Towards an Architecture for Active Reading

Christian Wittern
Kyoto University

Abstract

Electronic publishing, especially publishing for the Web, is slowly becoming the main avenue for the dissemination of scholarly edited versions of primary texts. While this is a welcome development and increases access to high quality resources for casual users, it also poses new hindrances to deeper engagement with texts published in such a way. This article discusses some of the requirements for such a deeper engagement with texts, both in terms of the needed functionality, which includes commenting, annotating, translating, but also linking and enriching the texts and in terms of the needed overall architecture and the desired modes of communication between readers of such texts.

Keywords

Digital Text; Publication; Online Collaboration; Virtual Research Environments

Christian Wittern is Professor at Kyoto University, Center for East Asian Informatics at the Institute for Research in Humanities. 47 Higashiogura-cho, Kitashirakawa, Sakyo-ku, Kyoto 606-8265, Japan. Email: cwittern@gmail.com.

Where we are

During the past 20 years, the digital medium has become the primary environment for the critical editing of texts. While the publishing is still partly done on paper, electronic publishing has also become increasingly common.

Much research and practical effort have gone into the development and maintenance of a digital format that could form a stable foundation for a text in the digital age. The results in the form of the *Guidelines for Electronic Text Encoding and Interchange* (Burnard & Bauman, 2008) have been widely adapted in the community.

Considerably less effort has gone into the question of how the texts thus encoded will be published. There is no common model and indeed not yet a credible theory of reading in the digital age, although several developments towards this have been attempted (see for example Siemens, Leitch, Blake, Armstrong, & Willinsky, 2009). It should also be noted that the interests of creators of the texts, which at the moment also determine how texts are disseminated, do not necessarily match, in fact most surely mismatch the interests of those interested in reading the texts, be they scholars interested in studying and close reading a primary text, or casual readers in search of reading material. It should also be noted that ownership and distribution gets quite a bit more complicated with electronic texts compared to the way it works with printed books (with which it is possible to buy, lend, borrow, and resell one's copy of a book as one pleases).

In the following article, I will first look briefly at acts of reading in the world of print. This will lead me to a proposal for implementing what I will call "active reading" in the digital medium, a proposal that includes both an architectural part and ideas for a concrete implementation. The texts in question here are digital critical editions of primary sources, intended to work as a primary reference for studying the work in question. While the material consists of pre-modern Chinese texts, it is hoped that this method will be general enough to allow a wider application.

TRACES OF READING

The physical copy of a book printed on paper offers on its pages not only the text for the consumption of the reader, but also serves as a canvas to keep traces of the interactions between the book and the reader. These traces can be simply indications of the rhythm of the text, as in Figure 1, or they can be notes the reader puts down when trying to understand a text, which are especially necessary if the text is not in the reader's mother tongue, as shown in Figure 2, or anything in between. One of my teachers used to paste additional paper slips into the page when he ran out of space in the margin, and in addition had a list of noteworthy textual locations. If a book is used in such a way for an extended period of time, such marks will form a record of previous readings and interactions with the book.
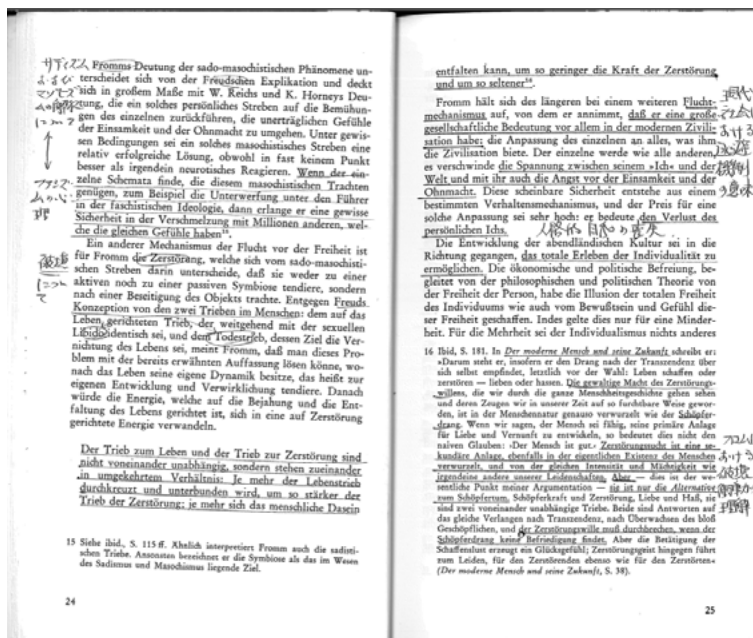
With the advent of digitally distributed editions of texts, the usefulness of such marks has been diminished. For one thing, in most of the distribution forms, it is simply technically impossible for the reader to leave traces of readings in a text. Even where this is possible, like in the case of the Mark Twain project, or where annotation can be achieved through third party tools like the "Awesome Highlighter" or the Open Annotation Collaboration (see below for websites), it is highly doubtful whether all

