

Webbed Footnotes: Collaborative Annotation on the Web

Scott Andrew Golder

A.B. Linguistics

Harvard University

June 2003

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning,
In partial fulfillment of the requirements for the degree of

Master of Science in Media Arts and Sciences
at the
Massachusetts Institute of Technology

September 2005

© Massachusetts Institute of Technology, 2005.
All rights reserved.

Author

Scott Andrew Golder
Program in Media Arts and Sciences
August 5, 2005

Certified By

Judith S. Donath
Assistant Professor of Media Arts and Sciences
Thesis Supervisor

Accepted By

Andrew B. Lippman
Chairman, Departmental Committee on Graduate Studies
Program in Media Arts and Sciences

Webbed Footnotes: Collaborative Annotation on the Web

Scott Andrew Golder

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning,
In partial fulfillment of the requirements for the degree of
Master of Science in Media Arts and Sciences

Abstract

More and more, web users are moving from simply consuming content on the web to creating it as well, in the form of discussion boards, weblogs, wikis, and other collaborative and conversational media. Despite this, the web remains largely read-only; web designers create sites that are designed to be consumed by the public, much like other, traditional mass media. In this thesis, I explore free, shared annotation as a means of making the web more writable. In doing so, I hope to empower readers to engage more deeply with web content by actively participating in its production, and to have a voice on equal footing with those of the media producers whose content they consume.

This thesis details the design and evaluation of Webbed Footnotes, a system for publicly annotating web documents. Though it is not the first such system, its design is novel in that it is sensitive to the needs of both the would-be annotators and the owners of the websites being annotated. In particular, annotators would like their additions to be highly visible, yet website owners demand that their sites be presented in the manner they intended. Webbed Footnotes attempts to fulfill both of these conditions by making annotations highly visible, yet ensuring that the underlying documents remain legible. If Webbed Footnotes can partially solve the tension between annotators and authors, then public, shared annotation on the web may have a chance for widespread success, leading to savvier and more engaged readers.

Thesis Supervisor: Judith S. Donath
Title: Assistant Professor of Media Arts and Sciences

Webbed Footnotes: Collaborative Annotation on the Web

Scott Andrew Golder

Thesis Committee

Reader

Wanda J. Orlikowski
Eaton-Peabody Chair of Communication Sciences,
Professor of Information Technologies & Organization Studies
Massachusetts Institute of Technology

Reader

Walter Bender
Executive Director and Senior Research Scientist
MIT Media Laboratory
Massachusetts Institute of Technology

Acknowledgements

This thesis represents more than a single project; it represents the culmination of several years spent studying and creating at the Media Lab, during which time I have had the opportunity to meet and work with many brilliant and creative people, all of whom have influenced my work for the better.

The impact that my advisor Judith Donath has on my work is immeasurable. The Sociable Media Group has been my scholarly home since I was an undergraduate; in the five years since, Judith has overseen two of my theses and has guided my development as a researcher. I have enjoyed and benefited from working with and learning from her. I suspect – and hope – that our collaborations will continue well into the future.

My fellow students in the Sociable Media Group are not only colleagues; they are also friends. I have enjoyed learning and working with them: Karrie Karahalios, Fernanda Viégas, Andrew Fiore, Ethan Perry, Aaron Zinman and Francis Lam. In particular, I would like to thank Fernanda for being a great officemate and friend during the past several years.

I am grateful to Walter Bender for serving as a reader for this thesis. Though he is very busy as director of the Media Lab, he has been generous with his time and ideas, and I am very thankful for this.

I would like to thank Wanda Orlikowski for being a reader for this thesis, and also for allowing me to participate in her seminar last fall. I am grateful to have had the opportunity to work with and learn from her. Her guidance has been invaluable to me in my scholarly development.

Behind the scenes, a great deal of administrative effort is necessary for academic and research programs to operate smoothly. I thank Pat Solakoff and Linda Peterson in the MAS Office for their advice and patience. In the Sociable Media Group, I am appreciative of the hard work of each of the administrative staff who were with SMG during my time here: Lisa Lieberson, Colleen Capodilupo, Dorota Wojtas, Veronica Lupampa, Elizabeth Hartnett, Lisa Sachetta and Mary Heckbert.

Finally, I would like to thank my fiancée Angie Boyce and my parents Stuart and Caroline Golder. Angie has contributed not only love and moral support, but also invaluable editing and statistical analysis. And my parents' lifelong love, support and good examples have made everything in my life possible.

Table of Contents

Abstract.....	2
Acknowledgements	4
Table of Contents	5
Table of Figures	7
1 Introduction.....	8
1.1 Annotation Old and New.....	9
1.2 Annotation and Power.....	11
1.3 Graffiti as Geographic Annotation.....	12
1.4 Roadmap of this thesis.....	14
2 Conversation and Annotation.....	15
2.1 Online conversation	15
2.1.1 Usenet.....	17
2.1.2 Weblogs	20
2.1.3 Wikis and Collaborative Editing	22
2.2 Annotation on the Web.....	24
2.2.1 Yeahbutisitart.com: Online Graffiti.....	25
2.2.2 ThirdVoice : Power and Controversy	26
2.2.3 UTok: Moderating Annotation.....	28
2.3 Moderation	29
2.3.1 Leader-Based Moderation	29
2.3.2 Collaborative Moderation	30
3 Webbed Footnotes.....	33
3.1 Overview	33
3.1.1 Motivation	34
3.1.2 Supporting Annotation.....	34
3.1.3 Building A Conversation Space.....	37
3.1.4 Messages in Space	38
3.2 Approval-Based Moderation	39
3.2.1 Underprovisioning.....	40
3.2.2 Herding and Information Cascades.....	41
3.2.3 Lurking’s Effects in Approval Based Moderation	43
3.3 User Interface Elements.....	44
3.3.1 User Identity.....	44
3.3.2 The Annotation	45
3.3.3 Button Bar	46
3.4 Navigation and Discovery	47
3.4.1 The Sidebar.....	48
3.4.2 The User Page	48
3.4.3 The New York Times	48
3.5 Implementation.....	49
3.5.1 User Authentication	51
3.5.2 Loading, Reading and Writing Annotations	51

3.6	Limitations and Extensions.....	52
3.6.1	Browser Extensions and Trust.....	53
3.6.2	Ephemeral Documents and Annotation Placement.....	53
4	Evaluating Webbed Footnotes.....	56
4.1	Classroom Study.....	56
4.2	New York Times Discussion Study.....	57
4.2.1	Subject Activity.....	59
4.3	An Analysis of Annotating Behavior.....	60
4.3.1	Examples and Analysis.....	60
4.4	Public and Private Audiences.....	65
4.5	Conclusion.....	66
	References.....	69

Table of Figures

Figure 1-1. A poster advertising the Mohegan Sun casino, with a graffito. It reads, "I am a locked and loaded buying machine with no off switch." "Off switch" is crossed out, replaced with "conscience."	13
Figure 1-2. A graffito in Harvard Square: "Curt Schilling is a two faced prick + a traitor"	13
Figure 2-1. A blog entitled "Bitch. Ph.D."	21
Figure 2-2. A text clip from a Word document using the "track changes" feature; the black text is the first author, the red the second, and the blue the third.	22
Figure 2-3. The WikiWikiWeb page at C2.com, the original wiki.	23
Figure 2-4. The "edit-page" interface for the WikiWikiWeb page.	23
Figure 2-5. Harvard Law School's Annotation Engine.	25
Figure 2-6. The homepage of Graffiti the Web at yeahbutisart.com.	26
Figure 2-7. My homepage, vandalized by yeahbutisart.com's Graffiti the Web.	26
Figure 2-8. Screenshot of ThirdVoice.	26
Figure 2-9. A screenshot of a UTok-annotated MP3.com.	28
Figure 2-10. UrbanDictionary.com uses a thumbs up and down moderation system to rate the quality of entries.	31
Figure 2-11. Slashdot.org uses a 1-5 moderation system to rate the quality of posted messages.	31
Figure 3-1. The Webbed Footnotes system.	33
Figure 3-2. The username and color selection form.	44
Figure 3-3. An annotation marker. Because this is not the currently-displayed annotation, its marker is hollow.	45
Figure 3-4. Annotation marker for the annotation currently displayed. When the author views it, it has no icons.	45
Figure 3-5. Annotation marker for the annotation currently displayed. Three options are presented in the form of "User page", "Reply", and "Approve" icons.	45
Figure 3-6. The checkmark displayed when an annotation has been approved, and the thumbs up button selected for approval.	46
Figure 3-7. A floating box containing an annotation text. When an annotation marker is moused over, the floating annotation box is displayed at the bottom of the browser window.	46
Figure 3-8. The Webbed Footnotes button bar.	47
Figure 3-9. The sidebar, shown on the right hand side of each page and detailing the annotations on that page.	48
Figure 3-10. The user page, containing a list of all the annotations by a single user.	48
Figure 3-11. The silhouette icon.	48
Figure 3-12. The "note" icon.	49
Figure 3-13. The Firefox web browser.	49
Figure 3-14. Database tables and their contents.	50
Figure 3-15. The Webbed Footnotes system architecture. From left to right: the user, the web browser (containing the Webbed Footnotes extension), the hidden browser frame, several PHP scripts residing on the web server, and the SQL database.	51
Figure 3-16. Webbed Footnotes with a highlighted paragraph indicating potential placement of a new annotation.	55

1 Introduction

Webbed Footnotes is an annotation system for the worldwide web that allows readers to annotate any web document with their own comments. In this, Webbed Footnotes is not entirely new, since the notion of annotating the web has been considered since the web began.¹ Indeed, annotation of physical documents is a very old practice, perhaps as old as writing itself. This is all with good reason, since annotation – modifying existing documents by augmenting them with additional content – allows readers to add new information to a text, for themselves or for others, and for myriad reasons.

Webbed Footnotes adds two important dimensions to the practice of web annotation. First is a design that I believe more fairly balances the needs of the annotator with those of the annotated, and second is a novel moderation system I call Approval-Based Moderation. Annotation takes up precious, finite space within a document and draws some focus away from the original text. As I will demonstrate, annotation in physical as well as electronic documents often results in a tension between the original document and the new content. The author and reader rivalry is one for control of position and dominance of voice. An annotation system provides functionality that is designed to empower the annotator, but a system that tips the balance too far may meet with disapproval from authors. The ideal system will be balanced in its support of clear, legible documents as well as ample opportunity for annotation.

Because a document is most often complete once it reaches the consumption stage of being read and annotated, the agents of change are the readers. It is at this point that annotation has the potential to overwhelm the original document. Plentiful annotations can add much beneficial information to a document, enriching and adding value to it. However, annotations also have the potential to be uninteresting, wrong or deficient in some other way; annotations lacking any positive qualities not only fail to add anything of value, they also detract from the value of the original document, because they obscure its message. Especially when the bar to annotating is low, it is important to promote quality in contributions from would-be annotators. A moderation system, in which either a single moderator or a group of participants evaluate added content, is a tool that puts a kind of quality check on the participants' contributions in order to maintain a high rate of quality in the eyes of the user community. Unfortunately, many collaborative moderation systems, in which all users can moderate by a form of voting, are ineffective unless used by a sufficient number of participants. Approval-based Moderation, which I detail more completely in Chapter Three, is an attempt to make moderation robust even when few participants make the effort to moderate.

A document often has precious little space for annotation; in paper documents, this space is often in the form of a margin or front or end page. Because the opportunity to annotate a book is often limited physically (because it has a single owner) or through rules (in the case of a library), parceling out this space to annotators is unnecessary. A web document, however, may have many readers, so screen space becomes a shared resource, and a scarce one. Moderation helps prevent unnecessary consumption of that resource.

¹ Web pages could be privately annotated, with annotations being stored on the user's own machine, in NCSA Mosaic, one of the earliest web browsers [7].

This is just one of the many differences between paper and electronic documents. The latter hold many more opportunities for continual and collaborative use and modification, but their basis for existence is the paper document; much inspiration for electronic document design has been drawn from its paper predecessors. Even metaphors like “web page” demonstrate the electronic document’s roots in paper.

1.1 Annotation Old and New

The very first inspiration for this thesis was the medieval manuscript and other early books. While texts were clearly very different from electronic documents in many ways, they provide inspiration with their tradition of annotation and continual growth of knowledge within texts. These books were rare and valuable, containing rarefied knowledge that bespoke the privileged status of the owners. These owners were scholars, and their interactions with their documents were in the service of their learning. In the margins and between the lines of their books, such scholars would add translations, summarizations and explanations, and would instruct their students to do so in their lessons [31]. Marking up texts was part of the process of learning, and these markings served as tools for focusing the reader’s attention and remembering what had been read. These techniques persist to this day, as students are taught to mark up a text as a central part of studying it [62].

Documents can be annotated for many reasons. Today, as in times past, annotation is used as an aid in learning. Because annotating a document makes it easier to recall particular points of interest, annotating one’s documents with highlights, underlines and summarizations [62] can be useful when that information is needed quickly or frequently, perhaps when in a discussion or in preparation for writing with several sources [58].

Reading a text actively, which is what one must do when annotating it, makes the reader engage with the text more intensely and personally. Annotation supports learning precisely because of the deliberateness that must go into annotation. In Chapter Three I discuss in greater detail other ways annotation is useful for learning. At this point, it is important to simply recognize the way annotation can make readers more savvy about the texts they are consuming. This is valuable in itself, because a more thoughtful reader is one who can transcend simple information consumption and reach a point where they are not simply being fed information, but are using the information to become more informed.

As a text is marked up by its owner, it develops a unique character and grows in educational and, in the case of many medieval-era texts, monetary value. The copy, though it may be like all the others, develops uniqueness through the contributions of its owner. The owner’s annotations are useful not only to him or her, but to those to whom the text is eventually passed on. A text passed to another carries with it the added knowledge of all who have annotated it. Through this “cumulative scholarship” [31], each holder is guided by the content of the text as well as the contributions of all those through whose hands the text has passed. Through the process of annotating a book and passing it on, a single-owner, single-user text becomes multi-user and collaborative over time.

In an environment like the web, where each document is multi-user – or, at least, multi-reader – annotation is not generally possible. Web annotation could be very powerful because it has collaborative benefits that are not afforded by traditional annotation. The ability to share

annotations widely and to vastly increase the number of potential annotators presents new opportunities for rapid and continual “cumulative scholarship”, given the speed of access and quantity of potential contributors. Webbed Footnotes seeks to provide this capability to readers of the web. Ideally, readers’ contributions would improve the quality of the texts while allowing them to engage with one another. However, an open-access environment allows all readers the opportunity to contribute, and therefore there is a strong possibility that their contributions will not improve the quality of the text, and may even detract from it. Some of the existing environments I discuss in Chapter Two, like Usenet, demonstrate that contributions by some users are perceived by others to lack value. However, I suggest that it is better to address this problem separately – with, for example, collaborative moderation – as the egalitarian quality of this system is most important.

Widespread shared annotation also presents new tensions between author and annotator. When a document is annotated by a single person, there can be a tension between two opposing voices, that of the author and that of the annotator [31]. However, a single-user book limits the distribution of that tension to only the annotator himself. The author has the first word, but the reader has the last word. This may be satisfying to the reader, but it is not problematic to the author, who never knows of the private rebuttal, which is, for most practical purposes, shared with at most a few fellow readers. However, when an annotation is presented to a mass audience, it has the potential to become influential to many other readers of a document. The author’s powerful role as the sole voice heard is threatened, because the author must now share the floor with potentially as many other voices as there are readers.

Now may be a particularly apt time to give web users the ability to become annotators, though they may not have taken advantage of the opportunity in the past. A recent study by the Pew Internet and American Life project showed that 44% of internet users have now contributed some kind of content to an online environment. Many of these contributions have come in the form of a review of a product or service; those who contributed ratings have also been more skeptical of information online [39]. This demonstrates that there is a relationship between being a savvy reader and actively creating online content. This is a promising finding that bodes well for an annotation system like Webbed Footnotes.

Electronic document annotation is likely to be a very different kind of practice than paper document annotation. Though paper documents are a useful metaphor, it is important to use those metaphors productively, but to not be constrained by them. Sellen and Harper, in a reaction against the excitement in academia and industry about the possibility of a future of electronic-only documents, wrote *The Myth of the Paperless Office*, in which they demonstrated that the excitement should not be about digitization for its own sake, but rather because of the affordances electronic systems can provide. Simply emulating paper documents electronically will not be the answer. There are many useful aspects to paper that might be retained in digital document system design, but there are also many aspects of electronic documents to be taken advantage of, such as flexibility of navigation and presentation [58]. For example, in Webbed Footnotes, users can now be given the choice of when to see annotations, or whether to see them at all. The electronic annotation systems I discuss in the next chapter seem to be constrained in some ways by the idea of paper annotation, especially the metaphor of the “Post-It.” In Webbed Footnotes, by contrast, I hope to take advantage of one key difference between electronic and paper documents, namely their ability to change dynamically.

1.2 Annotation and Power

There can be a somewhat subversive element to annotation. Annotation opens the possibility for members of the general public to correct what they perceive as falsehoods or omissions in widely disseminated media. When the author is firmly entrenched in a seat of power, annotation can take on a political dimension, because the ability to annotate can allow even the least powerful actor to gain social power and be heard.

