

Thinking about Interpretation: *Pliny* and Scholarship in the Humanities

John Bradley

john.bradley@kcl.ac.uk

<http://pliny.cch.kcl.ac.uk>

Our Digital Humanities community has been arguably in existence for about 50 years, and its pioneers – people like JB Smith and Rosanne Potter – expected right from the start that computing would have an enormous effect on scholarship. In fact, however, our effect on how most scholars work has been very small. Although tremendously innovative techniques have been developed by members of our community; few, if any, scholars from outside the DH community have taken them up. In the report on the *Summit on Digital Tools in the Humanities* held in September 2005 at the University of Virginia, the writers remarked 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”, and they went on to claim that humanists are on the “verge of a revolutionary change” (*Summit* 2005; pg 4) that they claimed would be triggered by digital technology.

This author, unfortunately, is not so optimistic as the writers of the *Summit* report about the future impact of the Digital Humanities when it comes to current research practice by non-Digital Humanists and sees little to suggest to him that anything significant that we are doing is even on most scholar’s research radar. However, although this claim seems to be rather overly optimistic, something else that was said at the Summit *did* relate more directly to my interests and beliefs: and it had to do with the recognition that there was a potential for new computer tools to support the development of scholarly interpretation.

Of course, we have many of the tools developed in the DH are intended to support certain aspects of interpretation development already. Indeed there is a stream of tool-building initiatives that started long ago with Text Analysis tools such as OCP and TACT and have continued up to today with the sophisticated tools such TAPoR and as those that have emerged from NORA. All these tools are based on the idea of using the computer to introduce new ways to view texts. The expectation is that these new ways will trigger new thoughts in the mind of the user, and in turn produce new interpretation insights. For certain kinds of research there is lots of evidence to suggest that this does indeed happen, particularly if the research is heavily oriented towards a careful study of the exact words used in the studied materials.

However, there is another way to think of tools to support interpretation – and from the discussion in the Interpretation subgroup at the Summit I think that this second way was closer to what many there had in mind: to assume that the interpretation emerges *in the scholar’s mind* simply as s/he reads his/her texts, looks at his/her art objects, etc – and that the computer can help by simply acting as a structured repository for these ideas. In this context there is no fancy displays, no seeming magic about what the computer has done to the text – perhaps no “wow” factor. Indeed, the computer doesn’t seem to *do*

very much of anything – its job seems to be to simply act as a clerk which merely records the scholar’s ideas and presents them back again whenever s/he wants to see them.

In some ways, TEI similarly focuses on the role of merely capturing an interpretation. From the very early days of TEI there has been a stream of thinking in the TEI community that claims that scholarly markup could be used to record complex scholarly ideas within a text. Indeed, many readers might well recall presentations by Michael Sperberg-McQueen and others (see Sperberg-McQueen 1991; p 36) for a particularly striking example of this) that tried to convince scholars that markup could be about more than simply recording non-controversial material about a text, and that it therefore deserved their attention. However, a case can be made that TEI’s orientation – towards markup of a text as the sole way to record and express these scholarly insights – does not entirely meet the needs of scholarly research activities, for reasons that, hopefully, will be more evident later in this article.

At the same time that TEI was under development quite a different approach to computer supported research of text was underway. Tools were developed within the Social Sciences to support their approach to textual interpretation called “qualitative analysis” (see some further discussion about some of the significance of these tools within the Social Sciences in Kelle 1997, and some of my thoughts about their significance for humanities scholarship in Bradley 2003). The first prominent piece of software of this kind was the redoubtable Nud*ist program, and this field has continued with more recent, and more complex, software like NVivo and Atlas.ti. *Pliny* takes up a similar theme. Indeed, several significant ideas from software such as NVivo and Atlas.ti appear in *Pliny*, but they have been adapted somewhat to create an environment that is intended to match more closely the needs of the humanist rather than the social scientist. There has also been work around tools to support interpretation in Computing Science. Work by Catherine Marshall and others and resulting in software such as VIKI (Marshall *et al* 1994) and VKB (Shipman *et al* 1999 and 2001) has been particularly influential on *Pliny*. Many ideas in *Pliny* are also related to those formulated many years ago within the *concept map* community (for an introduction see discussion in Novak and Cañas 2006).

One common thread that is shared among tools as diverse as TEI, Nud*ist and VIKI is the use of the computer as an assistant supporting the intellectual work of to the computer user. This attitude goes back even further than the beginning of markup and the TEI and has spawned many of the important ideas that we see represented in, say, “direct manipulation” tools as diverse as the word processor, the spreadsheet, or the drawing program such as Illustrator. One could argue that it began with the work of Douglas Englebart in 1962 when he formulated in his famous *Augment report* a view of tools that he called H-LAM/T: “Human using Language, Artefacts, Methodology, in which he is Trained” (Englebart 1962, pp. 8-29). The basic premise of H-LAM/T is that all tools that support thinking such as pen and ink, drawing surfaces, etc allow humans in some sense to do what they already do in the way of thinking about problems – but do them more efficiently – at least once they are trained in the use of these tools.

Englebart predicted that computing tools built according to his H-LAM/T principles would not seem *at first* to provide startling *new ways* to do things, and in a part of his report he presents a fictional visit to a place where Englebart's new tools – at the time of writing, in fact, still only existing as ideas in his mind – are in use. In this little scenario, a character Joe is showing a visitor around, and showing him how he uses the computer to support his work. The visitor seems a little disappointed that the role of the computer seems so insignificant– s/he is expecting exciting and radical “new things” – perhaps fancy graphics, say. Joe responds:

“You’re probably waiting for something impressive. What I’m trying to prime you for, though, is the realization that the impressive new tricks *all* are based upon lots of changes in the little things you do. This computerized system is used over and over and over again to help me do little things – where my methods and ways of handling *little things* are changed until, lo, they’ve added up and suddenly I can do *impressive* new things” (*D.C. Engelbart, Augment Report, October 1962 p. 89*)

This is one way to think about what the design of Pliny is trying to accomplish. Pliny does not replace the traditional way to do research by something radically new and different. Instead the computer sits seemingly in the corner and is almost invisible, helping the researcher do things the way s/he always does, but augmenting his/her innate abilities, the way tools do. It fits into a way of thinking about tools described in Bradley 2003 and Bradley 2005.

Existing Scholarly Practice

So, if we are to develop a piece of software that supports existing scholarly practice in the way Englebart describes, we have to develop some clear ideas about what existing practice is. This turns out to be rather difficult to do in any detail. Even as recently as 2005 Gabriele Griffin in her book *Research Methods for English Studies* (Griffin 2005) notes in the introduction that “Until very recently, research methods were not widely discussed in English studies ... – research was what you did, and the best you could hope for was a brief introduction to the vagaries of the library.”, and then goes on to say that “significant numbers of English studies academics in the UK” are still “surprisingly in- or possibly non-articulate about what they do to achieve ... results”. In her survey of studies from the Information Sciences perspective on scholarly practices Palmer (2007) similarly observes that there are “considerably fewer works” that focused on scholarly practice for the humanities than for the sciences.

