web page and the web page is placed here quite far to the right -- but the user can move this boundary to increase the size of reference area if s/he wishes.
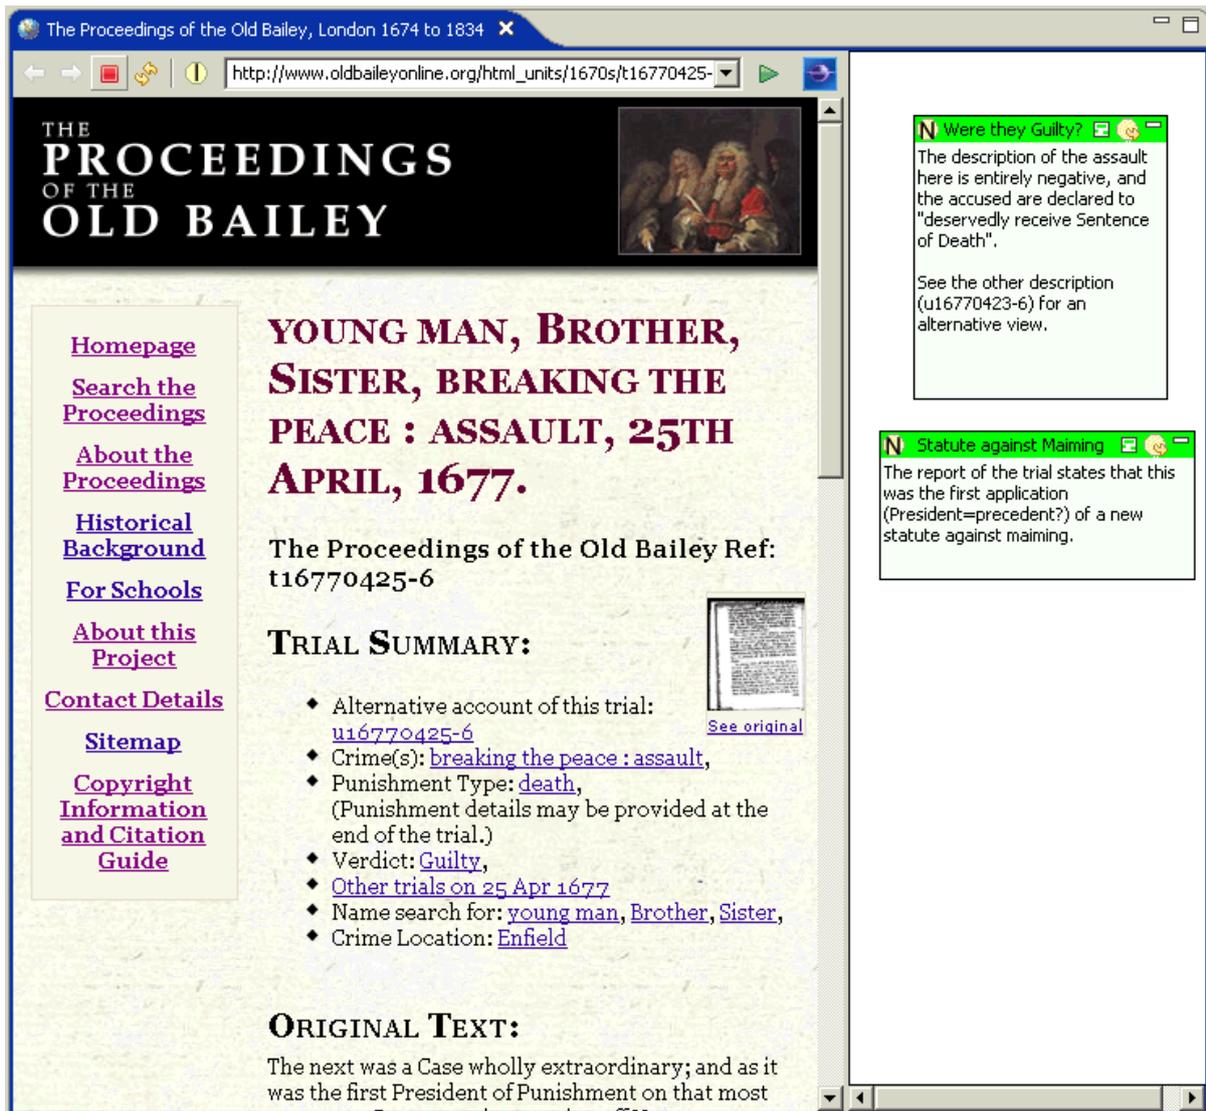
Figure 1

The notes are stored by *Pliny* on the user's own machine, and the website does not have to *do* anything to support *Pliny* -- pretty well any web page (at least any that doesn't use frames) can be thus annotated. *Pliny* will redisplay the current set of notes attached to this web page any time in the future if it is revisited using the same URL.

The *Pliny* prototype allows a user to go beyond note recording (as shown in figure 1) to annotate fully both images and PDF documents. See figure 2 for an example of an image annotated in *Pliny*. The image is the frontispiece of Giambattista Vico's *Scienza Nuova*, and Vico points out in his introduction that the image can be interpreted as a allegory of the topics he covers in his text. In our figure the user has used *Pliny* to label parts of the image that refer to several of those topics, and has also included a couple of comments about the significance of this image for his/her own study.
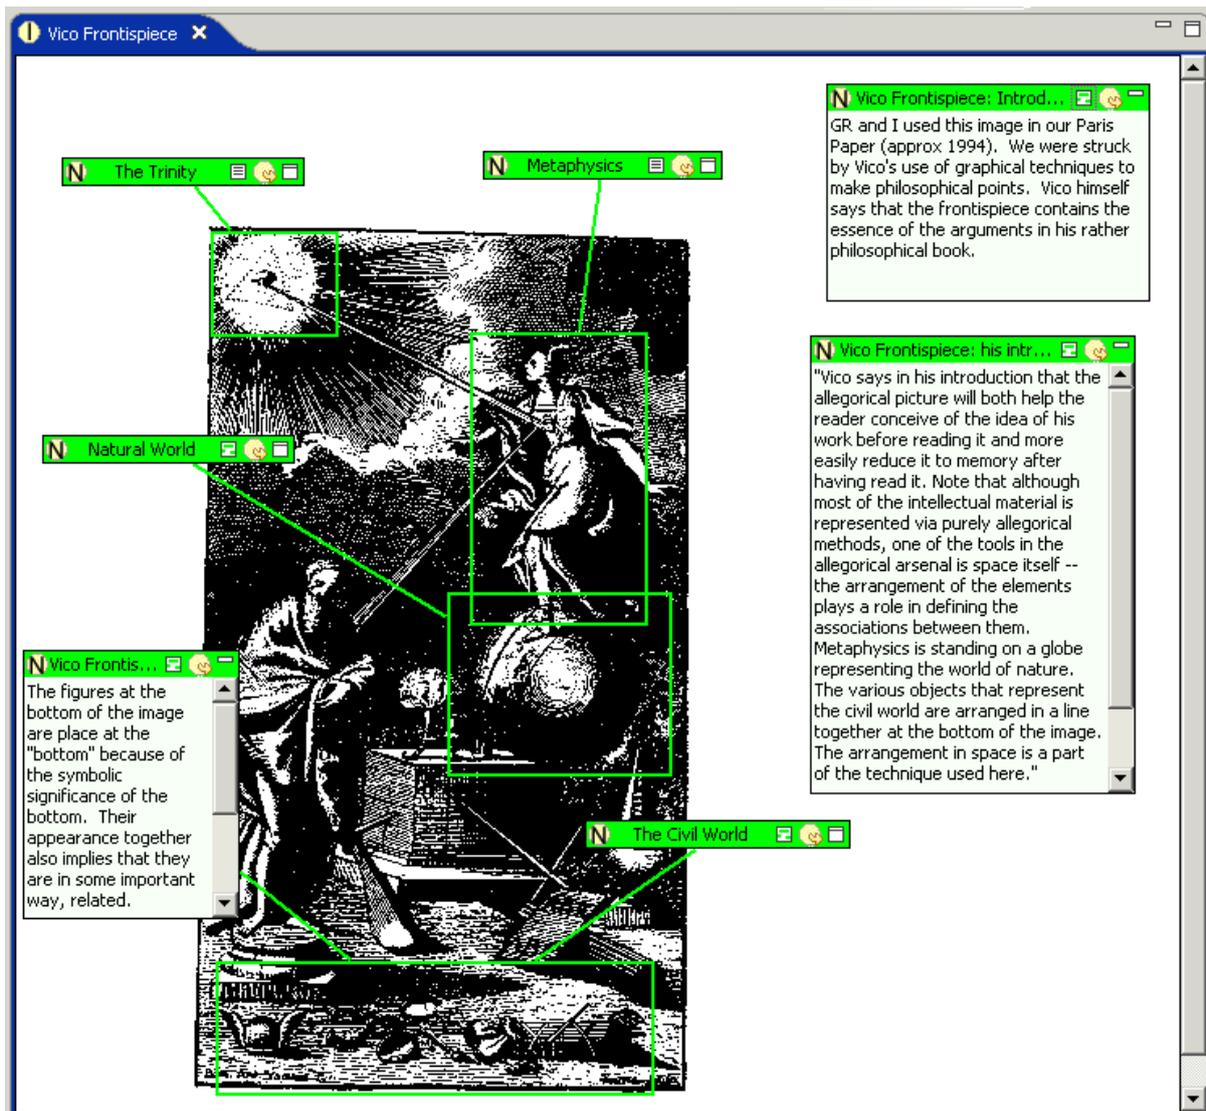
The Trinity