In traditional mass communication media, from print to radio to television, the voices heard and the opinions expressed tend to be those of powerful people, those who have the money necessary to disseminate information widely. To write, print and distribute a newspaper en masse each day requires large presses and raw materials, as well as a distribution infrastructure that can bring the content to the public. A television broadcast requires expensive transmitters and licenses from the government. All of this was supposed to change with the rise of the internet. As a distributed, decentralized network of computers, the internet gave rise to a utopian vision of replacing existing power structures by putting that power into the hands of the masses. The open source software community is one manifestation of this vision. Futurist Ted Nelson's *Computer Lib*, one of the earliest and most enduring texts on computing and culture, wrote in the introduction to the 1974 edition, "This book is for personal freedom, and against restriction and coercion" and "computers belong to all mankind," and with an almost revolutionary air reflective of the countercultural movement of the preceding decade, "Computer power to the people!" [43]. Nelson is representative of the belief that computers and technology in general would liberate and empower the masses against the establishment.

The mechanism by which this was supposed to happen was information. By making information cheap, it would be accessible to everyone. In his now-famous *Atlantic* essay from 1945 entitled "As We May Think," scientist and scholar Vannevar Bush predicted that information could be compressed and stored in a device he called the *memex* and, by virtue of being physically smaller, be cheaper. For example, he predicted the *Encyclopedia Britannica* might fit on a single sheet of microfilm and thus be distributed to all for only one cent. This is indeed a utopian vision, for several reasons, the least of which is technological. Storage and distribution of information comprises only a fraction of its cost; creating information can be extremely expensive. In order to be able to distribute information cheaply, it must also be produced cheaply as well. A full edition of *Encyclopedia Britannica* is expensive not because it consumes a great deal of paper, but because researching, writing and organizing information is time consuming and costly. In contrast, Wikipedia, a free, collaborative encyclopedia on the web is built using a tool called a wiki, or collaborative website to which any reader can freely contribute. Through the combined effort of thousands of volunteers, Wikipedia has become a very large container of freely available knowledge. There is animosity between *Brittanica* and Wikipedia; the latter says that the former represents the "old" way of doing things, is slow to change and represents the voice of the establishment. The former, in response, says that Wikipedia is a "faith-based encyclopedia" citing the lack of responsibility taken by the participants to ensure articles are accurate [40]. Certainly there are aspects of both arguments that make sense, but neither of the two serves all needs. On one hand, Wikipedia is important because it is free both in terms of cost and in terms of access to modify it. On the other hand, experts are experts precisely because they have rarefied knowledge that requires significant effort to attain; the *Brittanica* could not operate without being compensated for that expertise, and this sort of expertise is something Wikipedia cannot, and does not, make any claim to.

Nevertheless, Wikipedia is a highly successful, albeit imperfect, demonstration of the web's ability to decrease the cost of both the production and consumption of information, while making it an open, collaborative and collectively governed space.

Nelson's lifelong project and as yet unattained goal is the construction of what he calls Xanadu. Situated chronologically between Bush's memex and the Wikipedia, Nelson's Xanadu was also supposed to be a large information store. Largely inspired by the memex, Nelson envisioned information repositories that could be cross-linked to one another, where any document could refer to any other. Documents could be endlessly cut into pieces and reformed to create new documents, each piece referring back to its original source. This, Nelson reasoned, would lead to a resolution of disputes over copyright and would thus free all the world's information for public use. This was also the origin of the term "hypertext" [72].

Nelson's idea that documents should be able to be referenced from anywhere underlies a main premise of Webbed Footnotes, links are truly powerful only when anyone can create them. Being able to bring ideas from many sources together is key in the ability to defend a position in an argument. Therefore, in order to truly empower people to engage with texts and with one another, they must be able to connect those texts together for future readers. In creating Webbed Footnotes, for example, I envisioned annotations of news articles to contain references to alternate sources that present counterarguments. If readers can marshal evidence against a point within a document itself, it pressures the author of that document to be honest and transparent or, at least, empowers readers to warn others when an author is not. As I discussed above, annotation ideally will inspire more savvy readers; this is one way in which I hope it will do so.

Like any other tool, Webbed Footnotes bears the potential to be employed for more nefarious purposes. For example, annotations could be used for spam or vandalism, or could be used to harass and disrupt users across the web. Webbed Footnotes gives the power to annotate to all users, and that power may be used for good or for ill. Many of the problems that plague other online environments – spam, trolling, and so on – could persist in Webbed Footnotes as well. For example, many weblogs and online forums require users to "register" before posting in an effort to make them more accountable. Though it is impossible to control how user communities and the world at large might adopt such a tool, one might hope that, especially with a system of registration and moderation in place, that Webbed Footnotes' use would be overwhelmingly positive and that abuses could be handled by the community.

1.3 Graffiti as Geographic Annotation

In the physical world, an annotation is embodied through graffiti, a practice which further highlights the subversive side of annotation. Graffiti, often created with spray paint or heavy markers, can range from small "tags" in which the writer simply writes his name or nickname, to elaborate murals. In *Getting Up*, Castleman chronicles the lives of several graffiti artists in New York in the 1970's. Among the "artists" he meets, the monumental feat of painting a mural that covers an entire train car constitutes the highest achievement [12]. To these graffiti artists, what they were doing was indeed a form of public art, and they stressed aesthetics and style in their work. At the same time, there was a personal recognition aspect to their work; one way to achieve status was to have one's nickname seen in the most places around the city.

Their ultimate goal, it seemed, was to achieve some kind of recognition in a large, impersonal city, effected through the city's transit system.

Graffiti demonstrates the importance of place to annotation. Many graffiti, because they augment a specific item or location, can be meaningful only at that location. For example, many graffiti change the words of a sign in order to make it say something else. Often this consists merely of vulgarities, but sometimes a graffito makes a statement about the content of the item being augmented. In the example above, a poster advertising a casino's text is altered to suggest that there is something morally wrong with what the poster is advertising. In other cases, a graffito has meaning only in relation to the geographic and temporal locations in which it resides. The second example appeared in Harvard Square toward the end of 2004, after the Boston Red Sox had won the World Series and before the 2004 presidential election. Its text has some negative things to say about Red Sox star pitcher Curt Schilling, who actively campaigned for President Bush's reelection. This graffito has several ties to its geographic location and time period; its meaning and impact would be diminished if it were placed elsewhere. First, Bostonians were, by and large, proud of Schilling's role in winning the World Series. Secondly, he campaigned for a Republican candidate in a predominantly liberal city whose own Democratic senator was the opponent. Thus this graffito represents the city's mixed feelings toward Schilling. This graffito derives much of its meaning from and is tied strongly to Boston's political and social climate of late 2004. Therefore, it simply would not "work" the same way in another city or at another time.




Figure 1-1. A poster advertising the Mohegan Sun casino, with a graffiti. It reads, "I am a locked and loaded buying machine with no off switch." "Off switch" is crossed out, replaced with "conscience."



Figure 1-2. A graffiti in Harvard Square: "Curt Schilling is a two faced prick + a traitor"

Likewise, electronic annotations derive meaning from their locations. In many online environments that allow contributions from readers – message boards and discussion forums, product review sites, and so on – users' comments reside in their own space, separated from the content to which they refer. This breaks the relationship between the two and strips annotations of some of their meaning, making certain kinds of annotations – like the Schilling one above – impossible.

Graffiti add character and a unique touch to a space. The New York subway graffiti grew in popular culture to be emblematic of New York and its culture. When people walk through the streets of a city and see a graffiti like the Schilling one above, the feeling may be a sense of shared recognition, of shared knowledge and values. They inhabit a shared space, one in which

they are familiar and comfortable. In this way, I seek to turn an electronic document into a habitable space. By making a document annotatable, all of the document's readers see a segment of those who have come before and their reactions. The experience of the page is shared. This bond among readers has been examined before [e.g. 17,68], and Webbed Footnotes is designed to continue in that tradition.

1.4 Roadmap of this thesis

In this chapter I have discussed the importance of annotation in everyday documents, the need for giving readers a more active voice, and how these two things can go together. Annotation can be a way of bringing conversation into new spaces, making spaces more democratic and collaborative for readers.

In Chapter Two, I discuss the history and some of the theoretical issues of online conversation. Environments like Usenet, as well as tools like weblogs and wikis, each have features that have inspired aspects of Webbed Footnotes. Each supports conversation in a different way, and the positive as well as problematic aspects of each have affected the Webbed Footnotes design. As I discussed above, annotation is a kind of conversation. Accordingly, I will present in Chapter Two several examples of electronic annotation, as well as the positive and problematic aspects of each. I also discuss different kinds of moderation an electronic community might employ, including collaborative and leader-based styles, and why a community might choose one form of moderation over another.

In Chapter Three, I discuss Webbed Footnotes and my goal of treating an existing document as a conversation space and how I envision conversation growing within the space. Because moderation is so important in an online community, I have implemented a novel moderation system I call Approval-Based Moderation, designed to function properly even without many moderators, and especially suited to annotation systems. After these higher-level design concerns, I detail each aspect of the Webbed Footnotes user interface and its implementation. In the implementation section, I discuss the process of building Webbed Footnotes, how the client functions within the browser, how data is stored and transmitted, and other technical details related to the software.

Webbed Footnotes was tested on two groups of users, one large and one small. In Chapter Four, I present the results of those user tests. This includes statistical analyses of the participants' activities, qualitative analysis of the contents of their annotations, and selected feedback in which participants critiqued the system. Finally, I discuss the implications of Webbed Footnotes on online collaboration, and explore future directions.

2 Conversation and Annotation

A conversational environment based on annotation necessarily draws inspiration from existing environments, and learns from the successful and unsuccessful aspects of those environments' designs and use patterns. An asynchronous system, Webbed Footnotes draws from Usenet and online message boards. Because Webbed Footnotes' user-authored content is predicated on existing web documents, it therefore also draws from other systems that are similarly content-based, like weblogs and wikis. The novel elements of Webbed Footnotes, presented in chapter three, have implications for the future of environments like weblogs and wikis. These systems' current status, discussed in this chapter, will form the foundation for future directions and improvements, which are discussed in Chapter Four.

Document annotation and markup, especially on the web, is a well-trod space. Several commercial and noncommercial ventures have had mixed success, and several scholarly studies have analyzed their use. I discuss and critique several of these environments and the insights of other researchers, as well as my own.

An innovative aspect of Webbed Footnotes is the Approval-Based Moderation system. Discussed in Chapter Four, Approval-Based Moderation attempts to solve several of the dilemmas facing existing moderation systems. Several of these systems are discussed in this chapter, including their design philosophies and use, as well as scholarly analyses of them.

2.1 Online conversation

Online conversation is very different from face to face conversation. Participants are separated from one another physically and temporally, and so lack many pieces of information that are traditionally used in developing the social knowledge that forms the basis for social interaction. For example, others' perceived race, gender, socioeconomic and cultural status, provide clues about that person's identity and personality. Facial expressions and body language provide valuable information about another's immediate state of mind. Moreover, in a physical face-to-face encounter, another's presence is self-evident, the comments they make are unambiguously theirs, and the identity they can perform is somewhat constrained by these factors.

In stark contrast to this, people are physically separated and can have multiple identities online. They can often see others without being seen themselves, and can, to some degree, take on personas with characteristics very different from their own, and be a different age, race, gender, sexuality, and so on. This makes it possible not only to experiment and be free from some conditions of one's life, but it also frees people to do things that would incur social sanctions otherwise. This leads to two views of online life, one roughly utopian and one roughly dystopian. The utopian view holds that the separation of the body from online expression completely frees the individual from the constraints of sex, race, and so on and realizes a postmodernesque ideal of identity completely constituted through performance. In contrast, the dystopian view holds that, even in the online world, offline prejudices and power hierarchies persist and so the online world may be a hostile, male-dominated world. Citing the influence of feminist scholarship on early online communication, Kira Hall characterizes these as reflecting the divergence between theory-based and practice-based feminist work [27].

Utopian and dystopian views about the future of offline culture have existed for years, from dystopian novelists like Orwell and Huxley to real-world futurists like Negroponte [42] and Nelson [43,44], who predict utopian freedom of self-expression and information access through technology. The trouble with such predictions, whether by feminists, novelists or technologists, is their technologically deterministic approach, focusing as much on the outcomes (which are highly speculative) as on the processes and available tools (which are somewhat less so). One ongoing process is the separation of the individual through technology. The individual, unitary and corporeal, becomes according to Deleuze a “dividual,” a collection of separate pieces of information that constitute the whole [14]. While it is rather likely that this process of dividualization will continue, the state of society (online or offline) that will result is far less certain. That is, what norms and practices will develop around these new tools is unclear. Technologies and tools are adopted into environments – be they societies, communities, etc. – with preexisting values and practices; how these environments absorb or adapt to new practices will be highly dependent on how the new merges with the old. Dividualization, for example, provides a degree of freedom and flexibility; some may use that freedom to free themselves from an aspect of their identity they cannot escape offline – this is common with those in particularly stigmatized groups. Alternatively, people may use such flexibility to perform socially unacceptable actions and then dissociate themselves from those behaviors. How communities online react to the new possible choices among their constituents undoubtedly varies across communities.

Two of the most important practices in communication are the complementary processes of impression formation and impression management. Goffman [23] suggests that people tailor their actions in order to form desired impressions in the minds of onlookers. Individuals carefully tailor their social cues – those pieces of information that transmit social meaning – in order to make desired impressions; Goffman calls these cues *impressions given*. However, he notes that not all aspects of behavior are under one’s control. The process of impression management is limited by what one can manipulate (I cannot make myself taller) and is bounded by one’s own knowledge (I cannot know all possible interpretations of my actions). The impression actually made on the recipient, which Goffman calls the *impression given off*, is dependent on the recipient’s ability to interpret the information presented. This act of interpretation constitutes the process of impression formation.

Impression formation requires the ability to transmit and receive social cues. In everyday life, it is possible to perceive, however accurately or inaccurately, others’ age, gender, race, and personality factors. Many of these are visually perceptible, and are laden with cultural information. Style of dress, body language, and such, give clues about individuals’ tastes and cultural affiliations. Online, however, these cues are unavailable. People cannot see one another, so they cannot be sure of the identities of others, let alone any social information about them. This is problematic because social information about people strongly affects how others interpret their actions and statements. As a result, “miscommunication” and other misunderstandings are likely to occur when a lack of information about other participants leads to incorrect assumptions about them and incorrect interpretations of their actions. Jacobson’s study of perceptions of MOO participants and how those perceptions compared with reality demonstrated that even few cues can be used to fill out a detailed mental picture of what another person is like, and these mental pictures are based on prototypes, or general cognitive categories of, in this case, types of people [32]. These prototypes provide contextual

frameworks in which cues are interpreted; a cue's interpretation is highly dependent on that framework, and a mismatch can result in an incorrect interpretation.

Sproull and Kiesler [61] suggest that a lack of cues leads to several kinds of uninhibited behavior. They suggest that when many cues about others are available, behavior is "other-focused" and "controlled," and that when fewer cues are available, behavior is more "self-centered and unregulated." Among other things, Sproull and Kiesler observed that flouting of social convention is more likely when social cues are fewer. From these observations, they propose that mediated communication reduces the social cues transmitted to message recipients.

Walther [67] says that theories like this fall into a "cues filtered out" model, and notes that they have fallen into disfavor because they require there be a one-to-one mapping between cue and function, which is not necessarily true, and also because it takes a significant amount of time for meaningful communication to occur, an amount of time that was not provided when such early studies were conducted. Instead, Walther's theory of social information processing (SIP) suggests that when nonverbal and other channels of communication are unavailable (as in computer-mediated communication), communicators substitute other kinds of cues. He suggests that the critical difference between communication media is the rate at which they can transmit social information.

Because only a few cues are available, it is vital that the designer of an online environment display and visualize those cues in a way that does not mislead the participants by distorting the data in a way that misrepresents its meaning or exaggerates its significance. Graphical environments that seek to mimic offline environments often make the mistake of using "representational" graphics – an image of a sofa where online characters are intended to congregate, e.g. – which is actually more limiting than enabling, since it provides no actual purpose [16]. Likewise the garish carpeting of casinos is replicated in online cardrooms, where it is a distracting element for the players [25]. More importantly, individuals themselves must be represented in ways that do not ascribe to them behavior or traits they did not intend; rather, they ought to faithfully transmit only the information the individuals provide. Erickson [19] suggests that online environments "portray action, not interpretation." That is, interpretation of social information ought to be left up to the user.

Finally, it is vital to recognize that the mediated nature of online communication is emphatically not simply a hindrance to successful communication, but can rather be an enabling tool for richer communication. Most notably, the asynchronous nature of CMC gives participants the opportunity to reflect on the actions of others before acting themselves. Lags of time are common and expected; during such gaps, participants can marshal outside sources, do additional research, or more clearly formulate their thoughts before acting; an opportunity that face-to-face participants do not have – but often wish they did. It is my hope that in this thesis project, subjects will use asynchrony to their advantage in this way, bringing multiple texts together as they engage in discussion.