these tools, together with the supporting infrastructure will stay in place for periods of at least 20 or 30 years. A different approach to distribution and annotation thus seems necessary. It should be emphasized from the outset, however, that this new approach is not intended to *replace* existing models of scholarly publication, but are rather meant as *additional* avenues for publication.

Figure 1: Woodblock print of Zizhi tongjian (Sima, 1084) with interpunction added in red by an unknown reader

Figure 2: Page from Reif (1978) on Erich Fromm, annotated in Japanese by its previous owner

Wittern, Christian. (2012). Towards an Architecture for Active Reading. *Scholarly and Research Communication, 3*(4): 040143, 11 pp.

3

Traces of reading in digital texts

It goes without saying that all kinds of "traces of reading" that are possible in printed texts should also in some way be transferrable to digital texts. An interface that allows the reader to interact with digital texts needs to be created in order for this to be possible. My concern here however, is not with the implementation of such an interface, but rather with an architectural framework that would make the implementation of interfaces possible.

Margins in digital texts are theoretically without limits, so annotation of unlimited length could be added to a text. So among other things, it would be possible to scribble an entire translation of a text into such digital margins. This might seem over-stretching the notion of annotation but has the advantage that the text is automatically linked to the translation.

Whereas printed text might occasionally see revised editions or get a new postface in a later edition, they are for all practical purposes static texts, which do not get updated. Digital texts, on the other hand, are much more likely to be updated frequently, be it only to correct some misprints that have been noted since the last publication. In fact, as a reader of digital text, my fingers become literally itchy if I spot an error and would love to read it in the context of a system that easily allows me to update such a text, for the benefit of later readers. Any system that tries to solve the problem of digital annotations should thus take into account the fact that texts may change, should make sure that changes can be published and picked up by readers, and should also ensure that these changes do interoperate nicely with annotations a reader might have made to her version of the text.

Another key point, which distinguishes traditional, personal annotations from their digital equivalents, is the fact that the digital variant has the potential to be shared. Again, this depends largely on the protocol and underlying architecture of how such a system is set up, but it could be set up in a way that allows annotations to be either completely private, shared in one or more groups, or public. Groups of scholars could thus be set up to collaboratively annotate or even translate a text. An example for such a framework, although based on slightly different assumptions, can be found in (Gerber, Hyland, & Hunter, 2010).

Active reading in a digital context

The proposal made here is of an architectural nature and tries to point out how the goals outlined above can be achieved, namely, how scholarly editions could be distributed in a way that ensures long-term usability and free interaction of the reader.

There are different models that could achieve such a distribution. The model proposed here does not rely on the continued existence of a central authority or the continued existence of a sophisticated infrastructure and is thus designed in a similar way to the Internet itself. It goes without saying, however, that if the infrastructure goes missing, not all of the functionality will be preserved. A scholar already set up with texts will be able to continue to use them.

Distributed Version Control Systems (DVCS)

The enabling piece of architecture for this model is a so-called "Distributed Version Control System" (DVCS). This type of system is currently mainly used in software development, but other uses are already beginning (see for example Digital Papyrology, 2011) and will likely become more frequent.

Software projects typically are distributed across the globe with little or no hierarchy or control among the developers. Central repositories of source code with tight access controls, which had been used widely until a few years ago, proved to be a mismatch to this mode of operation. Consequently, DVCS, which supports a different style of organization, has been developed. There are a number of different solutions, some with a slightly different procedure. For more details, see for example Wikipedia (2012a), *List of revision control software*. Here I will use the program `git` (Loeliger, 2009) as an example, which has been developed by Linus Torvalds for the maintenance of the Linux kernel, as an example.