On the other hand, Claire Warwick, in her wonderful piece for the Susan Schriebman and John Unsworth's *Companion to Digital Humanities* (Warwick 2004), gives us a clue about how to think about the problem when she points out that much tool building in the DH has not taken into account what researchers *actually* wanted to do: to, as she says, “study texts by reading them”. Some writers describe about how scholarly research is done. See Richard D. Altick's (1963) book *The Art of Literary Research*, for example, which seems to be written for a scholar at the beginning of his/her career, and outlines in some detail some aspects of a methodology for notetaking and organisation. We also

have the work of the team (including Palmer mentioned above) that resulted in the 2001 report “Scholarly Work in the Humanities and the Evolving Information Environment” (Brockman *et al* 2001). This is a report from the Council on Library and Information Resources, that I have referred to in previous articles (see Bradley 2005), and which also both prioritises reading as a central scholarly activity, and describes briefly the experience of scholars who have tried to use existing digital tools to capture notetaking related to their reading and turn it ultimately into scholarly writing.

Brockman *et al* describes certain aspects of scholarly reading, pointing out that scholars read over a very broad range of primary and secondary sources, and that the reading ranges in intensity from skimming to intensive study, with key sources often getting particularly intensive attention, and perhaps being re-read periodically over a number of years. Reading across a range of sources, both primary and secondary, is often done minimize the dangers of producing too limited an interpretative view. It seems clear that any computing tool that supports reading and notetaking in only a single source – no matter how central to some particular scholarly community – will be of only limited benefit (although, of course, such a tool may well further the research activities of a community of scholars who are focusing on this text and use the annotation facilities provided as a means of highly focused communication). Furthermore, in spite of the growing number of scholarly digital resources available via the scholar’s browser over the WWW, humanities scholarly reading is still primarily of non-digital materials – materials on paper such as articles or books. Pliny provides mechanisms that support the reading of both non-digital and digital (particularly web-provided) materials.

At the other end of scholarship, of course, is the dissemination of results. In traditional scholarship this takes the form of a written book or article which is still today usually published on paper, and even if it is published digitally on (say) the WWW it generally takes a form – the e-journal – that is not very different from a paper format. But what happens between the reading and the writing? The trivial answer, of course, between the reading and the writing is the development of an interpretation. This three-phase process is summarised in Figure I.

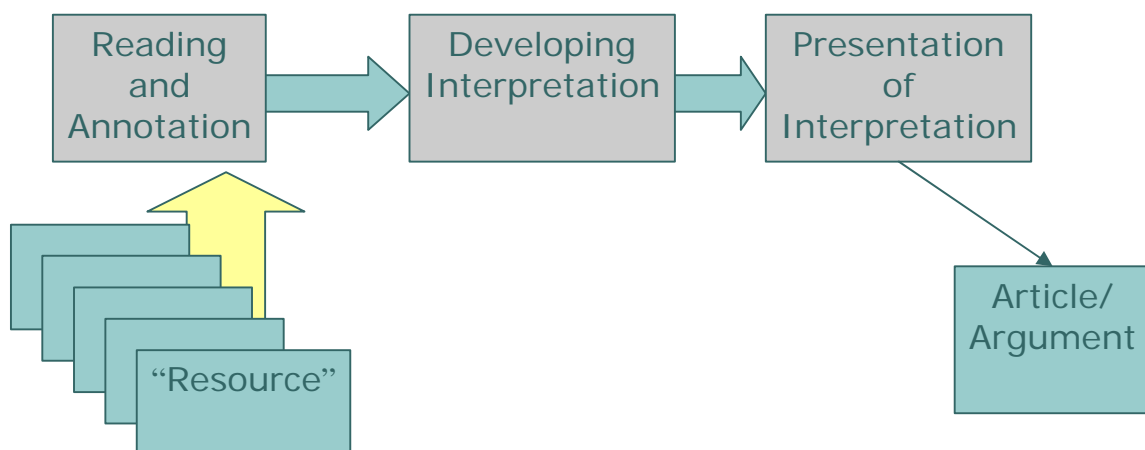


Figure I: Phases in traditional scholarship

For many scholars the development of an interpretation begins as they record what they find interesting in what they are reading in a note or an annotation. Brockman *et al* records (p. 7) that one of the scholars they interviewed stated that to really fix the observation in his/her mind, s/he needed to write a note about it. Altick, in his fourth chapter, also proposes a research methodology that begins with the writing of notes that record reactions to materials being read. Notetaking, although perhaps the first step towards an interpretation is not by itself sufficient, but requires a long period of time, juggling the ideas that have come from the reading until they begin to fall into place. Brockman *et al* also report on the practice of one scholar who described the process of preparing a paper to present his/her research as involving the juggling of bits of pieces of text that represent a number of related ideas until a coherent order appears. Notetaking, and this kind of juggling of notes to discover previously unrecognised patterns and relationships and to stimulate new ideas is one of the long established methods of scholarship.

Pliny as a thought-piece

Pliny is a piece of software designed to support note-taking and -organising. It is named after the classical author Pliny the Elder who was well known in his own day as someone who was constantly writing notes about things he was interested in. *Pliny* – the software tool – is a kind of software clerk that keeps track of the notes you put into it, and helps the user to manage them. Although it is currently available for others to use (you can find at the *Pliny* project's web site (Pliny 2007)), it remains nonetheless a kind of prototype or thought-piece that models an approach to help the user record his/her notes digitally, and then go on to use these notes as a basis upon which an interpretation can be created and perhaps more formally modelled. Hence *Pliny* is not *only* about holding and displaying annotations linked to materials, but also about manipulating them in ways that assist in the development of an interpretation of the annotated body of texts. *Pliny* provides digital support for the kind of activity that one did in a non-digital way when one wrote notes as one read on 3x5 cards, and then eventually tried to juggle them to discover connections between them. *Pliny* is not a web tool – it is a personal computer application that tries to work with resources, including digital ones such as those available over the WWW to support this task.

Pliny supports a second research agenda as well – it provides a model for the development of modular tools that integrate in ways that are richer than other models that have been discussed in the Digital Humanities community. This topic, however, must be the subject of another paper than this one another time.

The Three Phases of Research in Pliny

We will now take a look at what *Pliny* does to support the three phases of research shown in figure I, beginning with reading, or more precisely, notetaking and annotation associated with reading. Figure II shows *Pliny* in operation displaying an image with associated annotations and notes that a *Pliny* user has created.