2.1.1 Usenet

One of the oldest online communities, Usenet is also one of the simplest in terms of its features and interface. It is certainly the most open. Because it is long-lived, well-established and

technologically open and therefore easy to study, it is the benchmark against which other online communities are often measured.

Begun in the 1970s at Duke University, Usenet is a distributed, asynchronous system designed for sharing messages across computer networks. It began with several research laboratories and universities, but grew widely as internet use grew throughout the 1990s [28]. As of 2002, over 14 million people were actively posting to Usenet [59]. Usenet is a commons; it is not owned by anyone in particular, but rather is a shared resource available to all who wish to use it. Aside from a subset of groups that appoint moderators, most of Usenet treats all people equally from a technology perspective. Any roles that people adopt generally arise through the everyday interaction with their fellow participants, until a routine of behaviors generally acceptable to everyone emerges [26].

Usenet is organized hierarchically into conversational spaces called newsgroups. This hierarchy is roughly topical; Usenet's top level hierarchy divides newsgroups into eight categories, including *sci* (science), *rec* (recreation) and *alt* (other topics). These categories are further subdivided; a group like *rec.gambling.blackjack* exists in the *rec* hierarchy in the *gambling* subhierarchy. This is a way of organizing groups. Though each newsgroup in Usenet uses the same technology and interface, each group develops its own community and social structure based on the actions of its participants and the roles they enact within the group [26]. Some groups may develop a norm of being welcoming to newcomers. Other groups may value having strong group boundaries and so might be less welcoming. This condition may manifest itself in the way new participants are treated or how open the participants are to divergent topics of conversation. That newsgroup communities each exhibit their own cultures demonstrates that the same communication technology, when given to different groups, may be used in very different ways, depending on the group's interests, goals, abilities, and other factors.

Conversations in Usenet newsgroups are threaded; this means that each message is linked to the message that it is a reply to (if any) and to any replies it itself has received, so that each message is read immediately after the message to which it is a reply. Threading allows a conversation to diverge and split, with simultaneous development of many conversations that all originated from a single message. A user may skip a particular branch in a message thread, yet continue to participate in the remainder of the conversation. Computationally, the threaded structure makes it easy to analyze a Usenet conversation in terms of individual contributions, revealing such things as who talks to whom, who is replied to most, and so on. This ease of computational analysis has made Usenet very popular among online discussion researchers.

Like many kinds of online communities to follow, Usenet has the interesting characteristic of being unable to measure its readership. It is difficult to count the number of people posting to Usenet, since many people may use multiple or pseudonymous email addresses. However, it is possible to count the number of email addresses, and make an estimate from there. Measuring Usenet's readership is a whole different matter, since no readership data is collected. The practice of reading an online community's content but not posting oneself is called "lurking." These lurkers form an "invisible audience" [26] that most Usenet participants take for granted as being present, but nobody really knows. Goffman [23] suggests that "barriers of perception" bound social interaction; lurkers exist outside these barriers, so a newsgroup may have few

lurkers or may have so many that they outnumber the posters. It is not possible to know for sure.

In other environments, some techniques for estimating readership can be applied. For example, in most Web-based environments like bulletin boards or weblogs, a webmaster may look at web server logs, which keeps track of which documents were retrieved by web clients, in order to count how many times a particular document was retrieved. In all these environments, web-based, Usenet, or otherwise, readership has no impact on the community. In contrast, Webbed Footnotes readers have an impact on the environment solely through reading. Because reading content affects how it is shown to future visitors, readership logging is crucial. Webbed Footnotes is a closed system, and other closed systems that require users to maintain an account and log in may of course keep more accurate readership accounting.

Nonnecke and Preece have undertaken several studies on lurking in email lists, including lurkers' demographics [46] and their reasons for lurking [47]. One reason they cite is inhibition about being in a public space. Posting in Usenet is akin to jumping into the middle of the dance floor and demanding attention; there is no way to applaud from the sidelines. This is a design choice; Usenet's choice reflects the simplicity of its structure. Each post in Usenet is treated the same way by the Usenet protocol and most client interfaces. Each message is transferred separately and with its own headers, takes up a single line in the message list, fills the entire screen when it is being read, and so on. This prevents more complex kinds of participation. For example, anyone who wishes to make their opinion known must do so by authoring a whole post. There is no way to comment on someone's post, except through another post. Though this basic structure has served Usenet well, subsequent online environments have embellished upon the kinds of messages that users can author. The fact that Usenet users can publicly participate only by posting an entire message may contribute to some people's inhibition to do so. Perhaps if a more lightweight, less obtrusive method of public participation were possible, such inhibition could be decreased.

Takahashi, et al. [63] suggest that posting messages is not the only kind of activity a lurker can engage in, and underscore this fact through their preference for the term "active lurking." They observe that lurkers can make use of the information they glean from lurking in their own offline activities, or disseminate that information to others who are not part of the online community.

The debate rages on whether lurking constitutes free-riding, or taking advantage of a communal resource without contributing to it oneself. Kollock and Smith [35] argue that reading an online community's posts without posting oneself constitutes free-riding on the community's knowledge, but Nonnecke and Preece [47] suggest that lurking behavior is a valuable part of an online community, because distribution costs incurred through reading are not a limited resource in electronic groups. However, Kollock and Smith also note that bandwidth, or the information capacity of a channel, is a scarce resource and that overuse of bandwidth is another "social dilemma" affecting online environments [35]. Lurking can be

characterized as a behavior that, in the aggregate, prevents the overuse of bandwidth by not clogging the channel with too many messages.²

The issue of lurking extends well past Usenet into other online forums. For authors of weblogs, the conversation environment I discuss next, this is a personal, potentially ego bruising or inflating issue. The personal nature of a weblog means that the size of the audience represents how many people are paying attention specifically to the weblog author. A weblog author may often wonder, then, how many people are reading their writing, and who those people are. Rather than being a point of largely academic interest to a community, to a weblog author the makeup of an audience is personally deeply meaningful, despite being mostly unknowable.

2.1.2 Weblogs

A weblog or blog is a frequently-changing website owned or run by an individual or group, to which messages are posted in chronological order by those running it.

Blogs are designed to be topical. The date and time of each post is prominently displayed along with the post's title. Though they are asynchronous, they are intended to be experienced with some degree of immediacy. A blog must generally be updated frequently in order to hold the attention of readers.

A blog's technological underpinnings consist of a collection of web scripts that accomplishes the task of making it easy to chronologically add posts to the website. As such, this tool can be used for a great variety of purposes, depending on the blogger's interests. Personal blogs may recount what the author has done that day; politically-oriented blogs may discuss and opine on events currently in the news. When blogging, as the act of keeping a weblog is called, first received media attention, it was largely for its personal aspect [e.g. 49]. But during a contentious presidential election in 2004, political weblogs received a great deal of attention from large, well-established journalistic organizations. Though these two uses receive the most media attention, blogging tools can be used for any topic imaginable, and have even been used in classroom settings.

Weblogs have been compared to diaries and journals [49], soapboxes and op-ed-like journalistic media. The weblog is emphatically not a commons; it has a well-defined owner (or owners), the expression of whose viewpoints is the *raison d'être* for the site. Many weblogs allow comments from readers, but in each case the author's (owner's) post sits at the top of the page most prominently, and any comments from readers are listed below. This document layout imposes a clear hierarchy. The owner is more important than any visitor; it is *the owner's* personal space, and visitors are guests.

² Users seem to be aware of this problem. When posting messages in a shared channel, especially when the message's topic is obscure, users will sometimes apologize for their use of bandwidth and occasionally compare their message to spam.

Because a blog is so individual-centered, a blog post is a kind of solo performance by the author, in a way that a Usenet post, for example, is not. Visitors come to see the author “perform” through their posts, and where other communities rise and fall on the collective input of their participants, a blog does so based on the ability of its author to entertain, inform and otherwise captivate an audience. For this reason, some top bloggers, and many other bloggers in their own niches, become somewhat well-known and celebrity-like. Subtitled “least likely to pacify social conservatives,” the weblog known as “Bitch. Ph.D.” is authored by an anonymous professor who writes on gender issues and academia, and the overlap between them. She has developed a readership – a fan base – that appreciates her writing so much, her child, who she nicknamed “Pseudonymous kid”, receives gifts from readers as an expression of their appreciation.

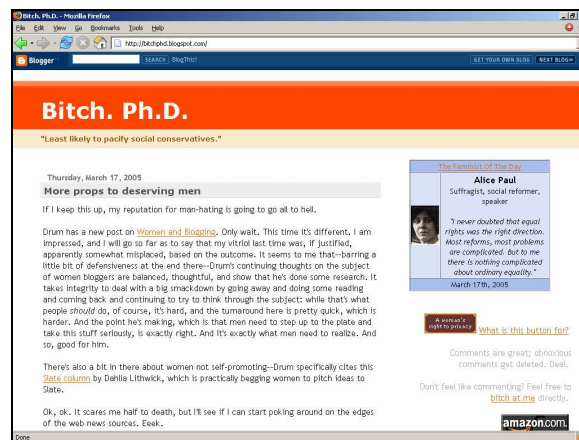


Figure 2-1. A blog entitled “Bitch. Ph.D.”

Because blogs are, at their core, simply web pages, their authors can make the most of the power of the web by linking to other web documents when talking about them. Linking to interesting websites or to news articles is a way of bringing together disparate pieces of information that, together, can be reorganized in order to argue a point that none of the individual pieces necessarily argues on its own. The poststructuralist semiotics concept of intertextuality explores the relationship between a text and other texts that are contemporary or antecedent to it. This theory suggests that each text shares ideas and language, whether directly and indirectly, with existing texts [33]. By bringing texts together, bloggers explicitly bring new interpretations to texts, more readily expose their flaws and prevent alternatives. Expanding on this ideal is the goal of Webbed Footnotes.

Many blogs also link to one another. It is common for blogs to contain a “blogroll,” or a list of other blogs that the author reads frequently. It is a way of signaling affiliation, but it also serves to direct readers to other blogs they may also enjoy. Additionally, blogs link frequently to other blogs’ posts while adding their own commentary or response, as a way of turning the network of blogs into a large, distributed conversation. With regard to Webbed Footnotes, this practice is especially interesting, because commentary is happening about and around documents, rather than within them. Webbed Footnotes seeks to make this third possibility a reality.

When reading news on the web, for example, it is not possible to see what, if any, blog commentary exists about it. One project that has attempted to remedy this problem is The Annotated Times, by Blogrunner.³ It searches the web for blog conversations linking to New York Times articles and reprints those posts in list, along with an excerpt from the article. It solves very well the problem of aggregating content that probably ought to be together from the start, but it does not present the content in a novel way; it looks like any other blog or message board webpage, with a topic (in this case, the article excerpt) at the top followed by a list of comments presented serially beneath it. Webbed Footnotes attempts to bring together

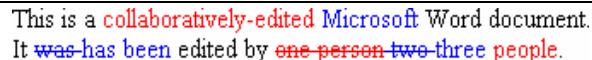
³ <http://annotatedtimes.blogrunner.com/>

documents and commentary in such a way that they are brought together by design, embedding each comment within the article or document. Whereas in the network of weblogs documents are treated as artifacts to be linked to and talked about and around, Webbed Footnotes treats the document as a conversation space, so that all conversation happens within the context of the document. Webbed Footnotes could certainly not be a replacement for weblogs, because their goals differ; the weblog, as I discussed above, is inherently personal, whereas Webbed Footnotes is article- or document-centric. What they share is a focus on discussion through exploiting and realizing the inherently intertextual possibilities of documents.

2.1.3 Wikis and Collaborative Editing

A focus of much computer-supported cooperative work (CSCW) research is the collaborative editing of documents, either simultaneously or asynchronously. In this work, the goal is often to allow a group to create a document that has a single, unified voice that represents the contributions, ideas and opinions of each contributor.

The ubiquitous word processing software Microsoft Word has a feature called “Track Changes” that allows successive editors to make document changes that are reflected visually on the document. New text is shown in a color assigned to the author.



This is a collaboratively-edited Microsoft Word document.
It was-has been edited by one-person two-three people.

Figure 2-2. A text clip from a Word document using the “track changes” feature; the black text is the first author, the red the second, and the blue the third.

This maintains the distinction among authors and their contributions, which is valuable during the editing process. When creating a document together, this distinction must be maintained so that authors can coordinate what each writes and can discuss and negotiate the content, as well as contextualize their further edits based on their knowledge of their collaborators’ individual perspectives. However, when “track changes” is turned off and changes are accepted (or rejected), all the text turns black and the author color is removed, removing the distinction among authors. Eliding differences among authors helps project a single voice, thereby expressing agreement among those authors and maintaining their credibility as a unified group. This is especially important in documents where maintaining a public face is necessary, such as a corporate press release or an academic research paper. A document that expresses a single voice carries with it a suggestion of completeness and finality, since the single voice implies agreement among all the authors with any disagreements or equivocalities among them resolved.

On the web, an increasingly popular tool for creating a collaboratively edited document is a wiki. Wiki software creates individual web pages, the content of which any visitor may edit. Visitors may add, change or delete any text, images, links, and other content on the page through a simple edit interface; once saved, their changes are reflected immediately on the page, and are visitable (and editable by) subsequent visitors.

Wikis were invented by Ward Cunningham. The front page from his first wiki, *WikiWikiWeb*, is shown below. The idea behind the wiki is egalitarian in the extreme; every collaborator is considered equal and the group of collaborators is often referred to as a community.

Responsibility for building consensus and being thoughtful and nice are stressed⁴. It is certainly possible, in an environment where all content is freely manipulable, for vandalism and destruction to run rampant. It is indeed the case that wikis are often vandalized. However, just as anyone is free to vandalize a wiki, anyone is free to undo the vandalism; since many wikis keep version history, it is often possible to do this with a single click. When a wiki has many users acting in good faith, there are many people who can undo the damage; therefore the burden of “policing” the wiki is distributed across the community rather than falling on one person’s shoulders.

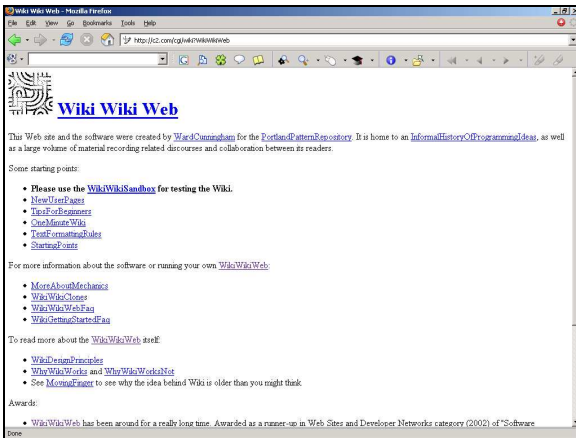


Figure 2-3. The WikiWikiWeb page at C2.com, the original wiki.

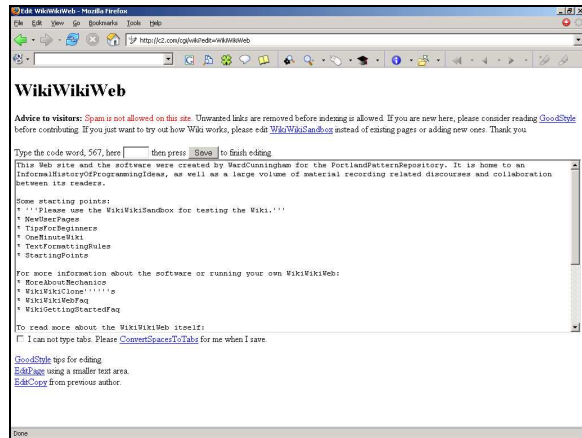


Figure 2-4. The “edit-page” interface for the WikiWikiWeb page.

One of the largest, most popular and most well-known wikis is Wikipedia, a wiki encyclopedia which contains almost 600,000 collaboratively authored articles in the English version, as well as smaller versions in many other languages.⁵ Because wiki software makes it easy for users to update the content of a document, wikis are useful as continually updated and refined information repositories, as the original wiki was for the Portland Pattern Repository, an association of software developers. As documents are updated, obsolete content is removed or replaced. An encyclopedia can be thought of as an information repository on a massive scale, and therefore using a wiki to create an encyclopedia is fitting.

The collaborative and egalitarian nature of wikis, however, detracts from Wikipedia’s credibility for many people. Since the authors are unknown, it is unclear where the information comes from and whether it – and, by extension, all of Wikipedia – can be trusted as a reference source. This is problematic because, for an encyclopedia, establishing trust that true, accurate and unbiased information is being presented is of prime importance. Critics contrast Wikipedia’s openness with the monolithic and authoritative nature of other encyclopedias such as Britannica, where articles are written by identified authors who are widely recognized as experts in their respective domains, and fact-checkers verify the accuracy of each article. Wikipedia advocates respond by suggesting that because so many people look at articles and can edit them, incorrect information will be caught and corrected.