`Git` distinguishes a remote repository, which is accessible by all developers from a local repository, which is only accessed by the user of the local machine. In addition, there is a place for the actual doing of the editing work, which I will call the "workspace." Both the remote and the local repository can and typically will contain several branches, that is, different versions of the codebase, just as a text might have different versions.

A typical workflow in development would be a developer will "cloning" the repository or a part of interest to him. This cloning action would copy the files together with associated information about the development history and the various branches in existence, from the remote to the local repository, and usually also create a workspace copy. He might then edit and change the content as he wishes, keeping track of this process by committing (i.e., saving) the changes back to the local repository. Whenever he feels like it, he might look at what changes were made in the remote repository and "pull" any changes he wants to incorporate, merging them into his local repository. When he is ready, he might "push" his edits to the remote repository, if he has the rights to do so. However, frequently he will announce his changes to his peers and invite them to "pull" them from his repository and into their own. The maintainer of the remote repository might also have a look and pull any changes that he likes. The changes might then be merged into one of the branches of the remote repository or continue to exist as a branch on the remote repository, which makes them available to all other users there, who might be interested in this.

DVCS FOR SCHOLARLY PUBLISHING

This workflow would translate quite nicely to work on electronic editions if they were made available as `git` repositories. Users could clone texts or repositories they were interested in, and from that time onwards pull any changes the publishers of the edition might have made. If they have made local corrections, annotations, or added more witnesses, this will not be overwritten, but can be merged, which will allow the users to work with their local versions as their research requires, while still keeping track of the remote changes. In addition, they could share their own work, be it annotations or corrections with the editors of the edition, or with groups of other researchers that share a similar interest. In the example above, the text will be

in one branch of the `git` repository and the translation will be in another branch. Additionally, branches could be set up to maintain other versions of the texts, such as those produced by other projects, or to reflect specific historic editions of the texts.

In addition to recording the changes, `git` could keep track of who made the changes, thus allowing for a very fine-grained attribution of the changes to the original editor. It would thus be immediately clear and verifiable who introduced a particular change.

DVCS CAN RECORD MULTIPLE EDITIONS

Another important advantage of a DVCS is that it can record different versions of a text in so-called *branches*. In the case of an existing XML edition it might be desirable to publish such a "branched" version of editions. The first step here would be to set up a repository that publishes the XML "master" version as it is. In the example used here, the master edition contains information about several other editions via a textcritical apparatus, encoded using the parallel segmentation method (see Burnard & Bauman, 2008).

In addition to this, the publisher of these texts, a community interested in these texts, or some other party can transform the texts to a "flat" version, where most of the markup is removed and all editions are represented in branches of the `git` repository, as shown in Figure 3, in this example the versions are represented by their sigle (a shorthand notation used in text critical editing to refer to one of the textual witnesses) in Chinese characters.

Figure 3: XSLT based transformation from TEI P5 to established text and multiple textual witnesses

If a new, additional edition is now added to this repository, instead of encoding a new edition tediously and error-prone in XML, this system makes it possible to simply re-create the edition in exactly the way it appeared, including page breaks, line breaks, additional prefaces or colophons, etc.

As well as recording the additional edition, the system can also be asked to produce a list of differences between editions. If this proves insufficient, a full collation can be performed and a collated edition produced.

**Figure 4: Cloning of branches from CBETA to private repository**

Figure 4 and Figure 5 show schematically how a repository can be cloned and a separate edition added to the private repository. Additional branches, which are meant to hold private "annotations," that is a translation into English and German in this case, are also created.