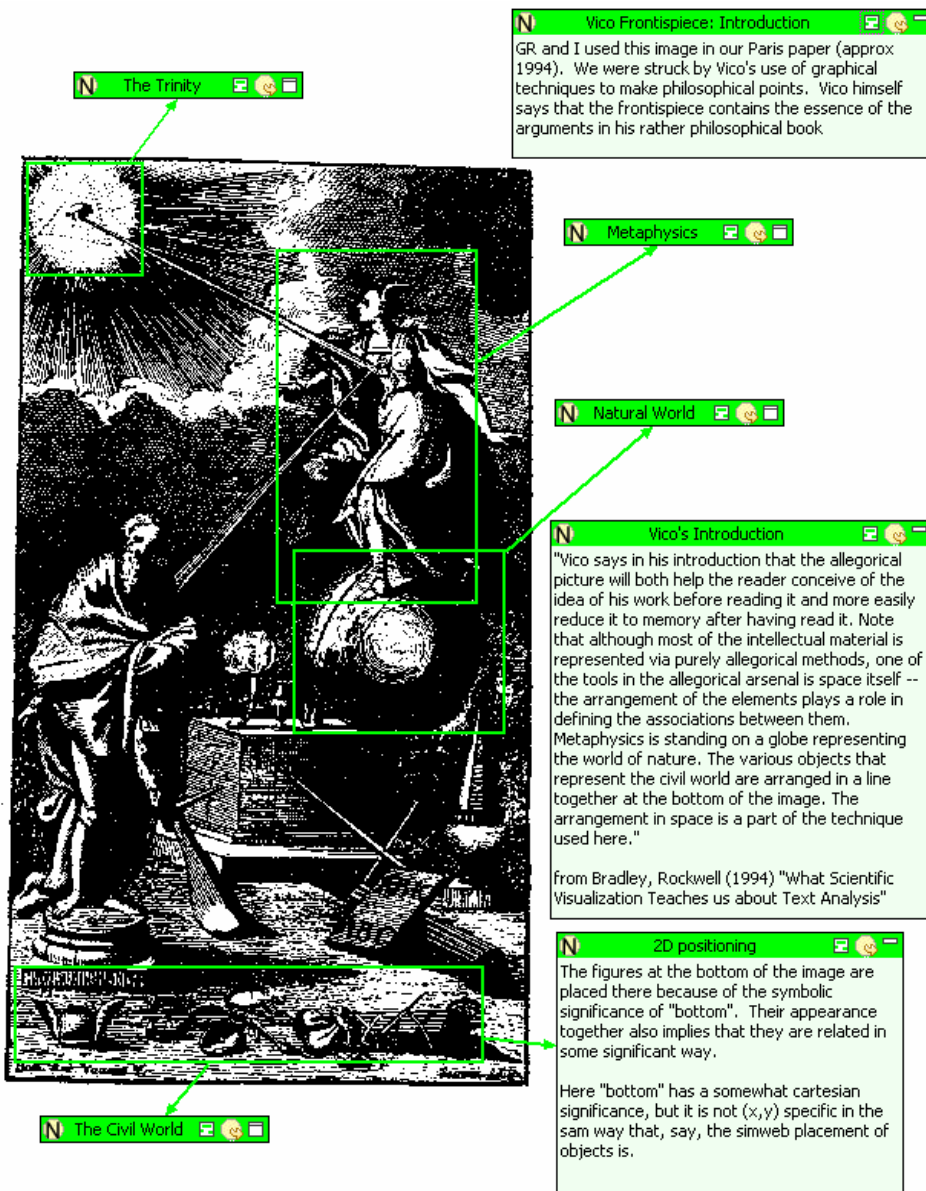


Figure II: Annotation of a image with Pliny

The image here is the frontispiece to Giambattista Vico's *The New Science* (1725). Vico remarks in his introductory text to *The New Science* that this frontispiece shows allegorically what his book is about. Pliny's user has loaded the image and then attached annotations to it that identify some of the allegorical elements in the picture (these derived from work Geoffrey Rockwell and I did for the ALLC/ACH 1994 conference – see Bradley and Rockwell 1994). Pliny's user has also included some explanatory text that, although not connected to any particular part of the picture, comments on it in a more general way. The resulting screen we see here looks perhaps a bit like a sheet of paper containing the image with personal material written on top of it. Like a piece of paper, each time the Pliny user returns to look at this image these annotations will also

appear. However, unlike a piece of paper, the material is more pliable. The user can at any time move materials around, add new notes, edit the text, delete an item or add to it.

In Pliny's terminology, the text the user has added appears in constructs called *notes*. All notes appear in a coloured box with a title area at the top (the particular colours used to display them is determined by Pliny's typing mechanism which is described later in this article). Although all notes have both space for a title and textual content some of the notes appear here with both a title and contents, some are displayed here "minimized" with only their titles visible. Minimized notes have been used here to hold labels for themes or concepts, but the user can easily switch at any time between either the minimized or un-minimized form of any note.

Pliny supports full basic annotation on images, so we have, in addition to notes, anchor objects – here rectangles that link the annotation note to an area on the image. If the anchoring was to be done to a bit of text in, say, an XML file, Pliny's current data model would allow the anchor's position to be described in a way suitable for this purpose – perhaps by means of an XPATH expression. Annotation notes are here visibly connected to anchors by connecting lines. Note that in this figure one of the anchors is attached to two annotations. Pliny does not restrict one note to one anchor, or vice versa. Although shown here supporting annotation on an image, Pliny also supports a similar kind of annotation mechanism for PDF documents.

Pliny *is* a prototype, and there are only a limited number of functions provided to support annotation. For example, the anchor in the current version of Pliny is always a rectangle – one cannot trace out a more complex shape to more exactly trace around a particular feature of the image. This is primarily because there was very little development time available to construct Pliny at all, and it was more important to develop several basic components of an ideal Pliny package that would most clearly reveal some issues about personal annotation than spend more time on enhancing one component or another. Pliny is open source and things like this could certainly be extended if need arose and time permitted. For a similar reason, unfortunately there is no tool in Pliny to specifically support the annotation of marked-up documents in TEI, even though TEI is an extremely important component in the Digital Humanities: to develop a suitable annotator to meet the general needs of TEI-encoded documents would have taken, by itself, more than the year of time available to build all of Pliny. However, even with these limitations, Pliny provides a useful, albeit basic, annotation environment for images that is perhaps more suited to *personal* annotation (the focus of Pliny) than *public*. Private annotation, after all, is by its very nature more informal than annotation that is meant to be public (see Marshall 1998 for more thoughts about this distinction), and if a rectangle does not, in fact, exactly trace out the area of interest, it is probably good enough for personal use.

Pliny supports notetaking for web pages by integrating the machine's standard web browser inside the Pliny screen, and keeping track of the web pages the user is viewing by monitoring their URLs. The browser Pliny uses (its host machine's main browser – usually FireFox or Internet Explorer) does not readily support anchoring of notes to locations inside a web page, so that is unfortunately not available. Instead, Pliny

provides an area to the right of the web page where notes can be recorded. Every time the user returns to the same webpage (the same URL) Pliny recalls the notes recorded on that page and displays them again. So, even though Pliny supports notetaking rather than annotation in full for web-pages, the area at the side of the page acts a bit like an extended margin in a printed book, where perhaps comments have been written.

Figure III shows how Pliny might be used to record notes about a non-digital, print object – here, an article by Susan Wittig that appeared in a 1978 edition of the journal *Computers and the Humanities* (Wittig 1978).