⁴ These terms are used on the Why Wiki Works page. <http://c2.com/cgi/wiki?WhyWikiWorks>

⁵ Wikipedia’s list of languages in which Wikipedias are in existence:
http://meta.wikimedia.org/wiki/Complete_list_of_language_Wikipedias_available

It is clear that many errors are caught, inaccuracies corrected and prose edited to flow more smoothly. Viégas et al.'s [65] study of Wikipedia document history demonstrates that though several aberrant behaviors take place, they are often very rapidly corrected. For example, they observed that "mass deletions" often occur whereby a vandalizer will delete all the text on a page, perhaps replacing the text with a line or two of vulgarities. However, such actions were repaired by other participants in only a few minutes. When a contribution is highly contentious, it might result in an "edit war," in which two or more participants go back and forth on whether a particular contribution should be in an article – one makes a change and another reverts to the previous version, the former makes the change again and the latter again reverts it, and so on. Eventually one party gets tired and gives up, or else a compromise is reached [65]. Though this neighborhood watch-like or community policing-like system clearly catches many transgressors, it is nearly certain that many more subtle errors persist, either because nobody has the knowledge to notice them, or nobody has taken the time to seek them out.

In a wiki, and especially in Wikipedia, documents are never complete. As an information repository, a wiki can never contain all the information in the world, and in no case could everyone agree on an optimal order. In the case of Wikipedia, as time marches on, new findings add to society's collective knowledge. Therefore, wiki pages may be in endless flux, always open to further addition, reinterpretation and revision.

The idea in Webbed Footnotes that every document ought to be forever augmentable comes from the wiki, and shares its belief that each reader should be able to contribute to collective knowledge. Their missions, however, are very different. Webbed Footnotes' goal is conversation through existing web spaces, whereas a wiki is a space unto itself and is not for conversation. Whereas a wiki seeks a collaborative yet uniform voice, Webbed Footnotes seeks to distinguish individual authors and their contributions, which is much more in keeping with a conversational environment. Furthermore, wikis are "walled gardens" – private spaces unto themselves in which users can write freely. Webbed Footnotes seeks to make the web – a preexisting space – into one in which readers can participate more actively.

2.2 Annotation on the Web

Annotation is a kind of conversation. As users annotate documents with their thoughts and ideas, they interact with one another, and an annotation environment begins to resemble other kinds of online conversational environments, albeit against a backdrop of preexisting content. Many of the characteristics of other online conversational environments apply to annotation systems as well, as users seek outlets for self-expression. For example, conversations may grow hostile as conflicting opinions clash, and environments with less power to sanction participants may encourage some to behave inappropriately. Indeed, the very purpose of most annotation systems has been to foster conversation among their users.

Webbed Footnotes is certainly not the first tool for annotating web documents through annotation. On the contrary, annotation is a well-trod space that has been explored by researchers and exploited by several commercial ventures. In this section I discuss some of this research and describe several of these systems, identifying aspects of them that were successful and other aspects that may be changed or improved upon. From each of these systems, some

lessons can be learned; lessons that were taken into consideration when designing Webbed Footnotes.

Bouvin [7] situates web annotation within a broader category of hypermedia augmentation, which he describes as including tools that integrate with web browsers or servers and add content not otherwise part of the web document. Other categories of hypermedia augmentation include tools for supporting discussion (of which annotation is part), link creation and structure creation. Bouvin cites the existence of several annotation systems, including ComMentor and CritLink, the former of which was designed to allow communication about documents in digital libraries [57] and the latter of which was [74] designed for structured conversation. Another tool in this category is Harvard Law School's Annotation Engine⁶, which adds references to annotations in-line, rather than by superimposing them like almost all other systems, including Webbed Footnotes. The annotations themselves appear in a sidebar. The document itself, however, appears rather cluttered, with references to the large number of annotations appearing within it.

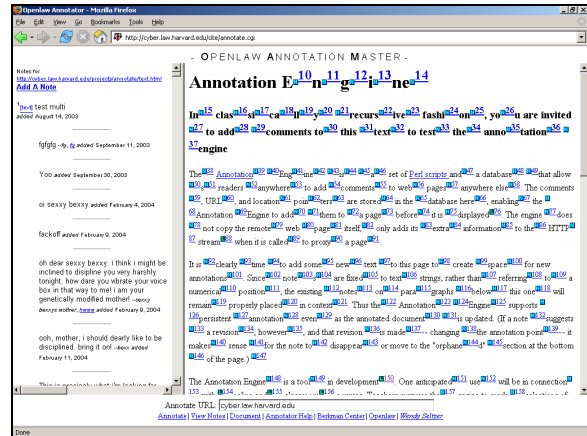


Figure 2-5. Harvard Law School's Annotation Engine.

In considering the design and use of any “hypermedia augmentation” system, it is important to balance the importance of the augmentation against that of the original document being modified. The user is visiting a web site for its own content, and the augmentation tool is supposed to do just that; augment – not supplant. As I discuss in this section, this principle can be violated, leading to controversial results.

2.2.1 Yeahbutisart.com: Online Graffiti

Yeahbutisart.com's Graffiti the Web is a very constrained web annotation tool. With it, visitors can superimpose content onto any target page. As the name implies, the metaphor of traditional urban graffiti is very strong here. However, this metaphor is implemented only very crudely. As Castleman [12] observed, graffiti artists were just that; though simple “tags” were common, the ideal was a large mural-like painting. In contrast, Graffiti the Web has an extremely limited “vocabulary,” consisting of a handful of pre-selected images – many of which are obscene enough that they could not be included in the screenshot below, as well as a single line of bold, red text or images of the user's choosing. This constrained vocabulary severely limits the self-expression of its users, and is representative of only the least imaginative and least prestigious kind of graffiti. Furthermore, the pre-selected images suggest that Graffiti The Web is promoting, through its design, simple defacement rather than commentary.

⁶ The Annotation Engine, from Harvard Law School's Berkman Center for Internet and Society, can be found online at <http://cyber.law.harvard.edu/projects/annotate.html>

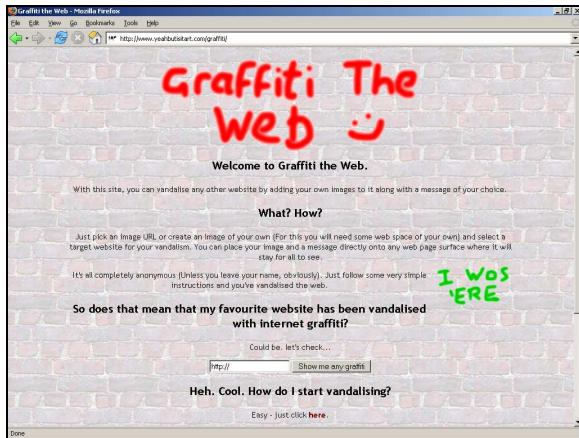


Figure 2-6. The homepage of Graffiti the Web at yeahbutisitart.com.

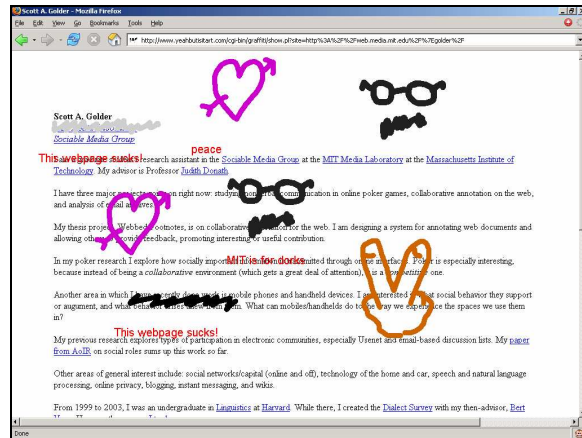


Figure 2-7. My homepage, vandalized by yeahbutisitart.com's Graffiti the Web.

In Graffiti the Web, all contributions are anonymous and, once placed, last forever. As a result, over time the most popular pages become completely obscured such that they are no longer even recognizable, let alone readable. Its “most vandalized site,” msn.com, is so covered in graffiti that, not only is the page’s original content completely illegible, most of the graffiti are as well. Graffiti the Web cannot be called a failure, however, since it has precisely achieved its goal of ruining target pages. Though its goal is antithetical to that of Webbed Footnotes, it is inspirational in that it demonstrates users’ interest in having some authorial power across the whole of the web, and is an outlet for some clearly pent-up desire to self-express.

2.2.2 ThirdVoice : Power and Controversy

The most commercially successful annotation system, ThirdVoice also encountered a great deal of controversy, controversy which was, ironically, most likely due to its very commercial success.

Third Voice was an annotation system that integrated with the web browser, allowing users to place Post-It-like annotations onto web pages. Third Voice was explicitly not intended to be used for defacement, but rather appealed to the same ideal of user empowerment as I do in Webbed Footnotes, desiring a “civic-mindedness that would keep corporations, government and the media honest” [37]. Despite their lofty ideals, Third Voice faced critics who believed that the software was being used to digitally deface the web pages to which annotations were attached. Third Voice countered, arguing that because their annotations resided on their own server, they were not altering the original document in any way and could not be considered defacement or destruction⁸. Some critics banded together and formed an organization called Just Say No To Third Voice.

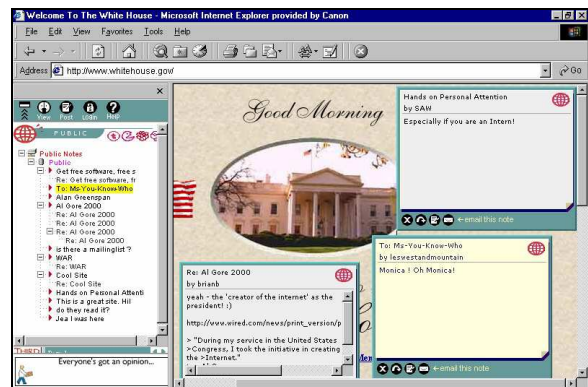


Figure 2-8. Screenshot of ThirdVoice⁷

⁷ Screenshot is courtesy Jeremy Bowers, <http://www.jerf.org/>

⁸ <http://web.archive.org/web/20000815054518/http://www.saynotothirdvoice.com/>

It remains an open question, legally as well as theoretically, what offline behavior is most analogous to annotating an electronic document. In one sense, annotating another's website in a critical way is a form of protest; it may therefore be most apt to compare this action to other protected acts of protest speech, like picketing or leafleting. Though critical annotations can be damaging to business, so can a visible public demonstration in front of a storefront or office building. Others might suggest that critical annotations are more akin to smashing a storefront's windows, suggesting that annotations are more significantly damaging and possibly are – or should be – illegal and a form of vandalism. A perspective that views annotations as more benign may compare the annotation to the restaurant review. A restaurant review does not mandate being read, but rather is available to interested readers who seek it out; likewise, annotation software is an adjunct to the web, and those who are not interested in reading annotations are in no way compelled to do so, nor are they even made aware of their existence.

Third Voice's assertion is true, that no actual changes to the original site or server are effected. Annotation is completely different from hacking a web server, which involves trespassing by forging access credentials like passwords. In contrast, annotation requires no more access to a web server than any other reader, since documents are simply retrieved, and any additional behavior takes place with a third party.

Several factors come into play in determining whether an annotation tool is more likely to face controversy or find acceptance. First, the scale of the system is important. A system that has a small number of users will have a smaller amount of annotation, and will also affect a small number of people, since annotation systems are technologically opt-in. The smaller a system's user base, the more the system could be seen as private. Likewise, the larger a system's user base, the more it could feel public and even ubiquitous. A system that is larger is also more powerful, and may be seen as posing more of a threat to those who might be upset by having their websites annotated.

Another factor that may affect whether an annotation system becomes an object of controversy is whether it is commercial. Many of the systems discussed above, including the HLS Annotation Engine, CritLink and even Webbed Footnotes, are research projects undertaken in academic environments. They are not seeking to make a profit nor push a corporate agenda, which may make their motives seem more pure and their actions more honest. In contrast, Third Voice and others like it have a distinct profit motive. They seek to grow a user base and profit from it, most often by advertising. It is understandable that a large website operator, who is either trying to sell its own products or else profit itself through its own advertising, would be upset at another entity trying to profit at their expense by advertising in their space.

The first two factors I discussed were social and economic, respectively; the third is design-oriented. Most annotation systems give the annotation a great deal of weight visually. Annotations are often large, solid-colored blocks that circumscribe the annotation's content. This makes the annotation very prominent on the page, but obscures a significant proportion of the underlying webpage. This is good if you are running the annotation system, but not so good if you are running the website the annotations are augmenting. Furthermore, even if you were an annotation system participant, if you simply wanted to read a web document and not any annotations on top of it, you may find it challenging to do so.

I suggest that the fundamental design flaw with existing annotation systems is that, by virtue of the heavy weight given to annotations, these systems foster not the conversation they seek, but rather interruption. The design elements of Webbed Footnotes have been specifically chosen in order to address this issue. By taking up a large amount of screen space, annotation systems stifle the underlying website's content. I suggest that an annotation system that ensures the accessibility and legibility of annotated websites would meet with less criticism than one that does not. In Webbed Footnotes, one of my prime goals was to create a system in which even a heavily annotated document remains legible.

Third Voice, and the conflicts it encountered, serve as examples, for the designers of future annotation systems. Any change in the balance of power in the control of mass communication channels is bound to be met with opposition. This opposition is bound to be expected and must be overcome in the interest of empowering users. However, this interest must be tempered with respect for the goals and needs of those websites being annotated.

2.2.3 UTok: Moderating Annotation

Unlike Third Voice, uTok, another commercial annotation system, did not integrate with the browser. Instead, an external client "listened" to the browser [75]. This had the benefit of saving some of the screen space otherwise lost to the third party application. However, uTok suffered from the same dilemma as Third Voice, whereby annotations obscured the documents beneath them. As discussed above, this can be problematic when wanting to read the document as well as its annotations.

uTok has the distinction among annotation systems of employing moderation. The system it used was a combination of up-and-down and numerical moderation. Users could vote "thumbs up" or "thumbs down" for a given annotation; on the server, however, these votes would aggregate into a 1-5 score, which would be displayed on the annotation as the appropriate number of "stars". When an annotation's score dropped too low, it would be deleted. UTok also rated individual users; the individual's rating is an aggregate of the ratings given to the annotations they authored [75]. The personal aspect of this rating system seems like an incentive to post messages that are appealing to others, because the quality of post reflects directly upon the author. One drawback to this and other moderation systems, in which a message must reach a particular threshold to be removed, is that such a threshold is often not reached, resulting in a large amount of not-quite-bad-enough content persisting longer than the user community might want it to.

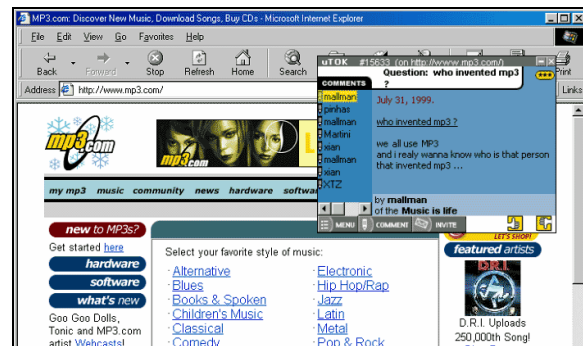


Figure 2-9. A screenshot of a UTok-annotated MP3.com.⁹

⁹ This screenshot comes from <http://web.media.mit.edu/~orit/utok.html> and is used with the permission of Orit Zuckerman.

Because the web is a very large place and dwarfs the number of annotations, the density of annotations is rather low. With billions of pages on the web, it would be unlikely that users of uTok or any other annotation system would randomly come across an annotated page, except on the most heavily trafficked sites. For users, this isn't a problem, since most users go to heavily trafficked sites by definition. However, for system designers, it is a problem because much of the user-contributed content might never be seen because the pages to which it is attached are so rarely visited.

Two aspects of uTok addressed the problem of annotation density. First was a simple notification mechanism that sent alerts to authors when their messages were replied to. Second, was the notion of interest groups that users could subscribe to and then receive alerts about messages assigned to that group. In general, seeing annotations on a page would be constrained by that page's often-limited traffic. Groups are an aggregation mechanism that help bring together the diffuse content of the annotation system, making it more likely that all the content added to it will be experienced by its users.

2.3 Moderation

When many users are contributing their thoughts, ideas and comments to a common space, not all user-contributed content will be highly regarded by others. Peoples' interests, knowledge and intelligence vary, and so too do the qualities of their contributions. Even at their best, people sometimes fail to be insightful, interesting, polite or well-informed. In short, much user generated content is unappealing to the majority of other users. For this reason, many online environments employ moderation systems that in some way judge the quality of users' contributions. The effect is that participants are shielded from lesser quality contributions. Not only does this system improve the participant experience, Resnick et al. [56] suggest that it creates a disincentive to contribute low-quality contributions, since few or no participants will see it.

Moderation systems can be divided into a taxonomy based on who performs the moderation, and when it is performed. Because these systems are related to structures of social control, they can also be compared to forms of government.

2.3.1 Leader-Based Moderation

Some moderation techniques involve having a clearly-established, technologically-enforced role of moderator, someone who has the responsibility of moderating content for the entire group. These people may be system administrators, in closed systems like web-based message boards, or volunteer participants in an open system like Usenet. Usenet-style moderation is akin to a republic, where democratically chosen leaders make decisions in the best interests of their constituents.