Metaphysics

Natural World

Vico Frontispiece: Introd...
GR and I used this image in our Paris Paper (approx 1994). We were struck by Vico's use of graphical techniques to make philosophical points. Vico himself says that the frontispiece contains the essence of the arguments in his rather philosophical book.

Vico Frontispiece: his intr...
"Vico says in his introduction that the allegorical picture will both help the reader conceive of the idea of his work before reading it and more easily reduce it to memory after having read it. Note that although most of the intellectual material is represented via purely allegorical methods, one of the tools in the allegorical arsenal is space itself -- the arrangement of the elements plays a role in defining the associations between them. Metaphysics is standing on a globe representing the world of nature. The various objects that represent the civil world are arranged in a line together at the bottom of the image. The arrangement in space is a part of the technique used here."

Vico Frontis...
The figures at the bottom of the image are place at the "bottom" because of the symbolic significance of the bottom. Their appearance together also implies that they are in some important way, related.

The Civil World

Figure 2

Finally, Figure 3 shows *Pliny* being used to record notes for a non-digital object -- in this case Susan Wittig's article *The Computer and the Concept of Text* (Wittig 1978). Here the Pliny screen is again divided into two areas. The smaller left area contains some citation information about the Wittig article. The right area (again called a *reference area* in *Pliny*) is the focus of interest for the researcher who has read Wittig and contains five notes (with the green headers) that the reader created while reading the article. Subsequent to reading the article the reader noted that his/her notes could be organised into two major thematic groups, and has done so by creating two containers to represent these themes (one entitled *Limits of [the] New Crit[ical] base*, the other *reading and semiotics*). They have their headers coloured here in light yellow (the significance of colour in the header is given later in this article) and the user has put the relevant notes in each of these two containers.
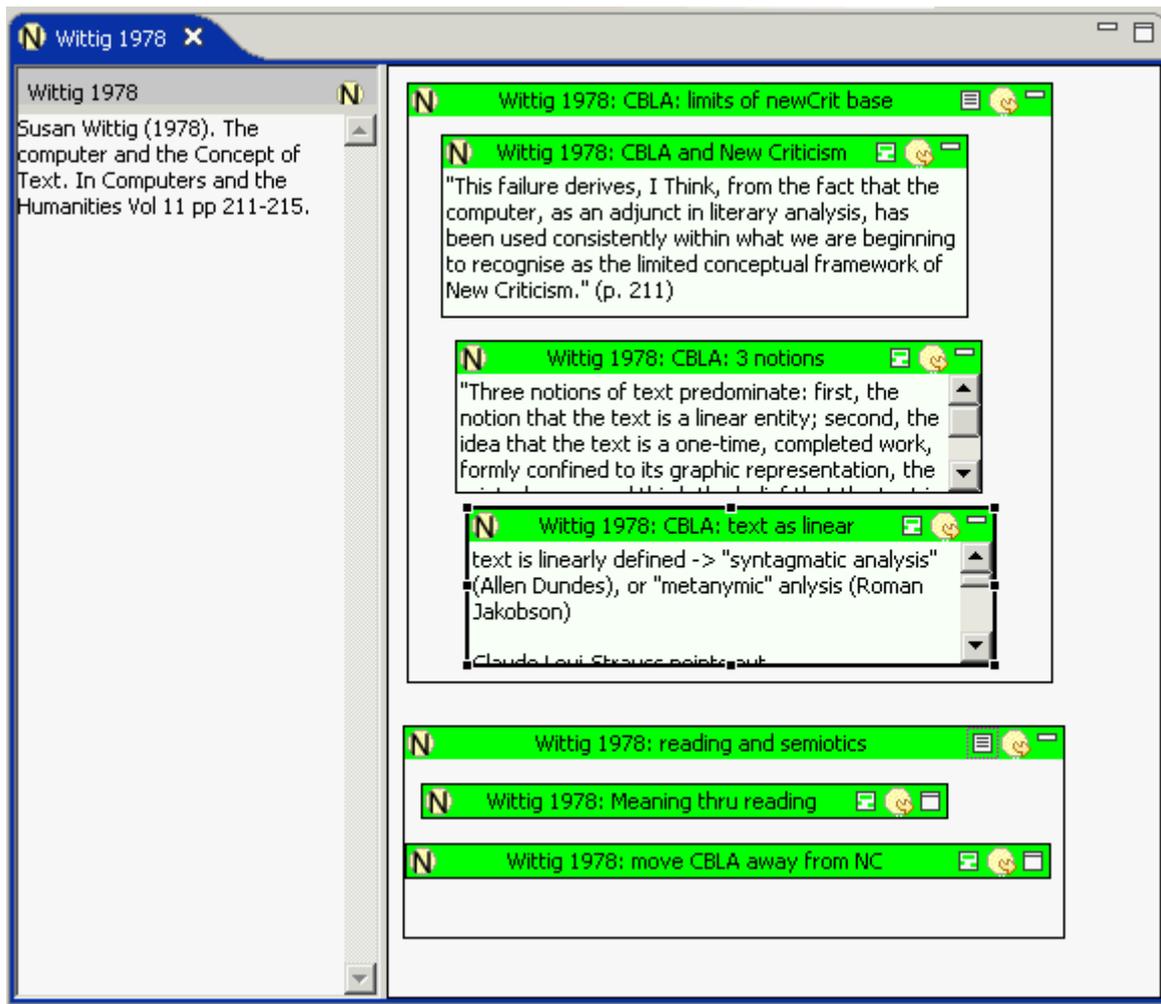
Figure 3

Digital objects that the user is working with are called by the general name of *resources* within Pliny. A web page, an image, a PDF file could be a Pliny resource, and Pliny code is designed to be extendable (indeed, in building *Pliny* I have taken extendibility very seriously indeed and it is, for that reason, built on top of Eclipse technology (Eclipse 2007)) so that code could be added to support the needs of other kinds of resources such as XML documents, bibliographic data, or audio/video files. In addition to being able to display the primary "content" of these resources (the web page content, the image, the pages of the PDF file), *Pliny* provides a place to record references to other resources that the user wishes to associate with the resource -- the reference or annotation area.