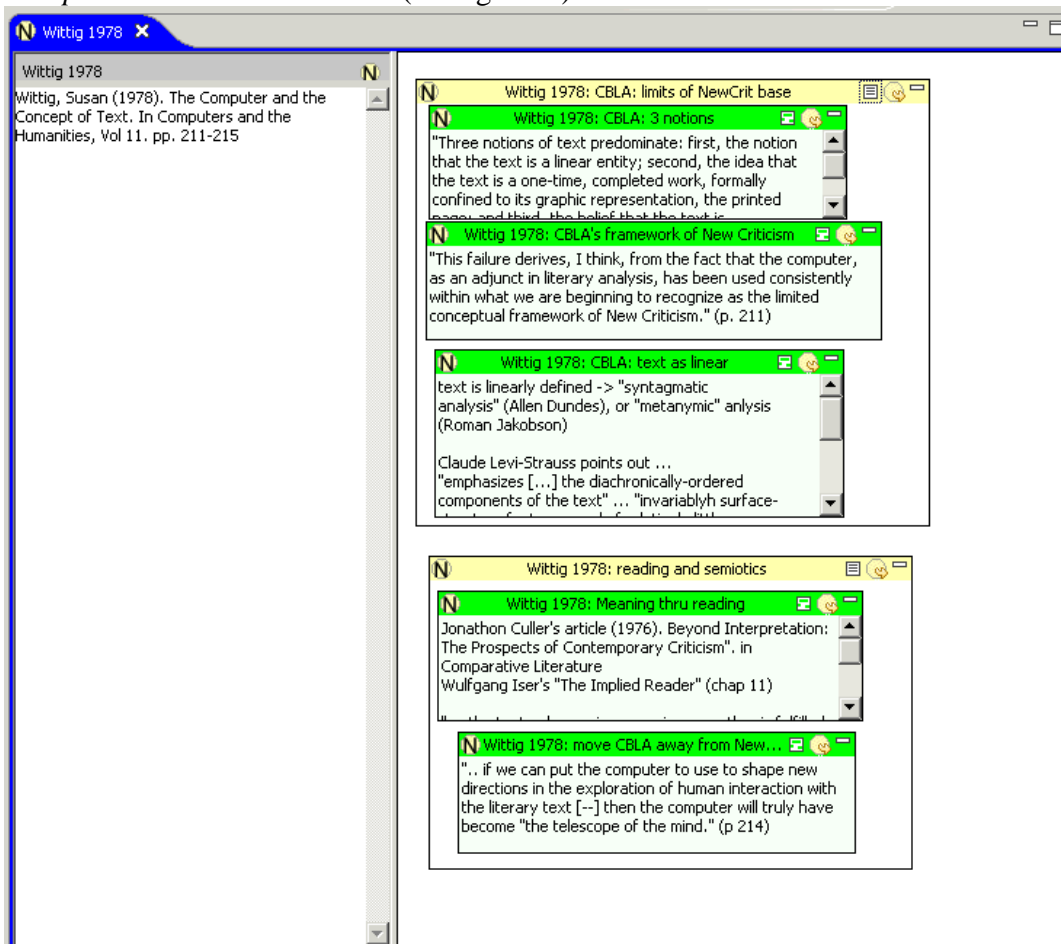


Figure III: Note-taking for non-digital resources

There is no digital object here to connect the notes to, so a Pliny note is pressed into service to simply hold a conventional bibliographic reference to the printed article as its textual content. Figure II shows notes appearing as annotation objects linked to another resource in Pliny (there, an image), and we have seen there also that notes have titles and textual content. In Figure III, however, Pliny's *Note Editor* is invoked to display the note in its own right – not attached to any other object. In the Note Editor the places on the screen where the note's title and content are displayed is different: one can see in this display the note's title (here a Harvard-style citation) and the note's content (the bibliographic reference) displayed in the left hand panel. In this figure one sees for the

first time that Pliny notes also have a 2D area for holding other notes (or, indeed, as we shall see later, references to other Pliny objects) about itself, and that this area can be pressed into service to hold notes that the reader creates to record his/her responses to the text as s/he reads.

Figure III shows the state of affairs after the reader has finished reading the article. Initially, as the text was read s/he recorded in notes things that interested him/her. These notes (five in all) are shown here with the darker headings. Then, upon further reflection, it struck him/her that the five original notes fell into two major groups related to two major themes. Three of the notes could be usefully put together in a group called “limits of the New Criticism base (for computer supported text analysis)”, and the remaining two notes seemed to be focused on a semiotic view of reading. Pliny provides containment objects to allow the reader to organise his/her notes in this way, and s/he then created two Pliny “containment objects” (shown here with a light-coloured header) and placed three of the notes in the first containment object for one theme, and the remaining two in the other.

We’ve now seen that Pliny supports note-taking for images, web-pages, PDF files, and also for non-digital objects. An important part of how Pliny is constructed is that it can be extended to support other kinds of data structures as well. Suppose, for example, one was doing analysis with a statistical tool such as Correspondence Analysis. An important way to use Correspondence Analysis is to examine the various charts that can be created from its results and examine the items (here words) that appear as outliers. This is, then, a natural activity for an annotation tool – perhaps in the way shown in Figure IV, which is derived from work that Geoffrey Rockwell and I did a number of years ago). Here, the graph has been interpreted by the user, and comments of various kinds attached to it.

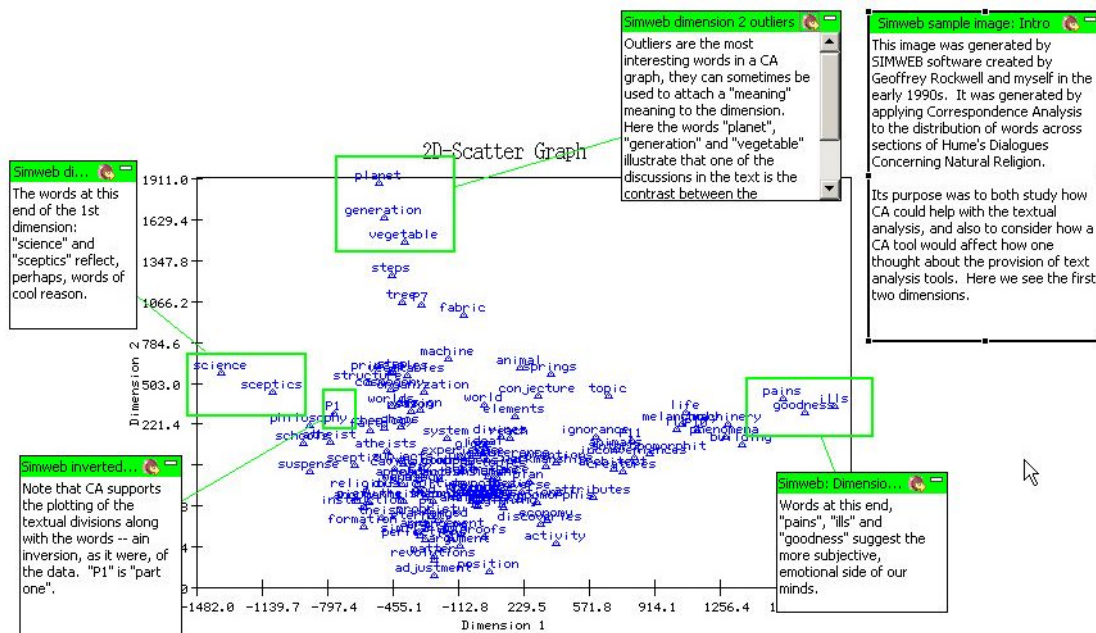


Figure IV: annotation of output from a Correspondence Analysis tool

Pliny is designed assuming that the user might wish to annotate any kind of digital object, not only web pages, digital images, or PDF files. There are many other kinds of digital data that would benefit from annotation supported through Pliny – files with time-dependent data such as video or audio, or perhaps annotation of TEI documents in XML. Any programmer can write a Pliny element that would support annotation of an audio file or a TEI document in such a way that the annotations recorded there would be fully integrated with annotations written for other data files and data types. This is an important aspect of Pliny’s design and is meant to recognise the observation started earlier that humanities scholarship must cross a great range of resources, and, at least by implication, types of data.