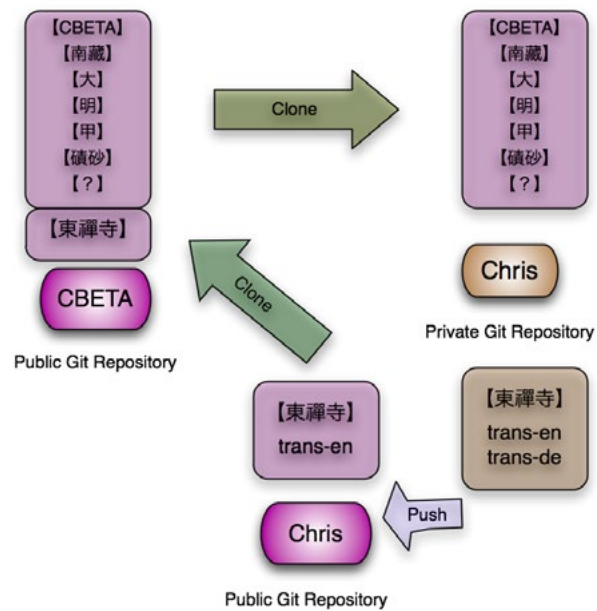**Figure 5: Adding new private branches in Chris repository**
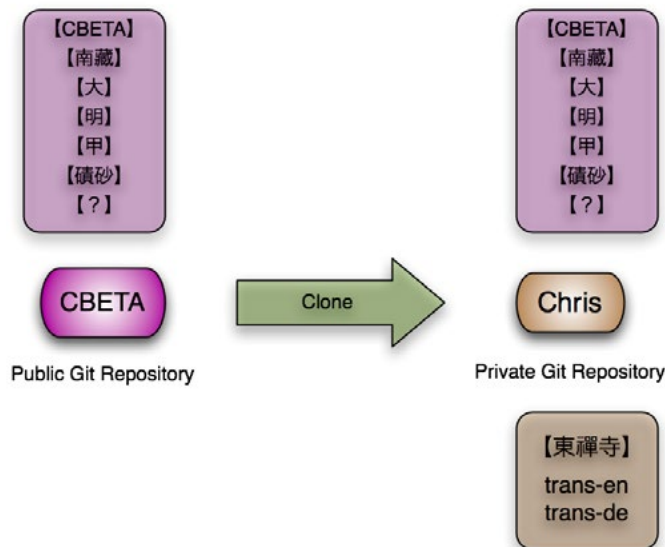
If the user Chris decides to publish from this repository the "trans-en" branch, that is the branch that holds the text with its English translation, and the new edition 【東禪寺, then Alice will be able to not only retrieve the public branches from the CBETA repository, but also the new 【東禪寺】 edition and the English translation, thus combining what is of interest to her in a new private repository on her own computer, as is shown in Figure 6.

Wittern, Christian. (2012). Towards an Architecture for Active Reading. *Scholarly and Research Communication, 3*(4): 040143, 11 pp.

7

Figure 6: Alice clones from the public repositories CBETA and Chris

Finally, CBETA might be interested in the new【東禪寺】edition added by Chris, and may add this to its own repository, which will make it more easily discoverable and might also indicate an endorsement of this new version.

Figure 7: CBETA adds branch【東禪寺】to repository

Wittern, Christian. (2012). Towards an Architecture for Active Reading. *Scholarly and Research Communication, 3*(4): 040143, 11 pp.

WRAPPING UP

The details of the format used for the editions and the application implemented as a prototype are given in *The Mandoku Project* (2012) but just to show how radically the term "annotation" is used here, Figure 8 shows a screenshot of the application with the text on the left hand side with punctuation and paragraphs added, a translation on the right hand column of the text, and additional notes between lines starting with :zhu: and :END:.

Figure 8: Text, punctuation, translation, and notes in the `chris-de` branch of the repository

With some simple conventions in using a publicly distributed version control repository, an infrastructure has been put into place that can serve as a backbone for the distribution of scholarly editions. While they might not have the same bells and whistles that users have become accustomed to in modern Web-based publications, they will require much less maintenance and will be serviceable long after the projects that initially created the texts has been discontinued.

There is however another precondition for this to work, which is not of a technical, but rather of a legal nature. In order for the above-described system to work, the cloning and copying mentioned above has to be possible not only technically, but also legally. That is, the texts need to be made available under licensing terms that allow, or maybe even encourage a re-use in the way indicated here, which seems to me to be the precondition of any scholarly discourse involving digital text.

Distributed ownership

As a researcher interested in a text, such as the 景德傳燈錄 *Jingde chuandeng lu* (CDL) used in the examples above, I would want to access all versions of the text, collate them, create annotations for certain parts of the text, link from the text to other, similar sections in other texts, and translate the text. In addition, I might want to mark names of places or persons in the texts and enhance them by looking up georeferences and biographical information. All of this activity will create new digital

Wittern, Christian. (2012). Towards an Architecture for Active Reading. *Scholarly and Research Communication, 3*(4): 040143, 11 pp.