In addition to supporting these non-Pliny specific digital objects as resources, *Pliny* has a resource type of its own which we've already seen called the *note*. A note exists only within the Pliny system (although *Pliny*'s exporting and importing mechanism make its contents available to other computer applications), and its primary content is (only) plain text -- in appearance like an ASCII file. In spite of this apparently serious limitation, it turns out that the Pliny note is very significant within *Pliny*. First of all, generally each scholarly note or annotation is often a rather short bit of text and can actually be well represented by simply plain text, and as we will see shortly most annotation activities involve creating Pliny notes and associating them with other Pliny resources to which they apply. Furthermore, as we shall also see later, like all Pliny resources, a Pliny note provides a reference area where other

notes or other Pliny resources can be referenced. A more complex idea, which in an article might result in a rather long block of text, might here be presented by a set of shorter textual notes that are associated together in a reference area. The idea that all resources have an associated 2D space where references to other resources can be included turns into an important concept in *Pliny*. It allows Pliny notes to be used to store not only notes and annotations linked to the source that stimulated them in the first place, but also structures that might be thought of as representing interpretative concepts -- there is more discussion of this below.

Getting people to annotate digitally

It is instructive to see that a significant body of work in computer science on annotation (see, for example Golovchinsky *et al* 1999 and Price 1997) arises at about the same time as the tablet computer; a machine with the ability to capture simple writing with a stylus on the computer screen. Why is this? It would seem that an important aspect of recording annotation and notetaking is involved in the gestures involved in doing it. In Bradley and Vetch 2007 we noted that if digital annotation was to be effective, it was critical that the action involved in annotating could be found to be a "natural" one which didn't interfere with the kind of thinking process that gives rise to material to annotate in the first place.

Discussions in articles which distinguished annotation on paper from digital annotations have from time to time drawn attention to this as well. Even before the tablet computer was well known, researchers recognised that digital annotation had to be easy to use, or perhaps feel natural, if it was to be effective. O'Hara and Sellen (1997) in their comparison of the affordances offered by online and paper documents noted that "... had the document been on paper, their [subjects'] natural tendency would have been to annotate the document in some way or other. However only one of the subjects attempted to do this on-line ... In doing so, this subject experienced a number of difficulties which interfered with the smooth flow of reading:

They then go on to quote one of their subjects who took part in their experiments:

"So the annotation was not as easy as all that... I think the whole process would have been a lot quicker on paper. Annotations are that much more flexible because you can write in the margins which you can't very easily do here. You have to establish a new text block and then have to write."

Pliny, perhaps unfortunately, does not support the digital ink interface of tablet computers. However, I have tried to recognise the significance of each user interaction *gesture* in designing the digital interface for *Pliny*. For example, *Pliny* supports drag and drop in many ways. Not only do users create notes that comment on something by dragging out a note area in a reference area, but they can actually drag some text (from, say, a web page of interest) into the associated reference area to readily create a note that contains this text. They can drag the URL from an external (to Pliny) web browser into *Pliny* to create a Pliny resource object for that URL and a reference to it on the current reference space. The gesture-like nature of drag-and-drop seems to be a good fit with the gesture-like nature of annotation in a book, and the need for the act of creating the annotation to not interfere too much with the reading that triggered it.

We believe that it is appropriate to focus so much on drag and drop as a suitable approach. There is some evidence in published HCI research that suggests that drag and drop often reduces the mental burden put upon a user when they are carrying out the action. See, for example, Rees 2001, where the author, while assessing the extent that the web browser interface might become the dominant interface on personal computers, points out both that drag and drop is not part of the repetroire of user actions that browsers support naturally, and that that drag and drop is "one of the most urgent requirements" if interaction with web documents is to be enhanced for users (p. 2). Lim (1996) proposes a concept of "automaticity" (p. 6) to explain how drag and drop reduces the amount of mental effort required to carry out certain actions, and demonstrated by experiment that users found the use of DnD significantly better for certain operations than, say, selection by Menu systems. Pliny's 2D central paradigm uses DnD by its very nature since the user is encouraged to organise notes by dragging them about in a 2D space -- when a user drops a bit of text into Pliny's 2D reference area it is obvious where the note created to store the text should be put -- at the point the user dropped the text! Using DnD to support the introduction of new materials from outside of Pliny seems quite natural.

Supporting interpretation development

We have now described how Pliny supports annotation and notetaking, and thereby supports the first aspect of humanities research. Let us now turn to the second phase -- the development of an interpretation of what one has read. It might seem that technologies such as computing ontologies (for an overview of OWL; one of the best known Ontology languages, see McGuiness and van Harmelen (2004)) might be highly relevant here, but it does not seem that their approach is as appropriate as it might seem at first glance. One of the difficulties is related to something that I observed in (Bradley 2003) -- that the development of an interpretation is a process which progresses from a fuzzy to a clearer sense of the issues and relationships. Instead, Ontologies are highly structured paradigms which require highly structured material that one would expect to emerge, if at all, near the end of the interpretation development process taken by an individual scholar. Unfortunately, the highly formal nature of computing ontologies provides little help while the model is still being developed: when the ideas it represents are more fuzzy or only partially defined.

Ontologies and other formal methods to express interpretative models have two problems further when applied to scholarly work. The first emerges from what has given ontologies so much prominence in the first place -- their connection with Berners-Lee's model of the next generation of the World Wide Web -- the Semantic Web. The semantic web focuses on enhancing information *retrieval* across a huge amount of data. As schraefel *et al* (2004) have noted, the Semantic Web's emphasis on machine-based reasoning actually draws one further away, rather than closer to the task of individual scholarship. Like them, I have thought of tools for scholarship as more closely aligned to Vannevar Bush's *Memex* model -- a tool for collecting, storing and organising information, and one based on a personal and probably rather informal approach.

The second problem is related to the first: that in order to support machine information retrieval effectively ontologies must be highly formal and complex languages, and the amount of work to formalise knowledge of the domain expert (here the scholar) is both substantial and highly specialised:

'Most often it is not the domain expert that formalises their knowledge -- because of the complexity of the modelling it is normally a specialist "knowledge engineer"' (Hill and Drummond 2005)

Unless a scholar has either extensive knowledge about the OWL technology him/herself, or has regular and extensive access to someone who does, s/he is unlikely to be able to take advantage of the formal kind of expression of his/her interpretation that an ontology mechanism such as OWL could offer. This is hardly the context in which an individual scholar can be expected to carry out his/her research. Ontology technologies have been successfully applied in the context of large community driven collaboration settings -- but they do not seem so appropriate here.

Instead of introducing such a formal (and technically challenging) approach to representing the interpretation, *Pliny* provides a 2D paradigm for organising materials where the user can explore his/her collection of notes or annotations, assemble them into thematic groupings, and relate these groupings together -- the beginnings, at least, of an interpretation. We can see in figure 3 that users are encouraged to think about how they might structure their responses to the text from very early on. As we have already seen, even immediately after s/he has read Wittig's article, the reader has noted that his/her observations fall into two categories, and s/he has used *Pliny* to record this.