Pliny and interpretation development

The previous section of this article has described how Pliny can be used to store notes and annotations created during reading (or study, more generally) of digital and non-digital objects. However, Pliny is not only a thought-piece about how one might support annotation during initial reading (the first phase of research shown in figure I); it also proposes a model of how one can take a collection of notes and annotations created during reading and work it over as an integral act supporting development of an interpretation – the second phase of research identified in Figure I. This is because Pliny supports the notes and annotations that a user creates in a way that allows them to be usable in other contexts – contexts that arise as the researcher thinks about the materials s/he has read. This section of the paper describes how this might work to support interpretation development as it might grow out of a note collection. The challenge in this is in finding ways to allow the user to express a growing awareness of the relationships between his/her notes. Once someone knows the structure of his/her materials, there are ontology-like technologies available such as Topic Maps or OWL/RDF to express this structure. Pliny, however, is meant to operate in the pre-ontological context. How can the computer support the emergence of structure expressed as relationships between notes, concepts, etc? Several of the techniques provided in Pliny are directly influenced by the ideas described in Marshall *et al.* (1994), for their VIKI editor, and the more recent VKB editor (Shipman *et al.* 2001).

Pliny’s support for what Marshall *et al* describe as supporting the development of an “emergent structure” can be best understood by thinking in terms of *affordances* – a term that originally arose out of the work of J.J. Gibson (see a brief summary at Gibson 1971) to describe what humans, seemingly inherently, understand of what can be done with objects in space. Donald Norman – the author of the influential book *The Psychology of Everyday Things* – describes his (somewhat different) usage of the term in Norman 2004, and the term has been taken up by a part of the Human Computer Interface community, where the focus has been on perceived affordances – what the computer user perceives as the potential kind of useful actions s/he can take on the objects on the screen.

Affordance for Interpretation: 2D space

2D space is a natural, perhaps obvious, organisational paradigm for notes and annotation when the object being annotated is spatial already, such as the image shown in Figure II,

or the pages of a PDF file. However, 2D space is applied to more than annotation in Pliny, and one of its central design elements (similar to VIKI's) is found in the use of 2D space as a working space for sorting and organising all kinds of materials. In Pliny the user is constantly encouraged to use a 2D space to layout his/her notes and other objects, and some of these notes might begin to represent emerging concepts. As Marshall *et al* (1994: p. 13) remarks, 2D space “supports the emergent qualities of hypertext structures”, and she then goes on to note that the spatial paradigm has “unique expressive qualities that take advantage of the human perceptual system ... and more generally, spatial intelligence”.

One of the features that 2D space provides is proximity. As I noted in Bradley (2003), exploring materials by laying them out on a 2D space is commonly accepted as a way to explore relationships – think of the approach of stacking note cards with seemingly similar ideas written on them into piles on the bed. The natural temptation is to put similar piles close together even if, as of yet, one doesn't know what the nature of the similarity is.

An example of using 2D space in this way arose when I was organising the notes I created from the Wittig article (see Figure III). Here, the notes have already been grouped together, but initially before I split them into two main themes and while I was trying to form an overview of what I had found interesting in Wittig's article I put the first three notes close together and near the top of the area, and the other two notes close together further down. I could see that the first three were somehow associated, and the remaining two were also associated, even though I initially couldn't put my finger on what the associations were, and suggesting this by proximity seemed entirely natural. The transition to the form you see in Figure III, with the two groupings subsequently identified and named, was then a natural one.

Affordances in Interpretation: naming and containment

Proximity, then, provides the first, and most informal, way to indicate association between objects. However, at some point it is possible that a more formal expression of this relationship, including the naming of it, will emerge – as it did for the Wittig materials. For this reason Pliny supports the ability to first explicitly collect objects together by using containment, and at the same time provide a place to name the collection in the object that is doing the containment. This is exactly what happened with the Wittig notes, and naming and grouping is, of course, one of the basic operations of developing an interpretation. The name acts as a label for an idea, and the collection of objects filed under that name are then associated with it.

Marshall's VIKI software provided a specific object called a *collection* (Marshall *et al* 1994: p. 17) for this purpose. Marshall's article notes that collections contain “an arbitrary spatial arrangement of any number of objects”, and explicitly recognises that the turning of the close placement of object into an explicit collection is a move towards a recognition and expression of an interpretative structure, although the article does not specifically observe that the naming of the collection is also of structural significance.

Pliny, in contrast, does not have a special object to represent a collection – instead, a collection object in Pliny is simply another manifestation of a Pliny note (which has, as we have seen, a 2D space associated with it already). An object that displays a note can display itself in two ways, either by showing the note’s content text – when it will look like an ordinary note – or by showing its associated 2D space instead – when it will look like VIKI’s collection object. Indeed, the two objects *Wittig 1978: CBLA...* and *Wittig 1978: reading and semiotics* showing in Figure III are displays of notes where the two note’s associated 2D space is visible instead of their simple content text. This expression of a containment item, then, as simply a particular kind of display of a Pliny note, rather than as a separate object, reflects my own experience of working with Pliny: sometimes I find that a simple note may begin by only having a name and a brief bit of text associated with it, but might later emerge upon further work to need to turn into a more complex object better represented by a containment display. The Wittig article note is a good example. Initially I might create the Wittig note and place a reference to it on a Pliny object called “Articles about Computer Supported Text Analysis”. Later, when I read the article I discover interesting material in it that furthers my studies and which I record with it in the manner shown in Figure III. The note that references the Wittig article no longer is sufficiently represented by a name and a bit of content text, but has taken on new structure itself – recorded in Pliny as a set of notes that record my response to the article, and then are subsequently further structured as I analyse my response and note that it represents two themes of interest to me.

2D space and visualisation

Although, in its current state of development, the set of objects Pliny supports is small, and its spatial model is also perhaps extremely simple, the combination of naming, containment and 2D layout provides for the creation of visual displays that seem to be useable as a basic representation of concepts and ideas. See, for example, the object created to represent the idea “Uses of space for study” in Figure V, as it appears when viewed with Pliny’s Note Editor.

First of all, the topic “Uses of space for study” is represented as a Pliny note. It has, therefore, a name which the Note Editor shows in the upper left corner. Because it is a note, there is provision for some text to briefly describe the topic. You can see it directly below the name. A more significant representation of the ideas associated with the “Uses...” note, however, is provided by its associated 2D spatial area.