Moderation by leaders is often a community protection tool [26], employed when an unmoderated forum has problems with disruptive or unwanted participants; such is often the case in Usenet. Because moderators are trusted to moderate according to the will of the group members, they must necessarily have a good understanding of what their fellow members' judgments about what to block or not block are. They must also be seen as trustworthy, fair and knowledgeable by the rest of the group; these traits are essential in being a good moderator.

In many cases, moderators are existing high-status members of the community, but there is a reflexive relationship between their status and their role. Their high esteem helps them carry out the moderator role, and the moderator role positively affects the high regard in which they are held [26].

An important distinction among leader-based moderation systems is that moderation can be performed either before or after messages are made available to all the users; this has important implications for how transparent the moderation is, and therefore how closely the user community can monitor the actions of its moderators. In the case of Usenet moderation, messages are approved of or denied by the moderator before the message is sent to the group. Group members never see the messages that are denied. In contrast, in many web-based message boards, moderation takes place usually only after a message is posted. This is probably a healthy choice; because web-based environments' moderators are usually system administrators with potentially autocratic power over the user community, if the member participants can see messages as well as the moderator's actions in a transparent way, then this transparency gives participants an opportunity to have their voices heard in an otherwise potentially repressive regime. A system administrator seeking to socially engineer a community in a web-based message board could remove messages, but it is likely that the removed messages would have been seen by some members, who could "blow the whistle" on a moderator who was believed to be moderating inappropriately.

Whether moderations take place before or after posting in a leader-based moderation system is likely an effect of the openness of participation within that system; systems with registration may prefer and indeed seem more likely to employ moderation after posting, and systems without moderation before posting. In Usenet, an open system where no user registration is required, moderation takes place before messages are posted and, in web-based message board with a registration requirement, it often takes place afterwards. There is potential for abuse in letting anyone post to an open system, but a registration system mitigates some of that risk; a system that ties participants' actions to an established identity decreases the likelihood for abuse [15]. Further, having well-defined group boundaries – as those provided by a user registration system – contribute positively to the ability to self-govern. These ideas are foundational to the development of reputation systems [35].

Weblogs provide an interesting hybrid. Though virtually all weblogs employ leader (i.e. blogger) moderation only, some require registration in order to comment, some do not require any moderation, and some require posts from non-registered users to be held for moderation. Many mailing lists also work in this way. What is taking place here is that those who are trusted within the community are seen as posing a lesser risk, but those who are unknown carry more risk, and so are subject to early moderation. Indeed, many moderated mailing lists function this way as well; the popular mailman¹⁰ list management software gives list administrators this option.

2.3.2 Collaborative Moderation

If leader-based moderation is like representative democracy, collaborative moderation can be compared to direct democracy, where each message is a referendum unto itself. Moderation is

¹⁰ Mailman is the GNU Mailing List Manager. <http://www.gnu.org/software/mailman/>

performed by all participants, whose votes are combined to decide the fate of a message. Generally, collaboratively moderated systems tend to employ one of two methods. One method has members vote on or rate others' messages using an up-down system and a message's score is the percentage of up votes. Instead of up-down, some systems employ yes-no or plus-minus; these vary only superficially, but have the same underlying effect. Another method uses a numerical system where members rate posts from 1-5 (integers only) and a message's score is the average of all these. GroupLens [56] and Slashdot [38] work this way, and other variations on this theme abound.

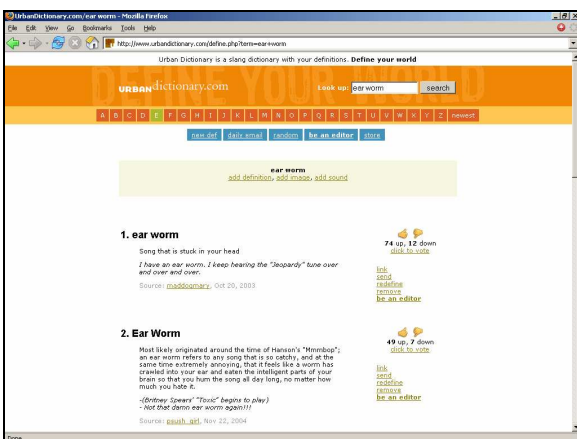


Figure 2-10. UrbanDictionary.com uses a thumbs up and down moderation system to rate the quality of entries.

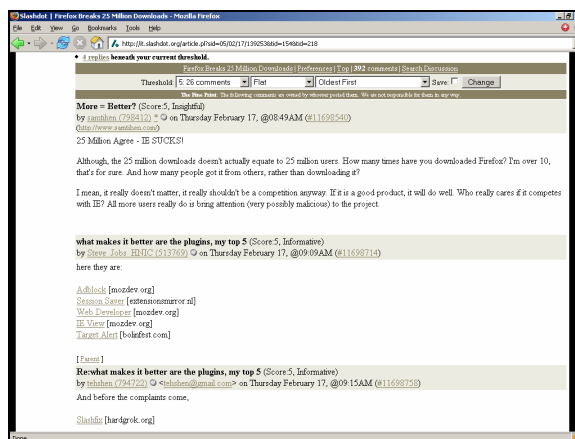


Figure 2-11. Slashdot.org uses a 1-5 moderation system to rate the quality of posted messages.

Because collaborative moderation considers the votes of many participants, collaborative moderation need not, and often does not, result in a definitive yes/no answer. Instead, when users' votes are combined to form a rating for a particular message, this rating affects the way in which the message is shown to others. For example, users may set a threshold beneath which low-rated messages are not shown, or highest-scoring messages are shown first, with subsequent messages decreasing in score. The general goal is that subsequent readers be somewhat shielded from those messages previous readers did not think highly of.

A democracy requires the participation of its constituents in order to function in a way that represents their wishes. Unfortunately, low voter turnout continues to be a problem in democratic countries like the United States.¹¹ There are many reasons for not voting, including apathy, belief one's vote does not "count," or the lack of desire to devote the time necessary to do it. Likewise, collaborative moderation systems suffer from "underprovisioning," or the lack of moderation data on with to rate messages, because an insufficient number of readers engage in moderation [38]. Their reasons for not moderating may be similar to those non-voters give for not voting, or may include some of the reasons lurkers give for reading but not posting: there are too many messages and the moderation task is burdensome, they have too little time to moderate and would rather simply read, or they lack confidence in their own expertise or judgment [46].

¹¹ The voter participation rate does not appear to be declining, according to political scientists Samuel Popkin and Michael McDonald; for the past seven Presidential elections, they report that voter turnout has varied between 52 and 60 percent. Nevertheless, they suggest that the country would benefit from higher rates of participation [53].

The underprovisioning caused by insufficient moderation presents a significant problem. When a message is rated with fewer moderation datapoints, the resulting rating is less likely to represent the opinions of the entire group, since the sample size is relatively small. As a result, moderation is either not possible or less effective until a sufficient amount of moderation takes place. This is a problem in environments like Slashdot, where only 28% of messages get one or more moderation votes [38].

One of the reasons for underprovisioning is that users can free ride on the ratings of other participants; there is a lack of incentive to provide ratings oneself, when one could as easily wait until others have done a reasonable amount of rating and benefit from their effort. The incentive to moderate includes the ability to have an impact on the content and practices of the community and, to some degree, to have a voice in determining what the community deems valuable, interesting, prestigious, and so on. For some participants, this incentive is not enough to make them want to invest time and effort in moderating.

An interesting theoretical question is, if lurking is not free-riding on the knowledge of the community, then why can lurking without moderating be considered free-riding? I suggest that while lurking provides a benefit (in the form of an audience that does not use the scarce resource of bandwidth), not moderating provides no benefit. Since moderation information is generally presented only in the aggregate, more moderation makes that information more valuable and does not take up a scarce resource like bandwidth. Furthermore, being a reader who posts messages is a fundamentally different role from being one who does not. It is reasonable for different role repertoires to entail different resource utilization. However, the reader role does not fundamentally change based on whether one moderates as well.¹²

Underprovisioning is only a problem when sufficient provisioning of moderation data is necessary for the system to function efficiently. If the moderation system can function robustly when only those most interested moderate, then not moderating does not constitute free-riding, because such an individual would not be depriving others of a more efficient system.

Webbed Footnotes employs a system I call Approval-Based Moderation, which is based on collaborative moderation but attempts to solve the problem of underprovisioning by functioning properly even when there is a lack of moderation data. It is discussed in the following chapter.

¹² That is, provided the moderations are not visible to others. In an environment where quantitative (1-5, +/-, etc) moderations are performed, these ratings are generally not shown to others individually, and even then not accompanied by the identities of the individuals who performed those moderations. This entire analysis presupposes that an individual's moderation choices are private. In contrast, if a user's individual moderation choices were made visible to other users, then moderation would involve a significantly higher amount of personal exposure and would be much more similar to posting.

3 Webbed Footnotes

In this chapter, I describe the design and implementation of Webbed Footnotes. After a high-level discussion of the design of Webbed Footnotes and its motivation, I present Webbed Footnotes' novel moderation system, which I call Approval-Based Moderation. Following this, I detail the design of Webbed Footnotes' user interface and its technical implementation.

3.1 Overview

Webbed Footnotes is an annotation tool that integrates with the web browser to allow users to write public annotations and affix them to web pages. Users' annotations are presented alongside the original text so that subsequent visitors may read the annotations while reading the document itself. Webbed Footnotes is designed to be conversational, and allows users to explicitly reply to one another's annotations, the text of all these interactions also being affixed atop the document. In this way, Webbed Footnotes is designed to enrich the reading experience by turning the read-only document into a conversation space.

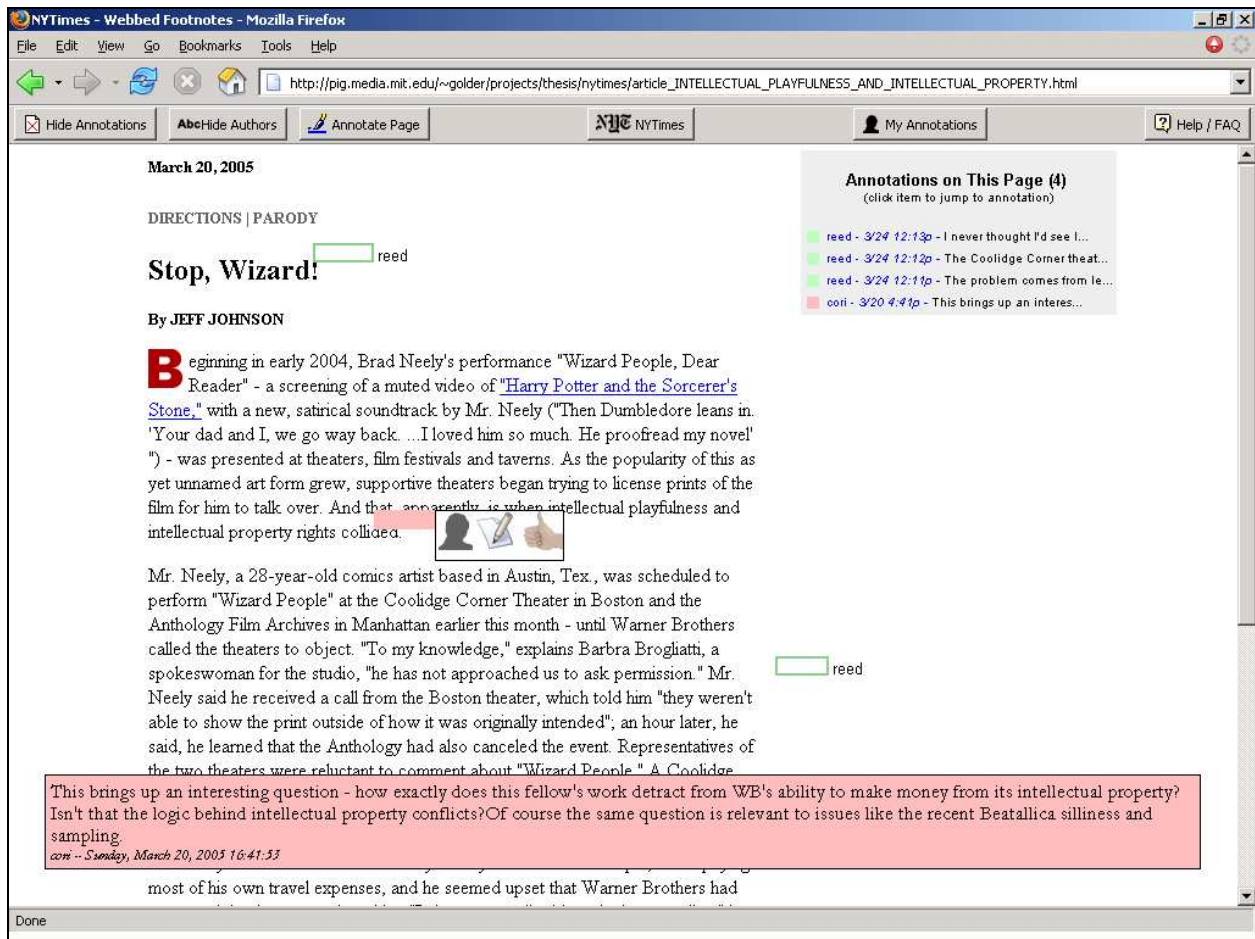


Figure 3-1. The Webbed Footnotes system.

3.1.1 Motivation

Webbed Footnotes is designed to allow web users to interact with a text rather than simply read it, and thereby interact with other readers *through* that text. Recall in Chapter One the discussion of medieval manuscripts and later books, and the commentaries left by readers within. At their best, such commentaries could provide clarifications of the text, correct errors within it, or build upon the knowledge contained within it. Albeit with some difficulty, the physical text – an essentially single-user tool – could, over time, become a multi-user, collaborative conversational environment.

The worldwide web is a multi-user collaborative reading environment that supports conversation in some environments. The distinction is subtle, but most documents on the web and other web pages are for consumption only – they are designed meticulously, with copy and other content carefully chosen by a clearly-defined author in order to deliver a well-crafted impression. There is a single voice presenting a single message. There is no room in most web environments for additional voices.

Where other voices are allowed, in weblogs, message boards, and other environments described in the previous chapter, these voices are set apart physically from their referents, divorced of the contexts in which they would otherwise occur. Their geographic isolation subordinates them to the main content of the document; most designers likely do this unthinkingly, but it nevertheless has the effect of using geography as one way of distinguishing the voices of the visitors from those of the creators.

Third Voice, discussed in the previous chapter at greater length, took a step in the direction of putting visitors' voices on the same level as those of the webpage creators. It understandably led to controversy over usurping power, but I suggest that the greater problem was that of legibility. An annotated page in Third Voice's software left the underlying document unintelligible, as it was completely covered up. Third Voice and its kin do not inspire conversation with a text; they encourage *interruption*.

Webbed Footnotes remedies this by allowing even a heavily annotated document to remain legible. Presenting annotations' text as footnotes rather than completely superimposing them over the document significantly decreases the degree to which the original document is obscured. Ideally, this will make the existence of users' commentary clear, thus empowering the reader-annotator, while preserving the dominance of message and design that website creators seek.

3.1.2 Supporting Annotation

In determining how to support annotation, it is important to identify what annotation is, as well as where it is used and why. I consider annotation to be the meaningful marking of a text. This can be writing, like marginal comments, or other kinds of marking like underlining, circling, highlighting, and so on. Such markings must, however, be meaningfully related to the text; this

excludes, for example, errant pen markings or the use of a book as scribbling paper for younger children, a rather common practice [31].¹³

Others have defined annotation in a similarly straightforward way, for example, as “marking made on a document at a particular place” [9]. In principle I agree with this definition, but take for granted that a marking is made at a particular place; no marking on a physical document can ever be placeless. Going one step further, I assume that the place is intrinsically part of the annotation’s meaning; that is, a marking on a page can have place but no meaning, but never meaning without place. In Chapter One the examples of graffiti demonstrated the importance of location and timing in creating the meaning in an annotation. Later, I discuss further the importance of the document as geographic conversation space.

People annotate documents for myriad reasons. Many of these reasons are purely personal. Readers record their own ideas in order to learn and remember the contents of the text in leisure reading [31] and also in work-related reading; Sellen and Harper suggest that, in work life, analysis, comprehension and review are prime reasons individuals mark up their documents [58]. More specifically, annotation behavior helps all readers see and understand document structure, as well as important facts within the text [62].

Documents, especially annotated documents, are useful in the completion of work tasks where workers need to bring together disparate sources of information; having many documents accessible at once, annotation allows workers to find the content they need at a glance [58]. Documents are often also used in support of discussion [58]. When in a meeting, for example, participants might refer to the contents of their personal copies of shared documents as evidence while making arguments. Having previously annotated a document, participants may be able to more quickly go back to those points they believe are salient, and therefore be more effective in argumentation. If annotation allows readers to make faster, more flexible use of their documents, either alone or in groups, then the power of annotation is that it leads to savvier and more effective users of information.