In fact, the Pliny user is constantly encouraged to think of organising his/her own notes and annotations in the 2D space of the reference or annotation area. Using a 2D space to organise notes, combined with the ability to group them under a named category, is not a new idea with *Pliny*, of course. Even from the time of handwritten notecards there are stories of people exploring possible relationships between their notes by laying them (or piles of them) out on a large, flat surface and moving the piles of card around on the table. The Apple Macintosh introduced folders for organising files that included a 2D space into which a file could be put -- and early Mac documentation even spoke about the organisation potential of this approach. Closer to Pliny is the Xerox NoteCards system, which was developed in the 1980s to "assist in the authoring process" and provided a model for authoring (Monty and Moran, 1986). Furthermore, there has been other work in computing science around the potential of a 2D space for organising and visualizing things. In Hsieh and Shipman's article on their software VITE (2000), they claim that "... people frequently sort items by rough notions of association or categorization" (Hsieh and Shipman 2000, p. 141), and then shortly thereafter they remark that "One natural organizational process has been found to center around manipulations of objects in spatial arrangements." (p. 141). See also a discussion of the use of 2D space to organise materials in Shipman III *et al* (1999), where a piece of software called *Viki* is described. In a subsequent piece on a more recent piece of software called the Visual Knowledge Builder. Shipman III *et al* (2001) presents this concept as a "visual information workspace for analysis", and continues to remark that "users non-verbally express formative interpretations through visual attributes and spatial layouts".

Given the widespread experience of this aspect of organising materials, there is surprisingly little discussion of it within the DH community, perhaps because so much of the technology promoted within the DH community is XML based, and is, deep down, thought of as hierarchical (although XML, because of its linking capability is not actually limited in this way). Indeed, the one dimensional hierarchy as an organisational paradigm goes very deep into the WWW, and into digital technologies of wide use within the DH community. As Fraistat, Jones and Stahmer (1998) (while talking about the significance of the WWW, and

hypertextuality in general, as a potential force for destabilizing the literary canon) state: "one cannot overlook the fact that the most common organizational strategy on the Web is the list." Of course a sequential-hierarchical ordering (extended with hypertextual links) encourages a more formally evident structured representation than, say, a set of paragraphs of prose text that one would find in a traditional scholarly article. However, a truly 2D space is naturally expressive in ways that are quite different from what can be expressed by sequential/hierarchical ordering. Surprisingly perhaps, one of its strengths here is in its ambiguity -- the significance of placing one object *near* another one is both expressive and usefully ambiguous at the same time. Given the balancing act of scholarly interpretation -- trying to juggle organisation against ambiguity, this is, perhaps a good thing.

A second important aspect in the design of *Pliny* is the recognition that the notes the user created while reading will at some point need to appear in more than one context. They will need to remain connected to the source document which stimulated their creation in the first place, but they will also often need to simultaneously appear in structures that represent conceptual or thematic categories that the Pliny user defines as s/he thinks about his/her reading. In Figure 3, the Wittig notes clearly belong to the Wittig reading, but they also fit into the researcher's developing conceptual model -- indeed, the Pliny user has already begun to organise these notes in this way when s/he groups them into the two major themes that interested him/her in the Wittig article in the first place.

Thus, although we might think that the notes we see in the reference area shown in figure 3 *belong* to the Wittig note, in fact *Pliny* uses in all its reference areas an intermediate "reference" object to display a note or other kind of resource rather than the note or resource itself. Any resource may thus be referenced from more than one other Pliny object. The reference area thus owns only a *reference* to the note/reference it displays, not the resource itself; and position of each reference object, its size (and its colour) are artefacts of its context in this particular owning reference area. By separating the context-specific reference to a note from the note, *Pliny* allows the same note or resource to simultaneously appear in other contexts.

The ability to refer to a note in a variety of contexts becomes relevant when, at a later stage in the research, the Pliny user begins to discover themes that s/he wishes to explore. The 2D space can then be used to visualise these larger themes and how they might relate to the evidence-based notes that one took while reading. Figure 4 shows an example of this. The Pliny user has developed an interest in the uses of 2D space as a tool for scholarly research or study, and has observed that a 2D space can be used in a rigorously "Cartesian" way (when the exact (x,y) co-ordination of the object is significant in the object's interpretation), and in a more mathematically informal way -- called here "topological" (where the exact position of objects is not relevant -- what is important is relative positions of objects, and the relationships between them that are then implied).
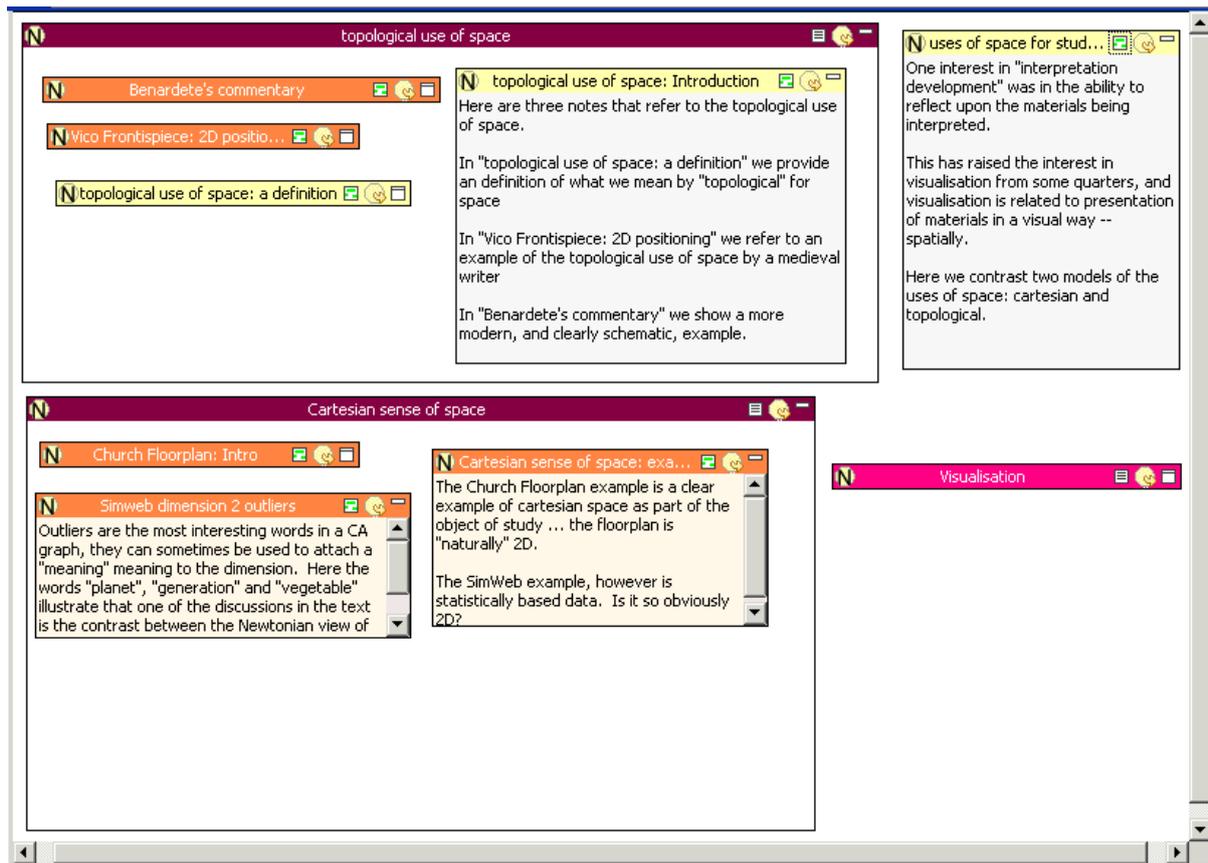
Figure 4

To record this, our researcher creates a new Pliny note called "Uses of 2D space for study", (shown in figure 4) and brings references to the various notes s/he collected earlier during her reading that s/he thinks relate to this topic. Like all notes in Pliny there is a textual area (to the left) and a reference area (to the right). The textual area here contains a brief description of what the note is about, but here the main interest is in the reference area, and the user has shifted the boundary between the text and reference areas far to the left so as to give the reference area by far the largest share of the space. In the reference area s/he recognises his/her thoughts about the uses of 2D space by grouping the notes that belong to the topological category into a Pliny grouping note or container called "topological use of space", and those that to a mathematical use of 2D space into the container called "Cartesian sense of space". S/he then thinks there might be relationships to issues of visualisation in all this, but hasn't yet sorted out exactly what it might be, so s/he has added a reference to a Visualisation note (acting as a holder for issues related to Visualisation) to this page. S/he then adds some further notes as a kind of commentary -- one in the topological group, and then one at the top level because it comments on the either issue of the *uses of 2D space for study* topic.