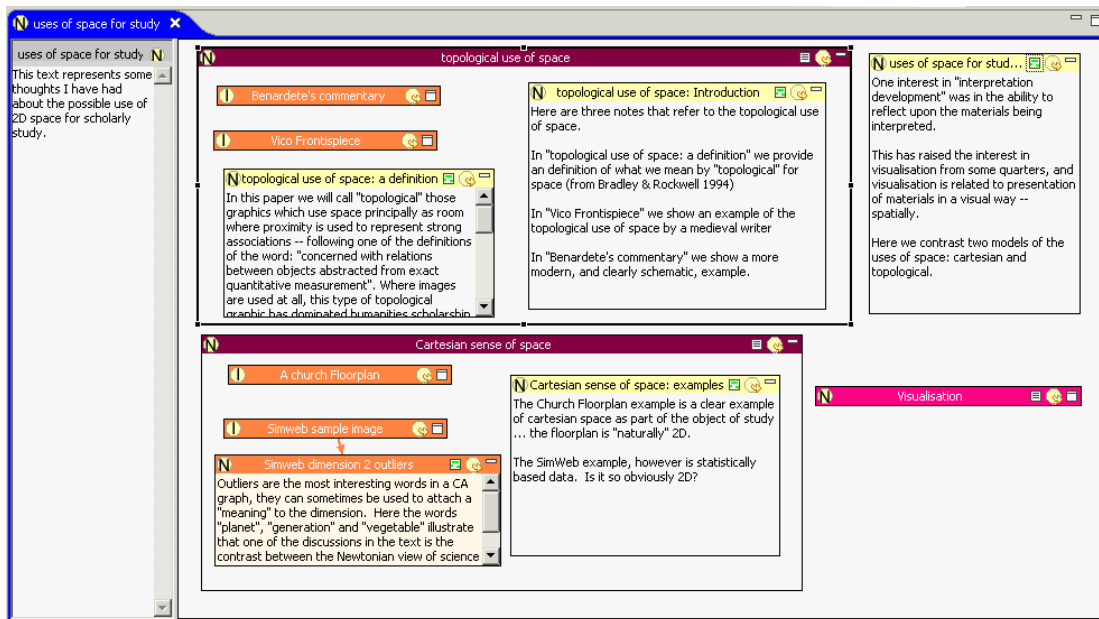


Figure V: The idea “uses of space for study” in Pliny

When thinking about how 2D space is used to support the study of some kinds of objects – the subject matter of this note – it became evident that 2D space could be used in two rather different ways. One is strongly mathematical and derived from Cartesian space – where the location of an point in the a 2D space can be usefully given by a (x,y) coordinate. For certain kinds of objects (such as floor-plans), the mathematical spatial position of the component parts is relevant, and it is also highly relevant for certain statistical data (such as what comes from methods like Correspondence Analysis). This way of thinking about space is the basis of its representation in much of the sciences and in engineering.

There is, however, another way to think about space: where the exact (x,y) co-ordinates do not matter so much, and what is interesting is the relative positioning of objects in that space. This approach to space, which Geoffrey Rockwell and I called *topological* in our unpublished paper for the ALLC/ACH 1994 Paris conference (Bradley and Rockwell, 1994), is how it is used in the Vico frontispiece (shown in Figure II) and represents associations between the items not only spatially but also allegorically. One of the tools it uses to make its allegorical point is spatial.

The spatial layout in Figure V represents these essential elements using the 2D space, containment and naming affordances of Pliny. Two large containers, entitled *topological use of space* and *Cartesian sense of space* are used to hold a mixture of links to examples of the use of space (note the appearance of the item *Vico Frontispiece* among the objects in the topological container – this is a reference to the annotated Vico image we show in figure II). Not all the objects are references to *examples of* the two uses of space. Some of the notes contained in the topological and Cartesian containers are commentaries on the ideas the containers represent, and a further small note outside the two containers provides a few covering words about the issues represented here. Finally, the use of

space in this way seems to have something to do with visualisation, and, although when I created this note I did not yet have anything to say about the nature of the association, it seemed useful to place a reference to a *Visualisation* object here to remind me that there seems to be some sort of connection.

Although much of the discussion about visualisation in the Digital Humanities refers to situations where the computer automatically creates some sort of image that allows the viewer to see new things in his/her data, here there is a sense of visualisation too, even though the layout has been hand-crafted. The use of 2D space shown in this Pliny note actually ends up acting as a sort of visualisation itself, in that it quickly reminds the user of the things that are involved in his theme “uses of space for study”. The 2D-layout, plus containment – even though created by the user rather than the computer – still has elements of computer visualisation about it – this in spite of the fact that the relationships between the objects displayed here are perhaps rather informally expressed.

Pliny Affordance: reference

We have seen Pliny’s 2-D spaces with boxed areas displaying notes or acting as containers. It appears that the notes that are displayed there *belong* to the note in which they are being displayed. However, each display of the note is not really the same as the note itself, but functions as a *reference* to the note. By setting this up as a note reference rather than the note itself, we allow the same note to appear in other contexts than the one where these notes were perhaps originally created – permitting a kind of, let us say, recontextualising of any note created in one context to also appear in a different one.

Although we have focused on using *notes* as the objects that are placed in Pliny’s 2D space, in fact a reference to any object that Pliny recognises can function in the same way – indeed, we have seen in Figure 4 that some of the objects laid out in support of thinking about the uses of space for study are not notes but instead are references to other kinds of Pliny objects. The objects named *Benedete’s commentary*, *Vico Frontispiece*, *A church Floorplan* and *Simweb sample image* are all references to images, which have in turn more annotations attached to them – indeed Figure I and IV show us the content of two of these. Although all these references are to images they could just as easily be to a webpage or to a PDF file – other objects that Pliny explicitly recognises. Pliny is extendable, so if other support for annotation of other kinds of objects (for example, video files) were added, then a reference to them could also just as easily appear here. The references in Figure V to the images act as hyperlinks, and if one double-clicks on the circled I icon at the left end of the reference object one opens the linked image and its annotations.

Since these objects are only references to the image objects rather than the objects themselves they are not actually owned by the note *uses of space for study*. As a consequence, any of these images could turn up in another organising context as well. For example, a note describing the use of allegory in book frontispieces in the 17th century would probably also contain a reference to *Vico Frontispiece* as well.