When people annotate for the purposes of sharing the annotated document, the relationship between reader/annotator and meaning changes. Using a non-textual marker for annotation, such as a highlight or underline or even asterisk, a reader marks off some text that is of interest to the reader. The additional information that is crucial – that which remains in the annotator’s head – is *why* that text is of interest. Is it particularly insightful? Is it unclear? Does it pique the annotator’s curiosity? Did the annotator mean for the annotation to be interesting to others, or only to himself? Subsequent readers can make only best guesses. In some ways this can be useful, if users seek to capitalize on the ambiguity of meaning for playful purposes [29], but can have more serious consequences if it results in miscommunication between annotator and subsequent reader. I suggest that the interpretation of any annotation is always and inherently up for debate, and even more so for non-textual annotations.

¹³ There are, of course, many cases that are not clearly on one side of “meaningfully related” or the other. For example, it isn’t obvious how to classify the use of a book margin to write a phone number or other opportunistic message. Anecdotally it has been reported to me that some people write such notes in the book so that the book serves also as a temporal record of what the reader was doing while reading that book. This “diary-like” use of books is intensely personal and, though it may not be related to the content of the book, it is undoubtedly related to the use of the book, or the book as an artifact, and cannot be disregarded, but nonetheless is decidedly different from more traditional marginalia.

When made available to other users, annotations on paper or electronic documents can serve as bookmarks or signposts that alert the reader that something interesting lies at that location. Wexelblat's [68] work demonstrates that the bookmarks that users record in a shared electronic space like the web can be immensely useful to subsequent visitors who are alerted to the presence of something interesting.

Bookmarking, structuring and commenting on paper documents are all very high-level goals for annotation. In order to carry them out, several lower level, technical tasks are performed. These tasks include writing prose, highlighting, underlining or circling text, and manipulating documents into logical groups or piles. Sellen and Harper observe that paper affords being marked by a variety of media (pencil, ink, crayon, and so on) and being stacked, stapled and otherwise arranged [58]. There is a great deal of flexibility in what can be done to paper, and there is also a great deal of flexibility in what users seek to do with it, as well.

When thinking about the design for Webbed Footnotes, the goal was emphatically not to reproduce digitally that which can be done to paper. After all, electronic and paper documents have very different affordances, or properties that determine how a thing could possibly be used [48]. For example, though they do not afford being scribbled upon, electronic documents do afford being easily marked with regularity and precision, something that is not true of paper [58]. They also have different audiences and vastly different existing usage patterns. They afford being used by many people at once across great distances, and are written for different purposes in mind. Paper annotation serves merely as inspirational metaphor.

It is important to note here that I consider a document prepared electronically for printing on paper to be a paper document, not an electronic one, because its electronic state is one of creation, rather than its use. The kinds of electronic documents I am discussing are those that are meant to be used as well as created electronically, especially web pages. Web pages that have already gone through the process of being authored are complete, or are at least complete to the degree that they are presentable and ready to be consumed by users. The annotations I care about are those that come from users who are not the authors.

Webbed Footnotes, being designed primarily for interpersonal communication through documents, seeks to afford annotation that supports structure, but more importantly, supports discussion; Webbed Footnotes' annotations, therefore, are designed to be mostly textual. Each annotation consists of two parts, the annotation marker and the annotation text. The marker is a colored rectangle that identifies the annotation as being present. The annotation text is shown toward the bottom of the screen only when the annotation marker is selected. The annotation marker serves as the visual cue that another user has augmented this space, providing an *at-a-glance* notification of annotation, and therefore potential of structure and interest. The annotation marker, however, is not designed to stand on its own; the salience of the text box presented to the user when the user chooses to annotate a page nearly begs for content, such that the annotation marker with no accompanying text may seem out of place. In practice, the empty marker – the conspicuous absence of content – could conceivably take on meaning of its own among a group of users and, while it would be a use not designed for, it would nonetheless qualify as a meaningful augmentation of the document.

Webbed Footnotes' actual use could present other problematic conflicts with the definition of annotation discussed above. If an annotation is not referring to the text – spam, for example – then is it an annotation? The software infrastructure would say yes, and it would be presented like any legitimate annotation. Other users' characterizations may vary – indeed, it is often difficult for some users to tell trolls from legitimate posters in other kinds of online forums [26]. From a theoretical perspective, we can say the page is being augmented, but only content can be annotated and, in such a case, the content is irrelevant to the augmentation, so such an augmentation could not be considered an annotation of the text.

Though users' goals are flexible and therefore their tools must be as well, Webbed Footnotes does not support free-form page augmentation in the form of purely non-textual markings. This means there is not support for underlining, highlighting, or drawing. This is done in part to keep separate and distinct the voices and contributions of the annotators from those of the documents' authors. A uniformly-sized, colored rectangle that displays text is unambiguously a user's contribution, but highlighted, underlined text or other irregular content may not be easily distinguishable from that provided by the author, potentially leading to greater author upset.

3.1.3 Building A Conversation Space

As users add annotations to a document and other users respond to those annotations, users may engage with one another's words and ideas just as much or more than they do with those of the document's original text. Though the unannotated document works well standing on its own, once heavily annotated, it can become more of a backdrop to conversation than a topic of conversation itself.

Another metaphor that applies here rather aptly is the coral reef. The tiny organisms whose skeletal structures form large coral reefs do so by building off of one another's skeletal structures in order to form the whole [41]. In some cases, these reefs are bootstrapped by the sinking of large, non-organic objects like oil rigs, which provide the foundations on which the coral begins to grow [64]. Once the coral organisms begin to grow on one another, they no longer rely on the oil rig structure which gave rise to it, but at the same time the two are so intertwined, they are inseparable. Though a coral reef can grow without the oil rig, of course, its presence stimulates that growth. In much the same way, in Webbed Footnotes, the document is a bootstrap for a self-sustaining conversation that may grow to no longer require the document, yet it is inseparable from that document, which provides the clues to its origins and growth.

Since Webbed Footnotes' conversations go away after some time through its moderation system, a document that was once highly populated with conversation may over time become devoid of conversation as the annotations on it are read and expired. However, what remains is the original document; this document is still full of its original content, and is just as likely as before to bootstrap a new conversation.¹⁴ Contrast this with a message board or newsgroup which, in its original state is empty, lacking any conversation. There is no content there to

¹⁴ Of course, temporal issues may make the document less relevant. A news article, for example, may be less likely to give rise to a conversation a year after being published than it was on its day of publication.

bootstrap a new conversation, and revitalizing a “dead” message board or newsgroup may be significantly more difficult than revitalizing conversation around a document.

Unlike message boards and newsgroups, a document is not necessarily designed for conversation. A person may read a document, think about it, and not annotate it for discussion at all. An unannotated document does not look the least bit out of place. On the other hand, an empty message board or newsgroup is deserted, that is without use, and it calls out for conversation; the unannotated document does no such thing.

Conversation in physical space also often happens in places that are not designed for it. Opportunistic conversations take place on street corners and other inhospitable locations [67] all the time, yet an empty street corner does not look out of place. In contrast, a nightclub with few, quiet patrons is disconcerting. The nightclub, like the empty newsgroup puts pressure on its (non-)participants to be self-consciously active, whereas the unannotated document, like the street corner, supports opportunistic interpersonal interaction, but does not demand it.

The pressure to perform in a demanding online space, whether populated or not, adds to the self-consciousness of its would-be participants, their discomfort “in public” preventing them from active participation [46]. The all-or-nothing nature of participation in an online environment, where one is either an active poster or else a reader with no effect on the group, significantly affects users’ motivations and decision-making regarding participation. In the following section, I discuss what the effects are of giving users alternative means of participation.

3.1.4 Messages in Space

In this conversation space, the stuff of conversation is the messages themselves. The geography of the annotation is often part of the message itself; a message that refers to a paragraph, sentence, or even word, must be near its referent if it is going to be coherent to readers. Because geography is so important, annotation placement is fixed once the author chooses it. Moreover, once completed, an annotation may not be edited, even by its author. This design choice has some definite tradeoffs; if a user misplaces an annotation or makes a typographical error, the inability to correct such a mistake may be frustrating. However, if a user were to seek to make larger changes, such as altering an expressed opinion in light of subsequent conversation, ensuring that the history of the conversation remains intact may be more beneficial, with respect to the user group as a whole. When encountering an online environment for the first time, it is beneficial to see the recent past experiences of those already in that environment [68, 30]. This is especially true in a conversational environment, where important cues are learned from those already present [26].

A conversation is an organic, continually developing exchange of ideas, and may include formal argumentation, disagreement and the changing or refining of opinions by its participants. Unlike the construction of a document, such as a Wikipedia article or a Word document, or a deliberative process in which agreement and compromise are sought, a conversation need not have content that is widely accepted or seen as timeless. Rather, much of the value of a conversation comes from the process of conversing and from the understanding that many different voices with sometimes competing interests are interacting. Much of the interesting information in a conversational space consists of trends in participants’ activity levels and in

changes in topic [26]; this is especially true when visualizing conversations [16]. Even unstructured conversations exhibit a coherent structure [29]; several attempts have been made to identify these structures computationally [60], as well as to encourage users to add that structural metadata themselves [73]. Especially in a discussion about timely events, each contribution to a conversation is part of the historical record, and Webbed Footnotes views this as important.

With messages in fixed locations across the system, messages are inevitably geographically diffuse. Additional non-spatial methods of reaching users' annotations proved necessary in order to increase the likelihood users would encounter one another's contributions. The two non-spatial aggregators in Webbed Footnotes, the Sidebar and the User Page, aggregate users' annotations by news article and by author, respectively. While each of these tools brings together annotations from across many locations, the annotations cannot be fully appreciated separated from their context. Therefore, these two tools provide links to each annotation's geographic location so that the annotation may be understood in its intended context.

Additionally, the fixity of the location of the message presents an interesting constraint on the moderation employed in the space. In many online environments, each message is moderated, either by other participants, or by an appointed moderator. In many of these environments, the presentation of messages changes based on that moderation; often higher-rated messages appear first. For all the reasons discussed here, it would be impractical to separate annotations from their locations in order to represent high ratings. Therefore, it was necessary to find another way to use moderation data while leaving annotations in place; the method developed for Webbed Footnotes is called Approval-Based Moderation.

3.2 Approval-Based Moderation

Webbed Footnotes is the first system to employ what I term Approval-Based Moderation, a collaborative moderation method in which content requires proactive approval in order to remain in the system. Unlike other collaborative moderation systems, discussed in the previous chapter, Approval-Based Moderation functions just as well when no users moderate as it does when many users do.

In Approval-Based Moderation, each annotation (or comment, post, message, etc.) has a predefined lifespan. This lifespan may be defined in a variety of ways, including absolute time (e.g. several days or one week) or number of "views", either overall or by unique visitors. As it is viewed (or as time passes), its remaining lifespan slowly decreases until it is expired, at which point it ceases to be displayed. For example, a message that has a lifespan ten views "long" will, after being shown to ten unique users, expire.

When a user "approves of" an annotation, its lifespan is increased. In Webbed Footnotes, for example, an annotation beginning life with a ten view lifespan, as all annotations do, will, after being approved of by a single user, have an effective lifespan of fifteen views. If each approval adds several more unique views to the life of an annotation, then one that is highly-approved of will be shown to more people than those that are not highly approved of. As a result, users are more likely to see messages that are highly approved than ones that are not, just like other collaborative moderation systems.

In Webbed Footnotes, Approval-Based Moderation is realized through the use of a “thumbs up” button on each message. Thumbs up and down have been used in several other systems, including ones discussed in Chapter Three, like uTok and Urban Dictionary. Approval-Based Moderation is designed to be less burdensome on users than other kinds of moderation. First, it is a much simpler choice – approve or not – than other systems, where one must choose among up, down or no approval at all, or worse, a selection along a numerical scale. Secondly, it is far more optional; since users know the system functions effectively without approval, then approval may seem like less a chore and more a choice than other kinds of moderation.

The choice to give annotations an initial lifespan of ten unique user views and incrementing lifespans by five views for each approval in Webbed Footnotes was a purely arbitrary choice. These numbers were chosen based on the expected size of the user community. Consider a community of ten people; it would not make sense to have expiration occur after ten people, because that would effectively include the entire population. Likewise, in a user community of one hundred thousand, expiration after ten unique viewers is probably too small. It is perhaps sufficient to propose that these numbers be definable by communities themselves, or their designated leaders, such that the numbers allow for a reasonable sample of the population to see a message before its fate is sealed.

Resnick [55] suggested that disappearing moderations might not be the best approach and suggested as an alternative that a probability be used that determines whether an annotation be shown or not. A message that is never approved of would have a very low probability and thus be very rarely shown but, as Resnick points out, it would always be possible to “rescue” an annotation [55]. This might work well in some scenarios, but it is likely that even a very high quality comment might eventually cease to be shown because it ceases to be timely. In such a case, a message probably ought to disappear because the temporal context in which it was written and meant to be appreciated has passed. However, there is significant merit in the idea that annotations can be unjustly disposed of if a mix of users who are not right for that message happens to see it early on.

3.2.1 Underprovisioning

One crucial difference between Approval-Based Moderation and traditional methods is that, here, underprovisioning is not a problem. As discussed in Chapter Three, in a traditional collaborative moderation system, in order to rank messages, a significant quantity of moderations must have been performed. However, many systems lack such a quantity of moderations. However, in Approval-Based Moderation, many messages going unapproved of is not problematic; rather, it is beneficial or, at the very least neutral, as compared to traditional moderation systems. I consider in turn three kinds of messages, mediocre, low-quality and high-quality.

First consider a message of mediocre quality. Such a message, ideally, would not be shown to many users. In a traditional system, very few users may bother to moderate it at all. Even when some users do, the message may get mediocre scores, neither significantly raising nor lowering its score in aggregate, resulting in many people seeing such a message and disregarding it. With Approval-Based Moderation, however, a mediocre message that nobody cares enough about to approve of will disappear soon after its initial lifespan is over, having never or rarely been approved of.

Now consider a message of low quality. This message would ideally be shown to no users. After being consistently rated as low by several users, in a traditional system such a message may be so low it is never shown again. In Webbed Footnotes, no user might think such a message worthwhile and so it will not be approved. The same end result will occur once its lifespan expires: it will not be shown to any subsequent users.

Lastly, consider a high quality message, one that would ideally be shown to all users. After being consistently rated high in traditional systems or approved of in an approval-based system, the same end result, a long-lasting, highly visible message, will occur. Being rated not at all, in either system, leaves the message no better off than any other message.

In summary, for high-quality messages, Approval-Based Moderation compares neutrally to traditional moderation, since positive action on the part of users is required to raise a message up in both cases. For mediocre messages, however, Approval-Based Moderation performs favorably compared to traditional moderation, since the desired effect of reducing the message’s exposure is achieved. Finally, in the case of the low quality message, Approval-Based Moderation performs favorably since both achieve the outcome of showing the message to few users, but Approval-Based Moderation does so with less user effort.

Message Quality	Desired outcome (show to ... users)	Number of moderations needed for desired outcome		Result with no moderations	
		Approval-Based	Traditional Moderation	Approval-Based	Traditional Moderation
High	Many	Many	Many	Failure	Failure
Mediocre	Few	None or Few	?	Success	Failure
Low	None	None	Many	Success	Failure

Table 3-1. Comparison of Traditional Moderation and Approval-Based Moderation.

The above table demonstrates the effectiveness of approval-based versus traditional moderation. One qualification to the above table is how a mediocre message fares under traditional moderation. In the table it receives a question mark, because it is unclear what even an infinite number of moderations would do to a mediocre message because, by definition, a mediocre message might be rated high by some but low by others.

An additional and more important but philosophical qualification regards the lone failure of Approval-Based Moderation, a high-quality message receiving no approvals. Can a message that receives no approval really ever be considered high? If nobody felt the message was worth expending the effort in order to approve of it, then the message’s quality must not have been that high. Thus by definition, a high quality message is one that has been approved of many times. This definition may seem circular, but I propose that it is so in order to reflect the importance of the collective judgments of the participants in rating messages.

3.2.2 Herding and Information Cascades

In Webbed Footnotes, “herding” behavior and information cascades are significantly less likely than in other systems. Though preferential exposure is a desired outcome of moderation systems, and therefore some messages will necessarily get more attention than others, herding

behavior is dangerous because of the strong possibility that bad decisions can be made [6]. Webbed Footnotes decreases the likelihood of herding for two reasons. Annotation placement is never affected by approval, nor is approval data ever shown to subsequent readers.

Herding behavior is an economic theory that suggests that through observing other participants' behavior, groups will en masse converge on a single opinion; this can be the result of being influenced by their peers' choices, succumbing to an innate desire for conformity or acting on incomplete information based on the choices of a few initial actors [6]. Social learning theory supports this, proposing that watching others is a primary component of individual learning in social groups [3].

Since herding behavior is caused by basing one's own decisions on those of previous actors, one way to prevent herding is to deny users information about past actors' behavior. This is what Webbed Footnotes does. Unlike other collaborative moderation systems, Webbed Footnotes does not display previous users' approvals (or lack thereof) to subsequent readers. No reader knows whether a particular annotation has been read once or a thousand times, nor whether he or she is the first reader to read it, or will be the last reader to before it disappears.