Each reference object shows in the top left corner an icon that identifies the kind of thing it refers to. So, the reference to the resource *Vico Frontispiece* displays the Pliny icon for an image resource -- the circled "I". This is the image object we showed in figure 2. The Pliny user can ask *Pliny* to navigate to any of the referenced resources by double clicking on the icon, or by selecting the one s/he wants to view and choosing "Open" from the contextual menu.

In *Pliny*, references to resources can be assigned a type, and the type is shown by the colour. Note that the typing refers to the *reference* to the resource, not to the resource itself, reflecting the fact that the particular context in which the resource appears is likely to affect what "type" it should be. The Pliny user can establish any types s/he wishes, and associate any identifying colour to each of them. Here, our user has chosen the purple heading to mean that the object is an instance of the thing named as its container -- here a "use of space for study". The orange headed items are *examples of* its container (so *Benardete's commentary* is an example of *topological uses of space*). The cream coloured heading means to the Pliny user that the note is here acting as a *commentary* on its container, and the Pliny user has chosen the pink colour to assert that the reference object might be in some way *related* to -- s/he thinks that *Visualisation* might be somehow related to the uses of space for study, but at the moment s/he is not sure what the relationship might be. Assigning a type to the reference further enriches the data stored in the diagram and proves to be useful later when the materials being built in Pliny are exported as structured data into a Topic Map. See some discussion about this later in this article.

Two further observations about figure 4:

- First, although the "uses of 2D space for study" materials are stored in one of *Pliny*'s notes, it is not a conventional "note" which would be focused on some textual content. Here, the note's text (shown in the text area to the left of the reference area) is clearly secondary in significance to the reference area. Although *Pliny*'s "note" object is being used to store the object, its interpretative significance is held more in the relationships (sometimes implicit) between the objects it is linked to than in the brief text that is also attached to it. It is, in some ways, beginning to be a representation of a *concept* rather than a *note*, and taking on a bit more structure as a result.
- Second, the use of 2D space to store the references to other related materials begins to look here a little like a visualisation of the concept itself. In both the digital humanities and in computing at large thinking about visualisation has generally been thought about in the context of computer-generated visualisation (generated ultimately through a Cartesian (x,y) data model) rather than, as here, something that is in some sense hand-made (and more similar to the topological significance of 2D space). To what extent is the fact that the data is laid out in 2D space (even though this is done by hand) helping the viewer develop a better understanding of the issues that the layout represents?

Keeping track of materials

It is one thing to organise an object representing a concept such as the *2D Space for Study* from a small number of notes and resources. As the number of objects grows, however, it becomes harder to keep track of them all so that one can marshal them into meaningful categories. In this section we talk about some of the mechanisms *Pliny* provides to help.

Resource Explorer

At the most basic level, *Pliny* provides a *Resource Explorer* that maintains a link (organised by type of resource and by the resource's name) to all resources (including notes) that it knows about. (see figure 5). In the figure we have expanded the "PDF/Acrobat" and "My Bookmarks" folders.
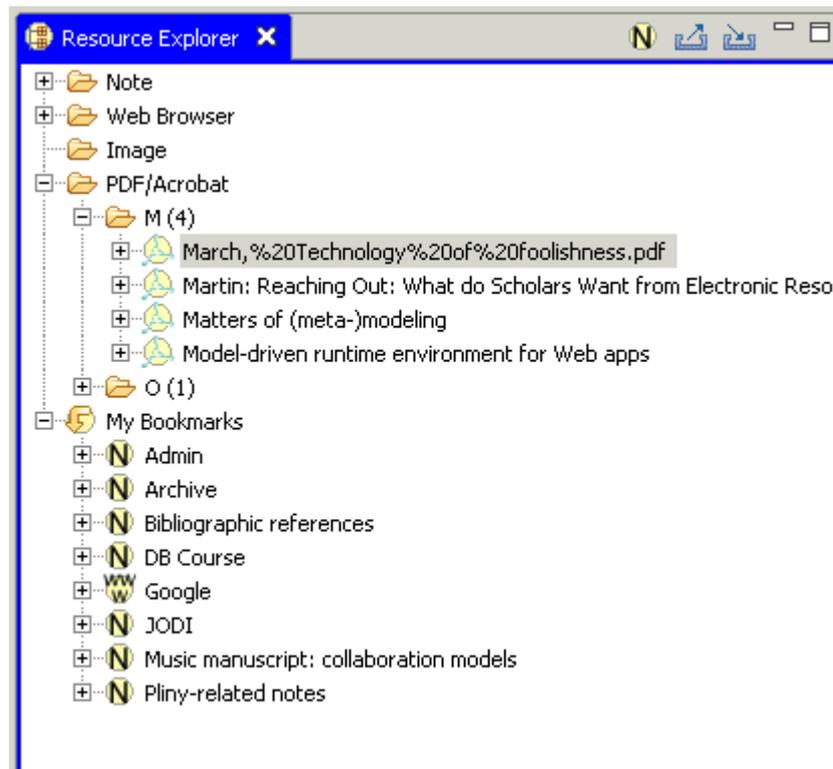
Figure 5

Most of the top level categories (*Note*, *Web Browser*, *Image* etc.) are kinds of resources that *Pliny* knows about. If the Pliny user added other components (as Eclipse plugins) that supported other kinds of resources, these other kinds of resources would appear here as well. A link to all the resources known to *Pliny* will appear filed under one of these type headings.

The *My Bookmarks* folder is different from these others, however, in that it must be consciously constructed by the Pliny user. The Pliny user stores references to any resource in the "My Bookmarks" folder that s/he wishes to access quickly. Thus, it acts as a set of starting points for the user -- not unlike the web browser's Bookmarks list. Often, the user might put very general notes (or other Pliny resources such as web pages or images) into the list of Bookmarks which can act has holders for other notes. In addition, perhaps s/he would link to resources in *Pliny* that provide good starting points into current research interests. My "My Bookmarks" folder, for example, has a mix of items, including (as you can see) an "Archive" note where material attached to now completed projects can be stored, a Bibliographic references item in which I put references to notes that describe items that I have read, a link to a note called JODI where I store materials relevant to the writing of this article, and a reference to the Google webpage -- here because I consult Google so often.

Containment View

The 2D reference object model, with its ability to optionally show notes as containers (shown in figures 3 and 4) provides simultaneously a 2D-spatial and an hierarchical mode of organising materials that allows (perhaps the word should be "requires") the user to keep

track of the notes s/he collects by hierarchically organising them in this way. Although the structure appears there to be hierarchical, as we already know in fact any resource (including a note) can be referenced from more than one other resource/note, so the Pliny user can simultaneously place any resource in more than one hierarchy. A single note, for example, can be referenced from a note providing bibliographic details about the article that stimulated the note in the first place (and thus, in a hierarchy involving things s/he has read), and simultaneously in a separate hierarchy consisting of the themes of interest to the user upon which note seems to touch.