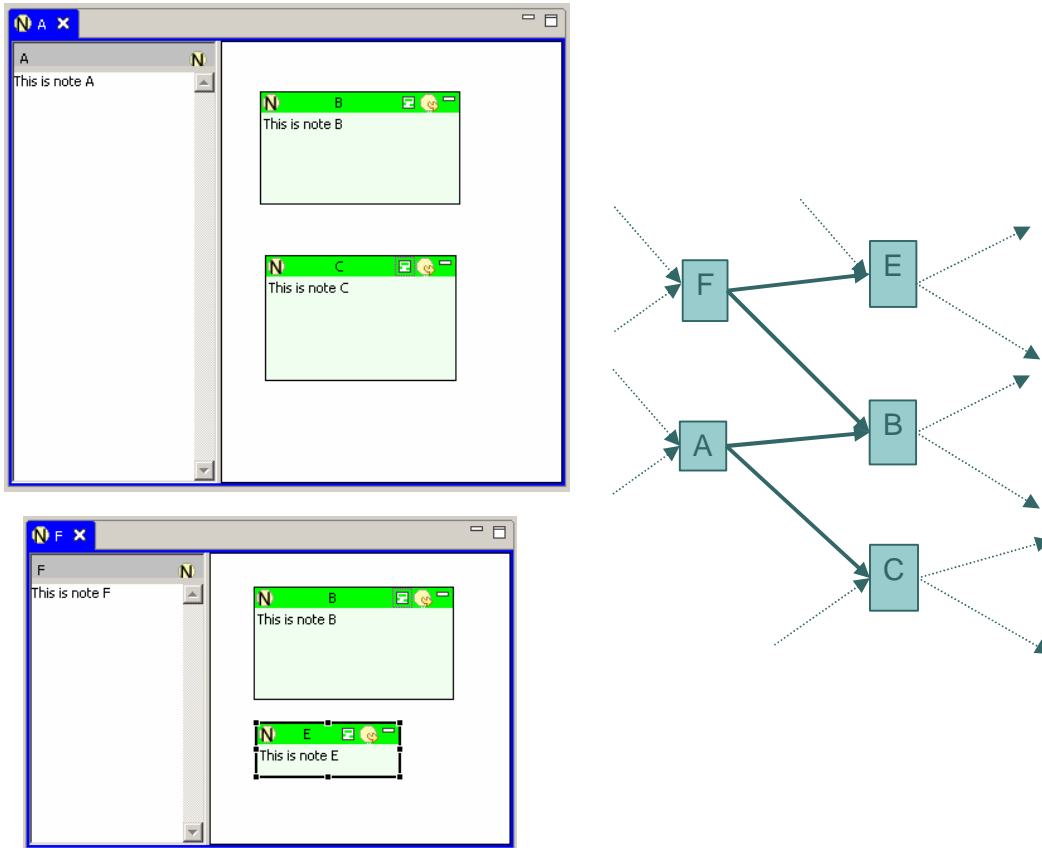


Figure VI: Containment and the Containment Graph

Figure VI shows some implications of the use of reference. On the left we see two notes (“A” and “F”) open in Pliny’s note area. “A” contains a reference to notes “B” and “C”, and “F” contains references to notes “B” and “E”. The right half of figure VI shows this containment structure as a directed graph, where the boxes are Pliny notes, and the arrows show reference by containment. Note the arrow from A to B, for example, and that since a reference to B appears in both A and F, there is a second arrow (starting from F) pointing to B. The fact that there *are* two arrows should make the point that Pliny’s containment model by reference allows the user to construct a set of objects that are connected in what mathematicians call a graph, rather than just a tree. Note as well that although we can see in the Pliny Note Editor that A contains references to B and C we cannot see what other objects might also refer to A, or, indeed, whether other objects refer to B and C as well – this hidden information is shown in the graph side of figure VI by the presence of the arrows with dashed lines.

The 2D view itself, then, while providing an important local view of relationships between items by showing that a reference to an object is contained in another, does an inadequate job of showing all the references that have been made between objects. Understanding these containment relationships by seeing them in a broader context is an important aspect of developing an overall sense of things. To help the user, then, develop this more general view of those s/he has built, Pliny can display a containment view which displays the current containment graph in the neighbourhood of a selected node in

a way similar to that shown in Figure VI. You can see almost all of one of these in figure VII.




Figure VII: A Pliny “containment view”

Here the user has initially opened the graph at the node *uses of space for study* that we have seen pictured in Figure V. The user of the view has control over how far out from this starting node the graph view extends, and in Figure VI the graph has been extended only a few steps beyond the immediate context. However, even then one can develop a better sense of how these particular Pliny object relates to others: for example, we can see from this that the *uses of space for study* item appears on two other items: *Research Items* and *Strategies for managing notes*, and that the *Cartesian sense of space* item appears not only on the current *uses of space for study* display, but also with the item *West Wycombe Materials-a*. Research tools currently available to support social network analysis might well be repurposed to assist in the study of Pliny’s containment graph.

Affordance for Interpretation: Types of reference

We have seen so far that the affordances provided by Pliny (2-D space, containment, naming and reference) provide mechanisms by which the user can enrich the structure s/he stores there. The fifth “affordance” provided by Pliny further enriches the expressive potential by making it possible to assign a *type* to a reference. The Pliny user can decide what types s/he wants to allow (identifying them by name), and Pliny’s Type manager allows for new types to be added, so that any categorisation scheme that the Pliny user finds useful can be created. Figure VIII shows the Type Manager with a few types displayed – one per line. Types have at least a name and a set of colours associated with them. The colour is used whenever a type is assigned to a reference so that by simply seeing the reference the viewer can tell which type of reference it is.

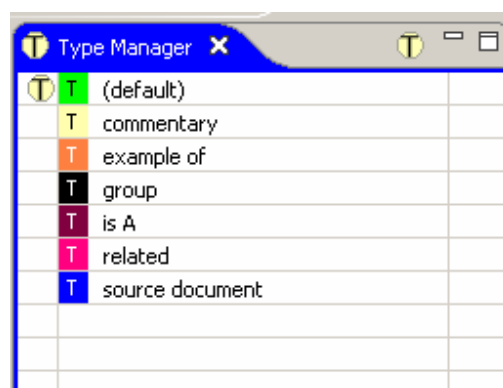


Figure VIII: Pliny’s Type Manager

Note that, unlike VIKI (where notes or collection objects are typed), the type is attached to the reference rather than to the object to which it refers. In effect, the type categorises the *relationship* between the containing object and the contained. This is because the same object can appear in different contexts, and the relationship between the object and its different contexts might also need to vary.

Pliny shows what type is assigned to the reference by colouring its banner. For example (although in print this will not be visible), the “is A” type displayed in figure VIII has been assigned by Pliny’s user the purple colour. In Figure V we would see the banners of the containment reference *topological user of space* and *Cartesian sense of space* displayed in this purple because the note creator has associated the “is A” type to these two items. By making this association the creator is asserting a kind of association between these two items and the containing item *uses of space for study*: that *topological user of space is a use of space for study* and that *Cartesian sense of space is a use of space for study* also.

Affordances for Interpretation: Gesture: Drag and Drop

Gesture is a rather different kind of “affordance” to support interpretation. It is, however, particularly important for annotation because evidence (see (Marshall 1998)) suggests that annotation works best when the mental burden to do it is kept as small as possible. Supporting gesture to get things done in Pliny is important because if it is too much “work” to create a note, then the process of creating it gets in the way of the thought that motivated it in the first place. Drag and drop operations are provided in Pliny because of research in computing science (see Lim *et al* (1996) that suggests that it puts very little mental burden on the user while it is being carried out. Indeed, Lim *et al* introduce the concept of “automaticity” to explain this.

Pliny supports drag and drop in several different situations. One can, of course, drag reference objects around on the 2D space to reposition them. However, one can also drag a reference from one 2D space to another to create a new copy on the new space. Furthermore one can drag a bit of text from (say) your word processor into Pliny to create a note from that text and conversely you can drag a note from Pliny to your word processor to get its text pasted into your word processor document. One can drag the URL from an external browser to get a reference to that page into Pliny.

Considerable more work can still be done to refine the gestures that Pliny supports, and one should not forget this aspect of supporting annotation activities when designing the environment to support them.

Publishing Results with Pliny

The third phase of research is the publishing of results, and traditionally publication results in some sort of text in the form of a narrative and/or an argument. I use Pliny regularly now to support this task. First, I create an item in Pliny that represents the presentation overall, and then I assemble the items I have either stored in Pliny already, and I create new items relevant to the subject at hand. This usually results in a rather dense cloud of reference objects on the presentation object’s space which must then be

brought into an appropriate order. While using Pliny to help me put together a recent presentation, for example, I had assembled this main object and placed 40 items on it representing what seemed to be relevant issues to be presented – a number of these items, in turn, contained within themselves references to other potentially relevant issues. As I juggled these items on the screen it became evident that there were three major themes to be discussed, so I was able to use containment and naming as a way of recognising these themes. Once again items were juggled between these containers, and new sub-groupings were formed within them until I had a general sense of how the materials could be usefully structured. Some items then seemed less relevant and although I didn't remove them from view, I used Pliny's typing mechanism to make it visually evident on the screen that they were perhaps not to be included in the paper, and I also placed them all off to the side.

It was then time to move from the planning to the writing stage. I can simply open the word processor and start writing, keeping at the same time the 2D overview I had created visible on the screen as well to guide me. Pliny's drag and drop facilities meant that any useful bit of text that I had in a note could simply be dragged to the word processor and automatically pasted in. Alternatively, one can ask Pliny to export the items and transform them into a hierarchically organised document that looks like a nested list of bullets. In either case – either when simply using a Pliny 2D space as a guide during writing, or when one explicitly exports the space into a list-oriented document – the problem is one of ordering. How does one transform the 2D concept map-like item into a temporally ordered, structured presentation or argument? There is further work possible here.