9

objects that ideally I should be able to make available to other readers interested in the text. But since I do my work as part of my job as researcher, I cannot simply give up ownership of these additions, but need to be able to tag them with my name and I need to be able to control how these objects are used. This will also serve as an indicator of trustworthiness (or lack thereof), based on the reputation I might have or not have among the groups of users of these texts. Use of the DVCS as indicated above exactly supports this kind of publication model.

The so-called Web 2.0 enabled many people to contribute to websites and enhance their content, but only a few of these sites did acknowledge the contributor and made the contributions trackable (see O'Reilly, 2005 for more information about Web 2.0). In many cases, registering at a site as a user will make all information created by this user owned and controlled by the site owner, not the user herself. A new way of enabling better sharing and adding of content, that does not force the user to give up her rights, is needed.

It might also be worth pointing out here that the model made popular by Wikipedia, which relies on one central website where only one version of a certain fact can be stated, does not seem to be very well suited for scholarly discourse, where a mulititude of opinions and theories will prevail without an arbiter.

One more thing

It should also be obvious that the model sketched here is not the only possible development we might see. At the same time as the Web matures into a hive of interlocking and thriving communities of users, some parts of the Web are cordoned off as "walled gardens" that operate under a completely different mode. One example of such a walled garden is Apple's App Store (Wikipedia, 2012b), which maintains a strict division between sellers and consumers and tightly controls what developers are allowed to sell. This is a completely different mode of operation and a completely different set of infrastructure elements (although they are ultimately built on the same foundations as the Web itself). This underlines the fact that there is no inherent infrastructure in the digital world and it depends on our imagination and the momentum it creates to determine future development. But in both cases, it is the sequence of the decisions of the users that influence the direction this takes at every moment.

These are not merely arcane technical issues of little relevance to end users, but do in fact influence the directions the whole digital medium and the Internet will take. Free flow of information is essential for the development of scholarship and religious practice alike, and any attempt to limit this freedom should not be allowed to succeed.

Acknowledgements

Wittern, Christian. (2012). Towards an Architecture for Active Reading. *Scholarly and Research Communication, 3*(4): 040143, 11 pp.

Websites

The Awesome Highlighter: http://www.awesomehighlighter.com

The Open Annotation Collaboration: http://www.openannotation.org

References

Burnard, Lou & Syd Bauman. (Eds). (2008). *TEI P5: Guidelines for Electronic Text Encoding and Interchange*. Oxford, Providence, Charlottesville, Nancy: TEI Consortium.

Digital Papyrology. (2011). *IDP Data available on GitHub*. URL: http://digitalpapyrology.blogspot.com/2011/01/idp-data-available-on-github.html [January 22, 2012].

Gerber, Anna, Hyland, Andrew, & Hunter, Jane. (2010). *A Collaborative Scholarly Annotation System for Dynamic Web Documents – A Literary Case Study*. International Conference on Asian Digital Libraries (ICADL), Gold Coast, Australia, June 21–25 2010.

Loeliger, Jon. (2009). *Version Control with Git: Powerful Tools and Techniques for Collaborative Software Development*. Sebastopol, CA: O'Reilly Media.

O'Reilly, Tim. (2005). *What Is Web 2.0*. http://oreilly.com/lpt/a/6228 [June 22, 2012].

Reif, Adelbert. (1978). *Erich Fromm. Materialien zu seinem Werk*. Wien, Austria: Europaverlag.

Siemens, Ray, Leitch, Cara, Blake, Analisa, Armstrong, Karin, & Willinsky, John. (2009). It May Change My Understanding of the Field: Understanding Reading Tools for Scholars and Professional Readers. *Digital Humanities Quarterly 3.4*, http://www.digitalhumanities.org/dhq/vol/3/4/000075/000075.html [January 22, 2012].

Sima, Guang 司馬光. (1084). *Zizhi tongjian* 資治通鑑. Woodblock edition from the Wangli period in the holdings of the Institute for Research in Humanities.

The Mandoku Project. (2012). URL: http://www.mandoku.org/ [January 22, 2012].

Wikipedia. (2012a). *List of revision control software*. URL: http://en.wikipedia.org/wiki/List_of_revision_control_software [January 22, 2012].

Wikipedia. (2012b). *App Store*. URL: http://en.wikipedia.org/wiki/Appstore [January 22, 2012].

Wittern, Christian. (2012). Towards an Architecture for Active Reading. *Scholarly and Research Communication, 3*(4): 040143, 11 pp.

11