Because of the use of resource references rather than the resource itself in Pliny's reference area, the connections between *Pliny*'s resources is not hierarchical even though the reference area encourages an hierarchical way of thinking about the relationships. *Pliny* allows its resources to be organised into what mathematicians call a *graph*, rather than restricting them to the more limited expressiveness of the *tree*. To help the user manage this, *Pliny* provides a Containment view which allows the user to see beyond the hierarchical presentation provided by the reference area to the network of references connected to any central resource. You can see an example of it in operation in figure 6.
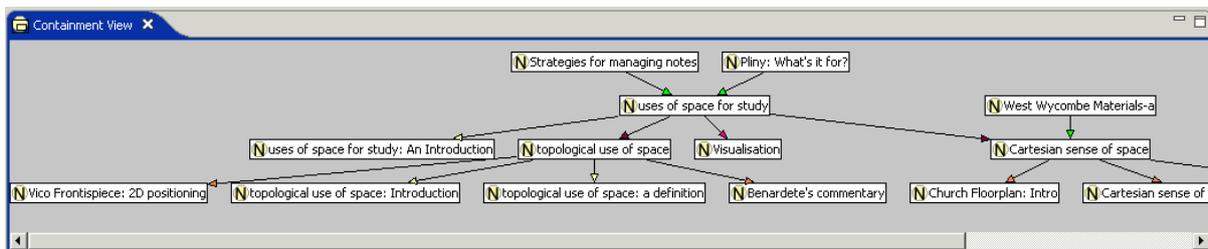

Figure 6

We have already seen figure 4's conventional view of a Pliny note about 2D space, and noted that we have used there the note's 2D reference area to present references to other Pliny resources that are connected to the subject of 2D space. Figure 6 shows the containment view (well, the center part of it at least -- it is too wide to fit in the illustration space provided here) when it centred on this same note. It can best be understood by comparing it to figure 4. In the containment view, resources are shown as boxes, and arrows connect a container resource to a contained resource.Â Although the view was started centred around the "Uses of space for study" note, it has been expanded to include notes that, in turn refer to it. Thus, we can see here that the "uses..." note is referenced in notes called "Strategies for managing notes", and "Current Issues". Furthermore, we can see, after a little further expansion, that the "Cartesian..." note appears not only on our "uses..." note, but is also referenced in a note called "West Wycombe Materials".

At present the Containment view provided by Pliny is rather crude. If Pliny became finished software, it would need to be tidied up, extended and presented in a better fashion. Even as it is, however, it illustrates another way to view the connections between resources, and suggests a way that helps a user to keep track of his/her materials, and to exploit the connections that s/he has already made.

Note Search View

As the notes begin to accumulate, it becomes possible to simply lose track of what you have. The classic way to find digital objects containing text is to ask the system to search for a word or set of words that they contain. For this reason, Pliny contains a Note Search view which allows the user to ask for a list of notes containing a word. See figure 7 which shows the Note Search view in the left frame.
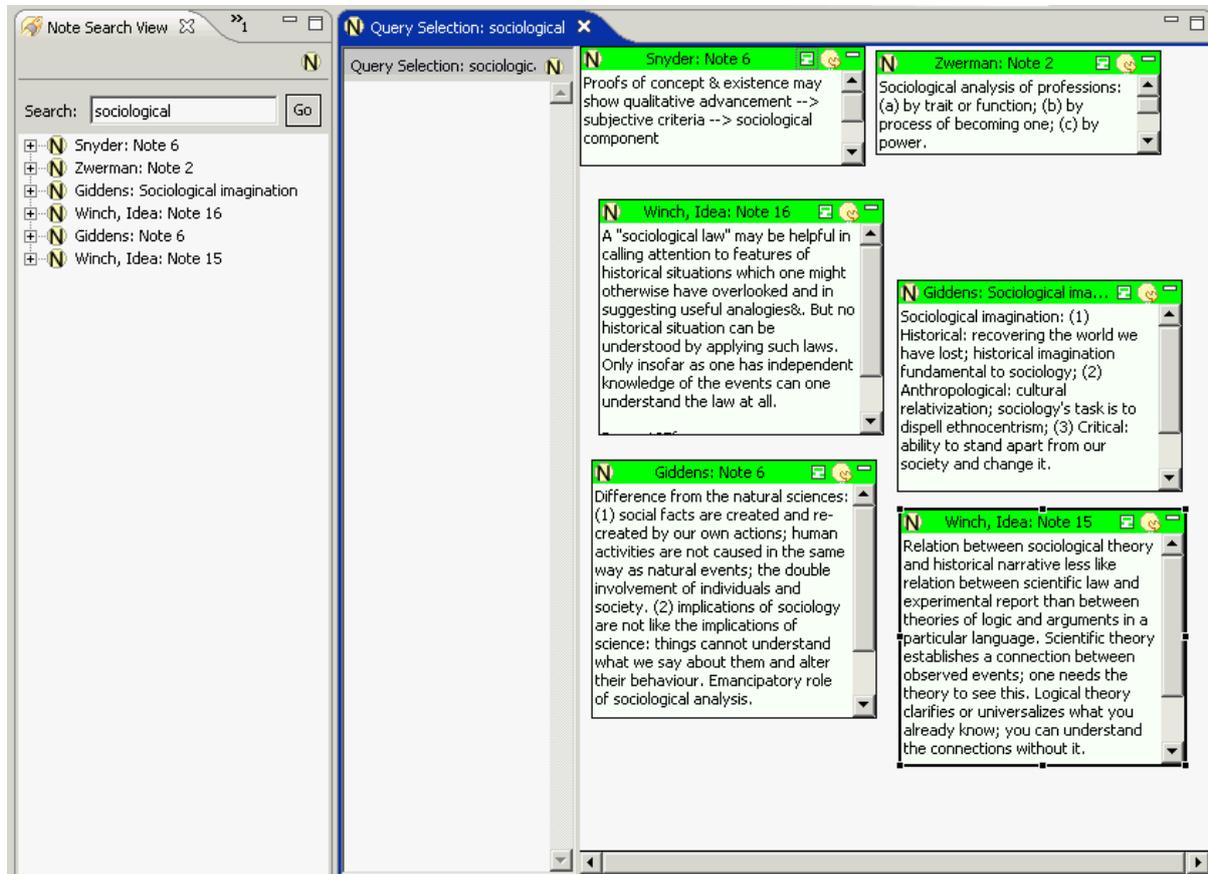

Figure 7

The user has typed in "sociological" as the search word, and Pliny has searched for this word in the text of all the notes Pliny is currently holding (searching is supported with the *lucene* (Lucene 2006) search engine), and found the word occurring in 6 notes, which are listed below the search field. Having found these items, the user can, from within the view, open one or all of them to see them in detail, or can drag one or more of them from this result list to a reference area of another Pliny resource to add a reference to it or them. In addition, by clicking on the little circled N icon (showing in figure 7 above the "go" button), the user can ask Pliny to create a new note, and put references to all the notes it found as a result of its search on this new note automatically. You can see the result of doing this in the larger right side frame of Figure 7. After the user pushed the make note button, s/he has rearranged the reference objects Pliny created so that they can be properly seen. Pliny has given this note the name "Query Selection: sociological". However, the user could choose to take this note as a starting point for a new concept of interest to his/her research. If so, perhaps s/he would rename it to "Sociological", and the user could then from time to time add references to other notes that seemed to be related to this topic.