In addition to producing scholarly output as a paper, Pliny provides for the possibility of a quite different kind of publishing in the form of a more formal representation of an interpretation – let us call these formal expressions “ontologies” (in the computer science sense of the word, although we are in reality being too informal in our use of the term here for this to be entirely right). As one can see, Pliny really works in a pre-ontological context, but by the time materials have been organised, grouped, named, and typed within Pliny one could argue that one is at least part way down the road towards a formal ontological representation.

Pliny's structures map to a significant extent onto the topic map model (Biezunski *et al* 2002) (for a working understanding of Topic Maps see Pepper 2002). Topic Maps contain *topics*, *occurrences* and *associations*. Pliny doesn't have a specific structure for topics, although as one can see from earlier discussions about Pliny, the Pliny note can take on this function as well as others. A Pliny note where the focus is on its textual content can represent a Topic Map's *occurrence*, and other occurrences arise from Pliny's ability to store references to web sites, images and PDF files (plus, if extended, other kinds of digital objects). Finally, Topic Maps related topics and occurrences by means of *associations*. Pliny's reference objects also define associations. Furthermore, Topic Map associations can be typed, in a way not dissimilar to Pliny's typing mechanism.

As a way of exploring these possible similarities, I have included an export mechanism in Pliny that takes materials in Pliny and maps them to a Topic Map format so that they can be subsequently explored using Topic Map software. The fit between Pliny and topic maps is not exact. Furthermore, Pliny is meant to support scholarship when it is still “pre-ontological” – before concepts and their relationships to evidence from sources have solidified – something that is rather foreign to the Topic Map paradigm. Indeed, to what extent should we assume that what researchers do as they interpret their texts will necessarily result in the end in a formal model that can be expressed using mechanisms such as RDF/Ontologies or Topic Maps? I hope that Pliny can provide one possible framework in which such a question could be considered.

The future of Pliny

Pliny is a piece of software but it is also a thought-piece that models and draws attention to some issues around scholarly interpretation and digital technology. Although Pliny, as a piece of software, is not terribly polished, there are a small number of people who are using it to assist in real scholarly work. I have encouraged them to keep in touch with me about their experiences. I have also had a couple of opportunities to try out Pliny with students in CCH’s teaching program.

As this article shows, Pliny is about personal annotation and note-management. However, it also explores another issue: extending the understanding of modularity in software tools by expanding the nature of interaction between tools beyond that reflected in the development of common file formats that a range of tools all share. I argue elsewhere (see Bradley 2007) that annotation, by its very nature, is not well served by the file-driven view of tool modularity and interaction. Pliny is built using a paradigm specifically developed to provide a richer repertoire of interactions between separately created tools – to encourage and support a larger range of interactions between software.

Pliny’s initial development has been made possible by the generous provision (by Harold Short, head of CCH, and by King’s College London) of research leave for me for a year to build Pliny in the first place. That is now over, sadly, but the resulting software exhibits the beginnings of a framework in which it will be possible to continue to explore Pliny’s two main themes, annotation and interpretation building on one hand, and tool making with richer interaction between tools on the other. There is certainly a lot to do here.

References

Altick, Richard D. (1963). *The art of literary research*. New York: Norton

Biezunski M., Newcomb S., Pepper S. (2002) (ed.). *ISO/IEC 13250 Topic Maps*. Second Edition. International Organization for Standardization. Online at http://www1.y12.doe.gov/capabilities/sgml/sc34/document/0322_files/iso13250-2nd-ed-v2.pdf (accessed October 2007).

Bradley, John (2003). "Finding a Middle Ground between 'Determinism' and 'Aesthetic Indeterminacy': a Model for Text Analysis Tools". In *Literary and Linguistic Computing* Vol 18 No. 2 pp. 185-207.

Bradley, John (2005). "What you (fore)see is what you get: Thinking about usage paradigms for computer assisted text analysis" in *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 (2007). *Making a contribution: modularity, integration and collaboration between tools in Pliny*. Poster paper at DH 2007 conference, Urbana-Champaign: University of Illinois.

Bradley, John and Geoffrey Rockwell (1994). "What Scientific Visualization Teaches us about Text Analysis". Presented at the ALLC/ACH conference in Paris, 1994. Online copy at <http://www.cch.kcl.ac.uk/legacy/staff/jdb/papers/paris.html> (Accessed Sept 2007).

Brockman, William S., Neumann Laura, Palmer , Carole L., Tidline, Tonyia J. (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/df095/> (Accessed March 2007).

Englebart, Douglas C. (1962). *Augmenting Human Intellect: A conceptual framework*. Stanford CA: Stanford Research Institute. Online at <http://www.bootstrap.org/augdocs/friedewald030402/augmentinghumanintellect/AHI62.pdf>

Gibson, J.J. (1971). "A Preliminary Description and Classification of Affordances", available on at <http://huwi.org/gibson/prelim.php> (referenced October 2007).

Griffin, Gabriele (ed.) (2005). *Research Methods for English Studies* Edinburgh: Edinburgh University Press.

Kelle, U. (1997) "Theory-Building in Qualitative Research and Computer Programs for the Management of Textual Data" in *Sociological Research Online*, vol. 2, no. 2, online at <http://www.socresonline.org.uk/2/2/1.html> (referenced November 2004).

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

Marshall, Catherine C., Frank M. Shipman III, James H. Coombs (1994). "VIKI: Spatial Hypertext Supporting Emergent Structure". In *ECHT '94 Proceedings*. New York: ACM. pp. 13-23.

Marshall, Catherine C. (1998). "Toward an ecology of hypertext annotation", in *HyperText 98 proceedings*, ACM: Pittsburgh PA, pp 40-49.

Norman, D. A. (2004). *Affordance, Conventions and Design (Part 2)*. Online at http://www.jnd.org/dn.mss/affordance_conv.html. (accessed October 2007)

Novak, Joseph D. and Alberto J. Cañas (2006). *The Theory Underlying Concept Maps and How to Construct Them*. Pensacola Florida: Institute for Human and Machine Cognition. Online at <http://cmap.ihmc.us/Publications/ResearchPapers/TheoryCmaps/TheoryUnderlyingConceptMaps.htm> (accessed October 2007).

Palmer, Carole L. and Meilssa Cragin (2007). "Scholarly Information work and Disciplinary Practices". In *ARIST Vol. 42 (forthcoming)*.

Pepper, Steve (2002). *The TAO of Topic Maps: Finding the Way in the Age of Infoglut*. Online at <http://www.ontopia.net/topicmaps/materials/tao.html> (accessed October 2007)

Pliny (2006). Project website <http://pliny.cch.kcl.ac.uk/>.

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).

Sperberg-McQueen, C.Michael. (1991), "Text in the Electronic Age: Textual Study and Text Encoding, with Examples from Medieval Texts". In *Literary and Linguistic Computing* (6: 1).

Summit on Digital Tools for the Humanities (2006). A report on the Summit on Digital Tools, University of Virginia September 2005.

Warwick, Claire (2004). "Print Scholarship and Digital Resources" in Susan Schreibman, Ray Siemens and John Unsworth (eds). *A Companion to Digital Humanities* Oxford: Blackwell Publishing. pp 369-370.

Wittig, Susan (1978). "The Computer and the Concept of Text". In *Computers and the Humanities*, Vol 11. pp. 211-215.