Because users know that their approval actions will not be seen by other users, they need not consider approving or not approving messages to be an act of performance. In interacting with others, one's visible actions create impressions in the others' minds about what kind of person that individual is; as a result, people are often self-conscious of their actions, deliberately behaving in a certain way in order to effect particular desired impressions in others [23]. In a moderation system that shields a moderator's decisions from others, each moderator may feel free from the coercion of social pressures and, in this way, moderation can be seen as a kind of secret ballot. Those who hold minority opinions may be more inclined to express themselves if they can do so secretly, as being in of a minority opinion holds some users back from active participation [46].

In addition to prominently displaying a message's moderation ratings, the ranking of messages according to those ratings is another way moderation data can be presented to users and lead to herding. Urban Dictionary, Slashdot, and other communities have this practice, in which the most highly rated items are shown first. Because not all users are going to read all the messages, the highest ranked messages will get the most attention, which is in keeping with moderation's goal of preferentially showing higher ranked content. However, such ranking has impact on messages other than the ones moderated. To illustrate, consider two equally-ranked messages; call them A and B. When A is moderated up, it is presented before B and is now more likely than B to be seen by users. Even though B was not moderated at all, it was effectively moderated down relative to A.

By contrast, Webbed Footnotes cannot reorder annotations since they are so deeply connected to the locations given to them at their time of writing. Therefore, highly approved annotation cannot "upstage" other annotations nor crowd them out. In the case just described, A being moderated up would not affect B's position nor its potential to be seen by subsequent users. In short, a messages lives or dies on its own, unaffected by being in the presence of several higher-ranked messages.

Even in the absence of herding behavior, feedback loops are natural outcomes of the preferential exposure moderation systems seek. Because the actions of the earliest actors are the first ones observed by subsequent actors, those early actors have a disproportionately large effect on the final outcome; this is called being a first-mover. Though theoretical economics applies these principles mostly to financial markets, they apply to the online world as well; an example more akin to that of moderation systems is the linking structure of the web. Barabasi [4] demonstrated that the growth of the linking structure on the web showed strong preference for the web sites that arrived on the web earliest. Likewise, messages highly ranked in traditional collaborative moderation systems or those that were posted earliest are most likely to receive attention from moderators; this is the case in Slashdot, for example [38].

Even in Webbed Footnotes, there is the potential for a feedback loop besides those associated with first movers. If heavily approved annotations live longer, then they are shown to more users, and therefore have more opportunities to be approved even more, thus living longer, and so on. Indeed, in Chapter Four, where the results of the user study are presented, I discuss the distribution of approvals among messages, and the evidence for some feedback loops.

3.2.3 Lurking's Effects in Approval Based Moderation

At the end of Chapter Two, I discussed differing perspectives on lurking and whether it is a problematic behavior, settling on the idea that it is not, because it prevents the overuse of the scarce bandwidth of a channel. In an Approval-Based Moderation environment, however, lurking ceases to be unproblematic and begins to look more like free-riding.

Moderation, unlike posting, does not overuse bandwidth or compete for the scarce attention of users. Because moderation information is presented in the aggregate when shown at all, one thousand moderations take up no more space than a single moderation. Therefore, not moderating fails to have any beneficial side effects the way not posting does.

A second reason reading without annotating seems more like free riding is that content becomes a scarce resource. In a traditional environment, because anyone can read messages without limit, content is not a scarce resource for consumers. However, in Approval-Based Moderation, since each message has a limited number of views before expiration, messages are now a scarce resource because their views are scarce. As such, a reader who consumes message views but does not replenish the stock (i.e. by approving some messages) could be considered a free rider.

If Approval-Based Moderation creates a social dilemma where before there was none, then there must necessarily be some greater benefit that makes it worthwhile. I propose that one benefit is that lurkers are made to feel that they have an impact on the community. Lurkers can read content in most communities and have absolutely no effect on it. However, in approval based moderation, traditional lurking has the effect of "using up" message views. This may scare some lurkers away from any kind of participation at all, but alternatively it may spur some on to further action, either in the form of approving or posting. Nonnecke and Preece [46] suggest that a main reason lurkers lurk is that they feel it is not necessary to post. If lurkers realize that their very reading affects the community, then they may feel that their impact merits further action.

In Chapter Four, in discussing the impact of Approval-Based Moderation, I show that several users who did not annotate on their own did in fact approve others' annotations. Like other environments [38], it appears that in Webbed Footnotes there are users who annotate, some who approve, and some who do both or neither. Since Approval-Based Moderation is designed to be less burdensome and more optional than other kinds of moderation, it is hoped that more lurkers will choose to do it.

3.3 User Interface Elements

In this section I discuss the details of the Webbed Footnotes user interface. Because Webbed Footnotes integrates with the web browser, it is important to ensure that the added functionality of Webbed Footnotes does not interfere with other browsing behavior. Though in principle Webbed Footnotes can annotate any page on the web and therefore annotation is always a possibility, in practice it is likely that the vast majority of web documents would not be annotated, nor would individual users want to annotate more than a small fraction of the pages they visit. Therefore, except for the small fraction of pages a user seeks to annotate or read annotations from, the Webbed Footnotes interface should be "out of the way."

Despite this, when a user does seek to annotate a page or to read others' annotations, the tools to do so should not be far away. Therefore, the browser with Webbed Footnotes is always in annotation reading mode; that is, when a webpage is loaded, any annotations on that page are displayed immediately, along with the webpage's original content.

One of the most important goals of Webbed Footnotes is to design the annotation system such that the original content of even a heavily annotated document remains legible. In Chapter Two it became clear with examples like yeahbutisitart.com and the HLS Annotation Engine that when a webpage is augmented with user-provided content, there is a substantial risk of the augmentation overshadowing and even obscuring the webpage's original content. The annotation marker, especially its hollow center, was designed precisely to prevent such obscuring.

3.3.1 User Identity

Each Webbed Footnotes user is represented by a unique username and a color selected from one of twelve possible choices. This username and color will appear on every annotation the user writes.

The username is a particularly powerful marker of identity, especially in text-based environments. Though Webbed Footnotes is not strictly speaking text based, in that there are other visual elements besides text, it is *conceptually* text based. That is, Webbed Footnotes is *about* text. It is the text-based document that establishes the geographic bounds for its primary interactions, reading and writing prose.

Participants in electronic social spaces often choose a single username and use it (or variations of it) consistently across many different spaces. Users can be playful or deeply personal in

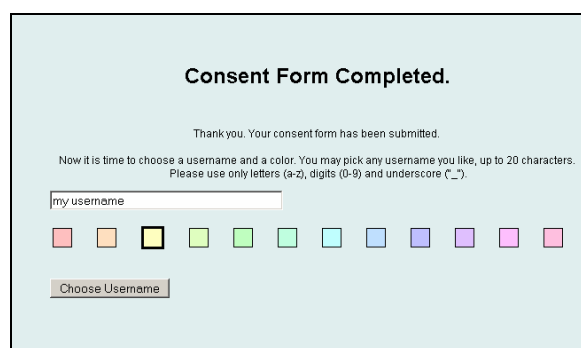


Figure 3-2. The username and color selection form.

choosing a username, often making references to their offline identity, location, hobbies or other interests [5]. The choice of a name is like putting on a public face, like choosing how one would like others to see them. An important part of establishing one's identity, online or offline, is making choices that are designed to convey a desired impression [23]. The username is one such choice.

After representing oneself with a particular username for some time, a user develops a reputation among other participants that is often tightly connected to that username. Moreover, it comes as no surprise that users often grow emotionally attached to, and identify with their usernames, after having invested time in choosing that name and in building a reputation that is associated with it.

In addition to the username, each user is represented by a color. Though the colors are not unique and people cannot distinguish among more than a handful of colors, nevertheless color can be useful for distinguishing among several users in an online space [16]. Here, an example of the usefulness of color is the ability for a reader to identify at a glance whether an annotated page is filled with comments from many different people, or with many comments from a single person. Color, in this case, distinguishes the lively conversation of an active group from the windy monologue of a zealous individual.

3.3.2 The Annotation

Because part of the goal of Webbed Footnotes is to allow even a heavily annotated page to remain legible, making annotations visible but not obtrusive seems to be an exercise in contradiction. The first step in achieving both of these traits was to, unlike most other annotation environments, avoid displaying the whole of an annotation directly on top of the page. Instead, each annotation is represented by an *annotation marker*, a hollow rectangle in the user's chosen color. Only when the user places the mouse cursor over the annotation marker does the marker become solid, demonstrating that that is the currently visible annotation.



Figure 3-3. An annotation marker. Because this is not the currently displayed annotation, its marker is hollow.



Figure 3-4. Annotation marker for the annotation currently displayed. When the author views it, it has no icons.



Figure 3-5. Annotation marker for the annotation currently displayed. Three options are presented in the form of "User page", "Reply", and "Approve" icons.

When an annotation is visible, if the author is someone other than the user who is viewing it, three options become available. First is a silhouette-like icon, which directs the reader to the annotation's author's User Page, which is discussed in greater detail below. Second is a reply icon, which allows the user to write a followup annotation. Lastly, the "thumbs up" icon allows the user to approve of that annotation.

The annotation marker's rectangle is always the same height. Its width, however, is variable and represents the length of the text in the annotation.¹⁵ This gives the reader an idea of how

¹⁵ Subject to a minimum width of 16 pixels and a maximum of 150. The minimum ensures that each annotation is visible to users and the maximum ensures that too much of the screen is not overwhelmed.

long an annotation is before reading it. Annotations that are replies to others' annotations are displayed in a single column immediately beneath original ancestor annotation. Therefore, this is not threading in its strictest sense, where each message is directly attached to its parent, but a looser threading in which related messages are grouped in a linear order. In this way, a conversation can be held together but without the visual complexity of threading.

The Approval button with the thumbs up icon indicates approval of an annotation. The thumbs up icon is then replaced with a checkmark. The checkmark indicates that the user has already approved that annotation and may not approve it again.



Figure 3-6. The checkmark displayed when an annotation has been approved, and the thumbs up button selected for approval.

As discussed above, annotation text is not immediately displayed on the screen. Instead, the user must move the mouse cursor over the annotation marker in order to read the annotation, the text of which is displayed in a floating box at the bottom of the screen, evocative of the placement of a footnote in a printed text. This solid box is colored in the color of the annotation's author. It contains the author's name and the date and time at which the annotation was written.



Figure 3-7. A floating box containing an annotation text. When an annotation marker is moused over, the floating annotation box is displayed at the bottom of the browser window.

By displaying the annotation text in a separate box but having the annotation marker interspersed with the rest of the document, two competing desires are simultaneously achieved. The desire to have user annotation share the location of the document's original content, and so to be on equal footing with it rather than be subordinate to it, is achieved through the annotation marker's location. At the same time, the annotation is set off as a distinct voice by having its text presented in a well-defined box at the bottom of the page.

3.3.3 Button Bar

The button bar is the only element of Webbed Footnotes that persists whenever Webbed Footnotes is installed. It sits at the top of the browser window, immediately below the address bar. The buttons on the button bar control all aspects of Webbed Footnotes; no menus are used.

The first two buttons, Hide Annotations and Hide Authors affect how annotation markers are displayed on the page. The first toggles the visibility of any annotation markers. The second, if the markers are visible, toggles whether or not the authors' names appear as well. The lone colored box is arguably less invasive than the same box along with the text name beside it, especially if the name's typeface, size and color are similar to the existing text on the page. Additionally, the authors' names may be less important to a user who is more concerned about the content of the annotations and less about who wrote them.

The Annotate Page button puts Webbed Footnotes into annotation mode. When this button is pressed, the user may click anywhere on the document to choose a location for a new annotation. Upon clicking the page, a text frame appears, in which the user can add the text of a new annotation.

During the New York Times user study, the NYTimes button redirected users to the study's customized New York Times cache. This was done as an alternative to resetting the user's homepage, which might have been seen as too invasive. Further, trusting users to bookmark the study's NYTimes page was too risky; many users might forget where the page is located and never visit it. In a version of Webbed Footnotes that can annotate any page, this button would necessarily be absent.

Finally, the My Annotations button redirects users to their own User Pages, discussed above, and the Help/FAQ button redirects users to an extensive help page with instructions on use.




Figure 3-8. The Webbed Footnotes button bar.

The button bar also serves as a constant reminder that the user is empowered to annotate any page. This constant reminder, however, is not always a good thing. In the follow-up survey, several respondents commented that one of the things they disliked most about Webbed Footnotes is that it took up what they considered to be an excessive amount of screen space. Those who felt this way appear to have a background in technology; anecdotally, for this group, screen space tends to be more precious. That said, it may have been possible in building Webbed Footnotes to make the buttons take up somewhat less space vertically, thus partially alleviating this concern.

3.4 Navigation and Discovery

In an environment where the most interesting content – that which is written by other users – is sparse and widely distributed among a large space, it is helpful to have methods of aggregating and indexing that content and providing tools for users to find and navigate through it. It has been shown that users consistently appreciate knowing about the actions of other users in relation to the use of shared or collaboratively maintained content [30,17,68].

In a printed book, one of the disadvantages of interspersing one's marginal comments throughout is that there is no easy way to keep track of them all, cross-reference them, and so on. As a result, many readers will collect their thoughts in a reader-authored index, either on separate sheets of paper kept in the book, or on the several blank pages that are left over from binding in the back of the book [31]. By aggregating one's comments, the reader can keep track of all the comments as well as record pointers into the text (e.g. page numbers) that show what a particular comment is in reference to. Likewise, in Webbed Footnotes, it is useful to be able to bring together references to all users' contributions while maintaining pointers to the contributions in context.

In Webbed Footnotes, user-authored annotations are attached to individual web documents. Absent any tool for aggregating them, users would have to spend a great deal of time searching for others' comments. Webbed Footnotes has two built-in methods for navigation and discovery of others' content, the Sidebar and the User Page. A third method was built for the New York Times user study.

3.4.1 The Sidebar

The Sidebar is a gray box that is found at the top, right-hand corner of each document. It lists all of the annotations in that document and displays the color and username of the author, the time the annotation was written, and the first few words from its text. Especially in a long document, it may be difficult to see where all the annotations are. It is, after all, one of the goals of the design to not make the annotations themselves overly obtrusive. However, there are times when seeing references to all of an article's annotations at once is helpful; for example, one might want to see what new annotations have been added since their last visit, get a sense of who and how many people have annotated this article, and so on, without having to scan the whole document. For these reasons, the Sidebar is added to each document.

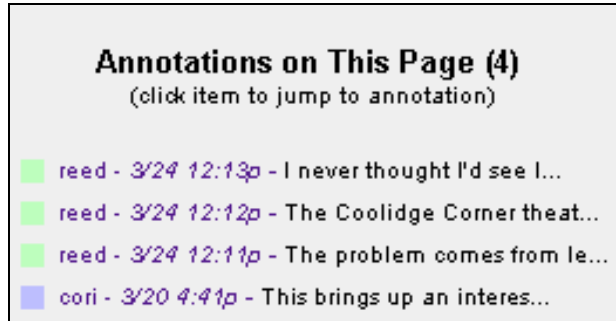


Figure 3-9. The sidebar, shown on the right hand side of each page and detailing the annotations on that page.

3.4.2 The User Page

The User Page is the second of the two aggregation tools built into Webbed Footnotes. Each user's User Page contains a list of all annotations written by that user, including the text of each annotation and a link to the article page to which the annotation is attached. The listing is in reverse chronological order, listing the most recent annotations first.

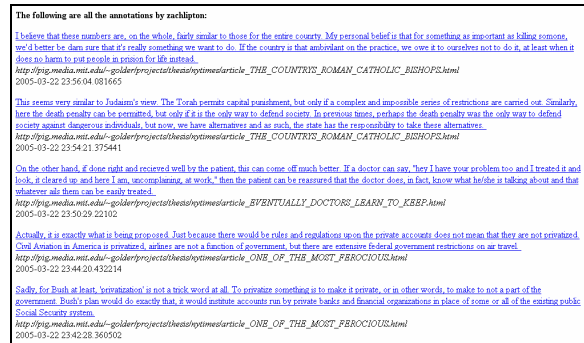


Figure 3-10. The user page, containing a list of all the annotations by a single user.

The User Page is accessible two ways. Users may always access their own User Pages with the My Annotations button on the button bar. This is important because in a multi-user system, users should always be able to see what information about them is available to others.

Additionally, any user's User Page can be accessed by clicking the silhouette icon that is displayed while mousing over an annotation marker. The silhouette icon redirects the browser to the User Page of the author of that annotation. With this method, users can better get to know one another's writings by being exposed to them all at once.



Figure 3-11. The silhouette icon.

3.4.3 The New York Times User Study

For the New York Times user study, which is described in greater detail in the next chapter, users were directed to a static, cached copy of the day's New York Times articles. Because there were far more articles than there were users, the likelihood of any one page having any annotations was quite small. If users were to seek out the annotations of others, they would have to look at a great deal of pages before encountering one. In order to combat this, Webbed Footnotes denoted on the index page, where all the articles' headlines were listed, how many annotations the page had. This was accomplished by placing a small "note" icon next to the

headline for each annotation. A page that had been annotated once, for example, had one “note” and a page annotated five or more times had five.