Publishing from Pliny

Having now discussed the first two phases or aspects of scholarly research (reading with notetaking, and note organising) it is time to discuss the role of Pliny in the third part of research: publishing.

Material might be published from a Pliny archive in two ways.

- First (and more traditionally), the Pliny user might want to write a scholarly article or monograph based on materials s/he has recorded in *Pliny*. If so, the 2D model that *Pliny* provides for organising objects has to somehow be turned into a stream of words -- temporally ordered and hierarchically organised into sections and subsections in ways that characterise much scholarly writing.
- Alternatively, one might consider preserving the more formal nature of the representation of materials in Pliny, and publishing the materials as a network of structured objects.

In both cases one must recognise that publishing involves, by its very nature, the taking of personal ideas and making them public. The Pliny user uses *Pliny* first as a tool to record *personal* annotations. What are the issues that arise when the user takes this personal collection and tries to organise it for public view? One of the issues noted in papers by Catherine Marshall and others (see, for example, Marshall 1998), is that annotation/notetaking for personal use is significantly different in character from what is done when the results are for public use. In (Marshall 1998) she characterises notes for private use as often more cryptic, more informal, and perhaps rather more loosely connected to their anchor than those that are intended for public viewing (this is Marshall's "explicit v. tacit" distinction rather than her "published v. private"). Thus, in moving from the personal space of *Pliny* into the public space of either a published article or a network of connected objects the user may well find that s/he has work to do -- material that is recorded for personal use has to be expanded, tidied up and tightened up for public presentation.

Publishing an article

I have found that when I am using *Pliny* to marshal materials that will appear in a prose article I am writing, I create a special main note for the article I plan to write, and then I drag references to materials I have that are relevant to the subject onto this article note. Eventually, I have in one place on my "article note" a set of references that represent the major themes of the article I plan to write. Then I can juggle the references in the 2D space of the note's reference area until I can begin to see how they will organise into a textual presentation. I have found this to be a highly useful way to plan out the contents of and ordering of materials in an article. However, where can I go from there?

At present *Pliny* provides a rather rudimentary mechanism to export this "article note" into a textual, hierarchical document that can act at least as the starting point for the article that must eventually appear in prose. The exported material often looks more like a set of nested bullet points than anything else. Although it is nice to start off transforming the materials I have assembled into an article by seeing all the materials I put in my Pliny main-article note together in a single text document, I always find that there is a substantial amount of text editing work involved in getting the cryptic note text I have exported from *Pliny* into a workable order so that the text begins to read like an article. In the end it usually happens that the connection between the article text that eventually emerges and the materials in *Pliny* has become perhaps rather tenuous.

It seems clear that *Pliny*'s 2D space model is useful when initially thinking about the material one is reading, and continues to help when these notes are being organised into significant themes or categories. It is *still* helpful during the early thinking about the article text, but currently does not help so much at that point when the note collection must be transformed into polished text for public presentation. We obviously need to think more about what tools could be added to *Pliny* that would facilitate the transformation of the materials from the 2D form into an ordered hierarchy of an article while the material is still in *Pliny*, and that would allow the materials to be both better ordered for presentation as article text, and simultaneously preserve the links from the text into *Pliny*'s network.

Publishing structured data

Earlier in this article I mentioned that Ontologies and similar technologies are perhaps not appropriate tools for a scholar who is in the process of developing an interpretative model for materials that s/he is working with. However, I expressed the hope in Bradley (2003) that as a scholar worked over his/her materials as a model that over time a more formal structure would emerge. As noted earlier, Pliny's containment model, which allows references to objects to be contained in other objects, maps to a mathematical *graph*, and a graph has a highly formal definition which relates well to aspects of ontologies and similar technologies. Furthermore, Pliny's ability to allow the user to associate a type (*is-a*, *example-of*, *comment-on*, etc) to a reference, provides a way to increase the formal content of Pliny's note/resource structure in ways that map well into Topic Maps as "associations" (for an introduction to TM, see Popper's *The Tao of Topic Maps* (2001)). Is it possible, then, that *Pliny* can support the creation of a formal enough model of materials to be useful when expressed in terms of an ontology or topic map?

Currently *Pliny* provides a prototype export mechanism that can express a collection of Pliny notes and resources into a Topic Map. Although the exporting mechanism preserves containment and reference typing it ignores topological aspects of Pliny's 2D modelling space such as "nearness" -- indeed, I assume that as the conceptual model firms up in the Pliny user's mind, the use of nearness to represent associations between objects will be replaced eventually by the *naming* of an association, and its representation by the containment of related Pliny references. Of course, at this early stage in *Pliny*'s development it is hard to tell if this is a reasonable assumption or not. To explore this further I have begun to assess the topic map representation that *Pliny* can create to see to what extent what comes out is actually a useful representation of the material in *Pliny* in this new Topic Map context. There is, however, much more work to be done here that, I hope, will result in time in another paper.

Closing Thoughts

Pliny raises two major themes related to computer support for scholarship:

First, that although the WWW has been an important driving force for much thinking about scholarship in the humanities, it can only be a part of the entire story. The web has focused the attention of the digital humanities scholar on the nature of resources to support humanities scholarship, and has resulted in work that has shown how resources for the humanities can be made potentially significantly more useful. However, with the WWW comes the browser, and with that the limitation of the browser user as "client" in the "client/server" model of the WWW. The limitations that browser technology imposes on the resource user (generally in

the interest of security) are strong and they often restrict the user's use of these digital resources to viewing and printing. There has been recognition of the need for a richer kind of interaction between resources and users in the growth of interest in more complex browser interactions that have been enabled by technologies such as AJAX. This way of "thickening the client" is, of course, enormously useful and has the potential to significantly enrich the user's experience of the resource. However, unless browser technology is changed even more substantially, the separation between a remote resource and a resource user must remain.

Pliny is not a thick client on top of the WWW. It is more of an attempt to blend remote and local materials, and blend the operation of reading and interpreting them on the local machine, under the control of the scholar whose materials are being managed. In placing itself primarily on the user's personal machine, it places itself close to where the actual research work is done. In this, at least, *Pliny* is meant to reflect actual conventional scholarly practice.

This leads us to the second message I wish to convey about Pliny. Much of the work on the development of tools for Humanists has focused on the power of the computer to transform data and present it to users in new ways, and ways that the human reader could not produce by hand except with great difficulty. Transformation tools ranging from the KWIC concordance to the recent excitement of the transformation potential of XML with XSLT emphasizes this ability. In contrast, Pliny seems to do very little in the way of data transformation -- instead it acts more like a "clerk" (a term Englebart used to describe aspects of his Augment system (Englebart 1962, section II.B.8), helping the user in the task of organising the materials, rather than doing much of the organisation for them. In spite of the fact that Pliny doesn't do the organisation for a user, it still does useful things and frees the user to think differently about his/her materials. In the same way, a word processor doesn't write for you, but is still a useful tool, and provides many new ways to think about the task of writing.

Pliny is still work in progress -- indeed, several paradigms expressed by Pliny (such as the use of 2D space as an organising tool) are based more on ideas gleaned by the author from the limited amount of available research on how scholarly work is carried out, and on personal discussions with a handful of scholars, rather than on, say, solid experimental trials. However, it is hoped that *Pliny* can at least draw attention to a set of issues that could, if further researched and developed, significantly enhance scholarly engagement with digital materials and would make available a technology that would truly encourage the scholar to knit together digital materials with his/her own ideas in ways that are typical of true scholarship.

References

- Altick, Richard D. (1963). *The art of literary research*. New York: Norton
- Bates, Marcia J. (1996). Document Familiarity, Relevance, and Bradford's Law: The Getty Online Searching Project Report No. 5. *Information Processing & Management* 32 (November): 697-707.
- Bradley, John (2005). "What you (fore)see is what you get: Thinking about usage paradigms for computer assisted text analysis". *Text Technology* Vol. 14 No 2. pp 1-19. Online at http://texttechnology.mcmaster.ca/pdf/vol14_2/bradley14-2.pdf (Accessed Sept 2006).

- Bradley, John (2003). "Finding a Middle Ground between 'Determinism' and 'Aesthetic Indeterminacy': a Model for Text Analysis Tools". *Literary and Linguistic Computing* Vol. 18 No. 2 pp. 185-207.
- Bradley, John and Paul Vetch (2007). "Supporting Annotation as a Scholarly Tool -- Experiences from the Online Chopin Variorum Edition". *Literary and Linguistic Computing* Vol. 22 No 2. pp 225-242.
- Brockman, William S., Laura Neumann, Carole L. Palmer, Tonyia J. Tidline, (2001). *Scholarly Work in the Humanities and the Evolving Information Environment, a report from the Council on Library and Information Resources* (Washington DC: Digital Library Federation, Council on Library and Information Resources). Online version at http://www.diglib.org/pubs/dlf095/ (Accessed March 2007).
- Busa, Roberto (1990). "The Annals of Humanities Computing: The Index Thomisticus". *Computers and the Humanities* Vol. 14 (1990) No 2. pp. 83-90
- Bush, Vannevar (1945). "As we may think". *Atlantic Monthly* July 1945. Online at http://www.theatlantic.com/doc/194507/bush (Accessed September 2005).
- Cunningham, Richard (2007). "Re: 20.475 fixing the MLA's problem, or what should the Town Crier cry?". A post to the Humanist discussion list Vol. 20, No. 478, Friday 2 March, 2007. Online through Humanist archives at http://www.princeton.edu/humanist/ (Accessed March 2007).
- Eclipse (2007). *Eclipse - an open development platform*. Online at http://www.eclipse.org/ (Accessed March 2007).
- Fraistat, Neil, Steven Jones, and Carl Stahmer (1998). "The Canon, The Web, and the Digitization of Romanticism." *Romanticism On the Net* 10 (May 1998). Online at http://users.ox.ac.uk/~scat0385/rcron.html. (Accessed January 2007).
- Gabriele Griffin (ed.) (2005). *Research Methods for English Studies* (Edinburgh: Edinburgh University Press).
- Golovchinsky, Gene, Morgan N. Price, Bill N. Schilit (1999). "From Reading to Retrieval: Freeform Ink Annotations as Queries". In *Proceedings of the SIGIR '99 Conference*. ACM. pp. 19-25.
- Hill, Duncan and Drummond, Nick (2005) *A Practical Introduction to Ontologies and OWL*. A printed tutorial prepared by University of Manchester (UK): CO-ODE group. See similar materials at http://www.co-ode.org/resources/tutorials/ (Accessed March 2007).
- Hsieh, Hao-wei and Frank M. Shipman III (2000). "VITE: A Visual Interface Supporting the Direct Manipulation of Structured Data Using Two-Way Mappings". In *IUI 2000 Conference Proceedings*. ACM. pp. 141-8.
- Lavignino, John (1997). "Reading, Scholarship, and Hypertext Editions". *The Journal of Electronic Publishing*, Sept 1997 Vol 3 No 1.
- Lim, Lai H, Izak Benbasat and Peter A. Todd (1996). "An Experimental Investigation of the Interactive Effects of Interface Style, Instructions, and Task Familiarity on User Performance". In *ACM Transactions on Computer-Human Interaction*. Vol. 3 No 1. pp. 1-37
- Lucene (2007). Apache Lucene website. At http://lucene.apache.org/ (Accessed Feb 2007).
- Marshall, Catherine C (1998). "Towards an Ecology of Hypertext Annotation". In *Proceedings of Hypertext 98*. ACM. pp 40-49.
- McGuiness, Deborah and Frank van Harmelen (2004). *OWL Web Ontology Language Overview*. Online at http://www.w3.org/TR/owl-features/ (Accessed Feb 2007).
- Monty, Melissa L. and Thomas P. Moran (1986). "A Longitudinal Study of Authoring using NoteCards". In *SIGCHI Bulletin* Vol. 18 No 2. ACM. pp. 59-60.

- O'Hara, Kenton and Abigail Sellen (1997). "A Comparison of Reading Paper and On-Line Documents". In *CHI '97 Proceedings*. ACM. pp. 335-42
- Palmer, Carole L. and Meilssa Cragin (2007). "Scholarly Information work and Disciplinary Practices". In ARIST Vol. 42 (forthcoming).
- *Pliny* (2006). Project website http://pliny.cch.kcl.ac.uk/.
- Popper, Steve (2001). *The Tao of Topic Maps*. Online at http://www.ontopia.net/topicmaps/materials/tao.html (accessed February 2007).
- Potter, R. G. (1988). "Literary Criticism and Literary Computing: The Difficulties of a Synthesis". In *Computers and the Humanities,* Vol. 22 (1988), pp. 91-97.
- Price, Morgan N. Bill N. Schilit, Gene Golovchinsky (1998). "XLibris: The Active Reading Machine". In *CHI '98 Proceedings.* ACM. pp. 22-3.
- Rees, Michael J. (2001). "Evolving the Browser Towards a Standard User Interface Architecture". In the report on *the Third Australian User Interfaces Conference* Melbourne Australia. In John Grundy and Paul Calder (eds). *Research and Practice in Information Technology*, Vol. 7. Australian Computer Society. pp. 1-7.
- schraefel, m.c., Leslie Carr, David De Roure and Wendy Hall (2004). "You've Got Hypertext". *Journal of Digital Information*. Vol. 5 Issue 1. Article No. 253, 2004-07-16. Online at http://jodi.tamu.edu/Articles/v05/i01/schraefel/ (accessed Feb 2007).
- Shipman, Frank M. III, Catherine C Marshall, Mark LeMere (1999). *Beyond Location: Hypertext Workspaces and Non-Linear Views*. In *Proceedings of Hypertext 99.* ACM, pp 121-130.
- Shipman, Frank M. III, Hsieh, H., and Airhart, R. (2001). "Analytic Workspaces: Supporting the Emergence of Interpretation in the Visual Knowledge Builder". Online at http://www.csdl.tamu.edu/~shipman/vkb/vkb.html (accessed May 2007).
- Smith, J.B. (1978). "Computer Criticism" in *Style* Vol. 12 No. 4 pp 326-356.
- Summit on Digital Tools for the Humanities (2006). *A report on the Summit on Digital Tools*, (Charlottesville: University of Virginia) September 2005.
- TEI (2006). *The Text Encoding Initiative*. Website at http://www.tei-c.org/ (accessed Feb 2007).
- *Tinderbox: The Tool for Notes* (2007). Software product. Website at http://www.eastgate.com/Tinderbox/ (accessed July 2007).
- Topic Map (2002). *ISO/IEC 13250: Topic Maps*. International Organization for Standardization: Joint Technical Committee 1 JTC1, Information technology, Subcommittee SC34, Document description and processing languages.
- Trigg, Randall H. (1996). "Hypermedia as Integration: Recollections, Reflections and Exhortations". Keynote address to Hypertext '96 conference, Washington DC. Online at http://www.workpractice.com/trigg/HT96-keynote/default.html (accessed July 2007).
- Wittig, Susan (1978). The Computer and the Concept of Text. *Computers and the Humanities*, Vol. 11. pp. 211-215.
- *Zotero* (2007). Software product. Website at http://www.zotero.org/. (accessed July 2007).