Though the “note” icon denoted which articles had been annotated, the headlines of those articles remained in their original order in the index; there was no sorting or re-ordering taking place. This was done in part to prevent feedback loops in which the articles that received annotation were disproportionately more likely to receive future annotation. Though such behavior is inevitable and arguably has a positive effect on interaction, in designing Webbed Footnotes it seemed desirable to encourage users to expose themselves to as wide a variety of the content as possible.



Figure 3-12. The “note” icon.

3.5 Implementation

Webbed Footnotes is an extension to the Firefox web browser.¹⁶ The Firefox browser is an open-source, free web browser released by the Mozilla Foundation.¹⁷ The Mozilla Foundation’s mission is “to preserve choice and innovation on the Internet.” Though Firefox currently has only a small share of the browser market, especially compared to Microsoft’s Internet Explorer, which has approximately 90%, it has been downloaded 25 million times and has a market share measured at anywhere from 5.7% to 8.5% as of February 2005 [34]. Firefox is being presented as a better alternative to Internet Explorer because of the latter’s perceived security flaws. Also, many propose open source software of all kinds as an alternative to commercial, closed-source software as a sociopolitical statement. These issues are important but, in the end, Firefox was chosen as a platform for Webbed Footnotes due to the ease of creating extensions for it.

Firefox is designed to be highly extensible by third-party developers. All extensions are written in JavaScript and XUL, or the XML User interface Language.¹⁸ XUL is an XML format for designing the user interface for web pages, and is also a project of the Mozilla Foundation. Because JavaScript and XML are both interpreted rather than compiled, the source code is included in each extension. Each extension is packaged in an XPI file, which is a simple archive based on the ZIP format, and contains a predefined directory structure containing the Javascript and XUL code, any supporting files like images, and an XML file containing several details about the extension.



Figure 3-13. The Firefox web browser.

The Firefox extension comprises only the user client portion of Webbed Footnotes. The server component of Webbed Footnotes consists of a web server, a relational database, and several scripts written in the PHP language. In principle, any web server or database application can be used; for the current implementation, Apache and PostgreSQL were used, respectively.

¹⁶ <http://www.mozilla.org/products/firefox/>

¹⁷ <http://www.mozilla.org/>

¹⁸ <http://www.mozilla.org/projects/xul/>

The PHP scripts collected and distributed, and the SQL database stored, user and annotation data. Of the nine tables in the database, four stored data necessary for the proper running of the service; the remaining five stored data for the analysis and evaluation discussed in Chapter Four, including records of webpage and annotation access. The following table describes the tables and their contents:

NAME	PURPOSE	CONTENTS
userinfo	Service	Users' personal info, including ID number, username and color.
annotations	Service	Text and location of annotation, IDs of its creator and the page to which it is attached.
urls	Service	URL string and the number of annotations on the page.
vet	Service	Annotation IDs and the IDs of the users who vetted them.
annohits	Analysis	Number of times an annotation was read, and by whom.
urlhits	Analysis	When each webpage was accessed.
userpage	Analysis	When each individual users' User Page was accessed, and from which annotation.
userpagelink	Analysis	When a link on a users' User Page was clicked, and by whom.
consent	Analysis	Full name and email address of participant, stored when signing the consent form required by the Committee on the Use of Human Subjects.

Figure 3-14. Database tables and their contents.

As described in this table a great deal of data is collected about users' behavior. In Chapter Four, this data is analyzed in order to develop an understanding of how users used the software, including how frequently and under what circumstances they read articles and read, wrote and vetted annotations.

In order to communicate with the database, the browser extension calls the PHP scripts in a hidden web browser frame. HTTP is a stateless protocol; there is no persistent connection between the client (i.e. browser) and the server. As a result, in order to pass data to and from the browser and server, it would ordinarily be necessary to reload the entire page, thereby reestablishing contact with the web server. The downside of this process is that it is very jarring for the user's experience to constantly be reloading a page whenever he performs some action, especially when the action is something as simple as reading or vetting an annotation. Reading an annotation is supposed to be a natural, flowing process that does not break the user's concentration. Clearly, reloading a page would have a very negative effect on concentration. Instead of reloading the page, we use a hidden frame to call the PHP scripts. This frame is like an ordinary browser window, except that it is invisible.

An event handler keeps track of when each window (the visible one in the browser or the invisible frame) loads a new page. When the main window loads a new page, it calls the appropriate PHP script in the hidden window. When the hidden window loads a PHP script, it calls the appropriate function in the extension's JavaScript code, usually to modify the contents of the main window.




Figure 3-15. The Webbed Footnotes system architecture. From left to right: the user, the web browser (containing the Webbed Footnotes extension), the hidden browser frame, several PHP scripts residing on the web server, and the SQL database.

The system architecture figure above depicts all the communication steps that take place between the various components. In summary, the user interacts with the web browser and the Webbed Footnotes extension. The extension uses the hidden browser frame in order to call the PHP scripts, which stores and retrieves data in the database.

3.5.1 User Authentication

As discussed above, users select a unique username and a color when signing up. In addition to assigning a username and color, the signup form also assigns a unique ID number to each user. This ID number, along with the chosen color and username are stored in two places: the SQL database on the server, and the Preferences system of Firefox. Firefox contains a Preferences system that stores browser settings, like the default download directory and options for handling popup windows and cookies. This system also allows third-party extensions to store information in key/value pairs much like the Windows Registry. Webbed Footnotes stores three keys: `annotation.username`, `annotation.userid`, and `annotation.color`.

Of these three browser-stored preference keys, the only one that truly matters is the user ID. It is the presence of the user ID that tells the extension that the user has signed up and to not redirect the browser to the signup/consent form. Moreover, with every call to a PHP script, the user ID is sent, in order to keep track of the user's actions. Sending an invalid user ID, either a nonexistent one or one invalid for the action performed, will fail gracefully, preventing any modification to the database. This step partially ensures that users will not be able to pass malicious data to the server and affect the integrity of the database.¹⁹

An unintended consequence of the signup process is that each user account is tied to a single Firefox installation, i.e. a single computer. This prevented users with multiple computers from being able to use Webbed Footnotes from each. This was annoying to some users, who reported creating additional accounts for the purpose of using Webbed Footnotes from multiple computers. In hindsight, a username/password system would have been much more effective and would have prevented this problem. In the data analysis performed, an attempt has been made to, where possible, find multiple accounts representing a single person and merge them.

3.5.2 Loading, Reading and Writing Annotations

¹⁹ An additional note on data integrity: All data sent to each PHP script is checked for consistency. For example, a new annotation can be added to the database only when it is accompanied by a valid URL for a page that exists in our database, and an annotation may be vetted only when the vetting script is given the vetter's user id, the annotation's ID, and the IDs of the annotation author and the page to which it is attached. Needless to say, this process will not deter the determined, malicious attacker, but it should be sufficient to stop the casual curious user.

Every time the browser loads a new webpage, it asks the annotation server whether that page contains any annotations. This is accomplished, as mentioned briefly above, through the use of a JavaScript function called an “event handler.” The event handler monitors both the visible and the invisible window frames, and detects when a new webpage has loaded – i.e. a “load event” has taken place. As its name implies, it handles this event by taking the appropriate action.

In the case of the main window, when a new page is loaded, a PHP script is called in the hidden frame, with the ID of the user and the URL of the new page the user has visited. This URL is added to the database (if necessary) and the database is checked for any annotations to that page. If any exist, they are “printed” in a delimited format to the body of the script’s output (recall this is in a hidden frame). When the script’s output – the annotations for that page – are finished loading, the event handler is again invoked. This time, it recognizes that the page loaded is in the hidden frame and is the one containing annotation data. The annotation loader is then invoked, which reads the data from the hidden frame and displays the annotations on the screen.

The annotations, like all other content Webbed Footnotes augments pages with, are in the form of “floating” <div> tags. The <div> tag is an HTML element that defines a rectangular “division” in a document. A “floating” <div> tag is one whose location on the page is not in-line with the rest of the document, but rather “floats” over the document at prespecified x,y coordinates. Each annotation marker’s hollow rectangle constitutes a single <div>. An additional <div> tag is used for rectangle containing the three icons displayed upon mousing over an annotation marker, and another for displaying the annotation text at the bottom of the page. These <div> tags are always present, but are marked “invisible” when not being accessed. The sidebar, displayed on any page that contains annotations, is itself a <div> element; it is always visible.

When an annotation is read by mousing over its annotation marker, a PHP script is called in the hidden frame that alerts the server that the annotation has been read. Likewise, when an annotation is approved by clicking the thumbs up button, another PHP script called in the hidden frame alerts the server of the approval. In this way, any action taken is silently recorded in the database.

Similarly, newly-authored annotations are sent to the database via the hidden frame. When the user writes text for an annotation in the HTML form provided, the Submit button calls a PHP script in the hidden frame that stores the new annotation in the database. One additional step is necessary, however; the new annotation must be displayed immediately on the page so that users can see the immediate results of their actions. Therefore, the new annotation is drawn “manually” on the page, using the same process as described for other annotations above, without having to reload the page.

3.6 Limitations and Extensions

As with any project, Webbed Footnotes has areas in which it could be improved upon. A number of interesting problems and challenges remain; two of them are discussed here.

3.6.1 Browser Extensions and Trust

Undoubtedly, the requirement that users use a specific browser and download and install software that is an extension to that browser presents a rather high barrier to participation for many users. Even users who are technologically literate and capable of performing the necessary downloads and installations might balk at the time and effort involved. It is likely that the technologically savvy users, rather than the naïve ones, might be even more skeptical about the installation of yet another piece of software, especially in the web browsing domain, which has a reputation of being the locus for unsafe or malicious activity in the form of “malware” and “spyware.”

However, at some point, the user must pay some cost, in effort, time or disk space, in order to receive the benefit of some tool. If Webbed Footnotes is to be a tool that has scope across the entire web, then it must be ever-present in the browser. For this version a browser extension made the most sense. Only a relatively small user group was sought, and a group of early adopters – those most likely to use Firefox – was an acceptable choice. They would be the most likely, it seemed, to be willing to participate in something new.

Early in the development of Webbed Footnotes, a proxy server was considered as an alternative to a browser extension. A proxy server is a web server that sits between the user’s web browser and the rest of the web; it can pre-cache, filter and, if necessary, augment content before delivering it to the user.²⁰ The proxy server could add all the necessary JavaScript code and annotations to the webpage before it reached the end user. A proxy server, however, is invisible, as compared to the browser extension, and fewer users have likely heard the term “proxy server” before or know what one is. Asking users to install a piece of software seemed to be a great deal more achievable than asking them to change an obscure browser setting.

When a third party is modifying one’s web browsing experience, a measure of control is given over to that third party. Whether a browser extension or a proxy server, persuading the user that the third party is trustworthy is paramount. Steps in instilling that trust include having the backing of a respected institution – in this case, MIT – as well as having highly usable software and a professional-looking design [21].

3.6.2 Ephemeral Documents and Annotation Placement

Attaching an annotation to a webpage becomes significantly more challenging when the webpage in question is itself constantly changing. Since the web’s emergence over a decade ago, there has been a transition from web pages being mostly static HTML documents to being dynamic documents that provide new and/or custom content each time a user visits it. This presents several technical and interaction-level challenges for annotation systems.

If a page’s content is ephemeral, should annotations of that page be likewise ephemeral? It seems that, if the content that an annotation is modifying is moved or absent, then annotations to that content should either move or disappear as well, otherwise the annotation would be

²⁰ Proxy servers are commonly used to filter content, in the case of public libraries, or speed up the delivery of content, in the case of large commercial networks.

lacking its context and would likely not make sense. This presents the technical problem of identifying which portion of the content the user intends to annotate.

Even on pages that have static content, one problem to be contended with is that each web browser renders pages slightly differently. Varying screen resolutions and browser window sizes may result in text wrapping differently on each user's screen.

Annotations in Webbed Footnotes are drawn on the page according to the x,y coordinates the annotation author selected when writing it. This is done with no regard for what part of the content is being annotated, and so Webbed Footnotes is currently susceptible to both the problems of dynamic content and shifting text wrapping.

In order to avoid these problems, however, Webbed Footnotes supported only a predefined set of static pages in its two user trials. Each document in the user studies, detailed in the next chapter, shares a similar structure, consisting of a 100-pixel left margin and a 500-pixel column of text, with any remaining horizontal space as white space in the right margin. Even a very low screen resolution of 640x480 or a browser window that has been resized to take up less than the whole screen would most likely be able to represent the text in its intended form without wrapping lines differently. This step partially ensures that each user will see the document the same way, regardless of the details of the configuration of the computer and the web browser. This is called the "frozen" document method [9]; it is effective for what it does, but it does seem somewhat limiting.

Possible solutions to both of these problems rely on the ability to identify which portion of the page's content the user intends to annotate. This is difficult. Consider an annotation placed in white space between two paragraphs when one of those paragraphs goes away. It may be unclear which of the two paragraphs the annotation was indented to augment, the one before it or the one following it. The relationship between the document's HTML structure and its visual appearance is often unclear even to savvy users; for example, an HTML table cell may or may not extend well into the whitespace beside where its text ends. Therefore it would not be enough to anchor an annotation to the HTML element closest to its x,y coordinates, lest annotations become attached to the wrong element.

One way to alleviate this problem is to visually show the user which HTML element a new annotation would be attached to. When in annotation mode, as the user moves the mouse about the screen, highlight in a bright color the HTML element immediately below,²¹ while still allowing the user to choose an x,y point. In such a case, the point would not be relative to the page as a whole, as Webbed Footnotes does currently, but instead would be relative to the top left of the HTML element. This would not be a perfect solution, as the width and height of that element may continue to render differently on different browsers, but would be an improvement over absolute placement on the page.

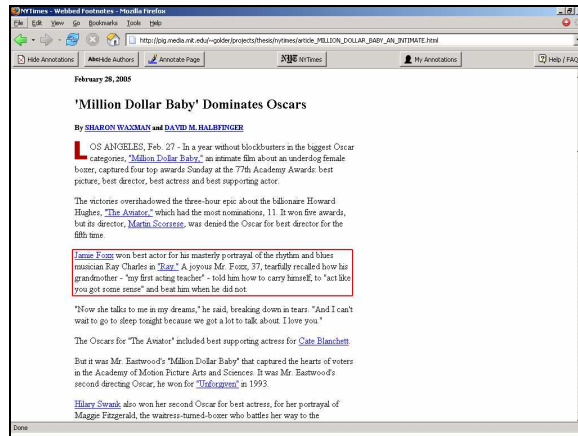


Figure 3-16. Webbed Footnotes with a highlighted paragraph indicating potential placement of a new annotation.

After determining which HTML element an annotation is intended to be attached to, detecting the presence of that element becomes easier. A hash signature can be computed for the contents of that element. On subsequent page loads, if no element that hashes to the same value is present, we deem the context for that annotation to be missing, and do not load the annotation. This practice has been employed by several systems [9] and are indeed more “robust” than simple coordinate-based positioning.

²¹ The Firefox DOM Inspector does something similar to this. The DOM (“Document Object Model”) Inspector is a developer’s tool for visually exploring the structure of an HTML document.

4 Evaluating Webbed Footnotes

Two user studies were performed in order to understand how Webbed Footnotes might be used by actual conversational groups for engaging in discussion and in shared interpretation of a group-relevant set of texts. The two user studies consisted of a small graduate discussion seminar, wherein the relevant text was a set of three assigned readings, and a group of news and current events enthusiasts, wherein the relevant text was a daily-updated set of news articles from the *New York Times*.

These two trials differed in size ($n=7$ and $n=97$) but, more importantly, they differed in group composition as well. The classroom study consisted of a group of students who knew one another already, having interacted in the seminar room weekly for the past several weeks, and who knew they would continue to do for at least several more weeks. Many of these students are labmates and attend the same university, live in the same city, and so on. Their ongoing exposure to one another as scholars and peers means they are committed to living and thriving in their social group; it is a group they cannot abandon easily, nor would they want to. Because of this commitment, they are more likely to seek positive interactions with others because to do otherwise would carry too great a social cost [24].

In contrast, the second, larger user study consisted of a group of strangers recruited on the internet. Unlike the students, this group of people do not know one another, nor need they have any expectation of encountering one another again. They have no commitment to one another, and so they have what Goffman calls “a problem of resolve” [24]. Initially, at least, there is little keeping them using the system, so there is little cost in abandoning the system and the other members for any reason. Indeed, encouraging participants to continue participating proved to be a problem. Further, since abandoning the system carries little cost, there is also little cost incurred from having negative social interactions with one another. In such an environment, then, the likelihood of conflict would be predicted to be higher. However, in this user study, there did not appear to be significant conflicts among any users.

The two groups also have come together for different purposes. Many purposes are cited for annotating text documents, but two among them are collaborative scholarship and the expression of personal opinion [31]. Collaborative annotators like seminar participants seek the former, ultimately desiring a shared understanding of a text and of the questions or deliberations it entails. Even where disagreements on positions occur, learning about and consideration of those positions is the goal. In the latter group, where no ultimate goal of educating the participants exists, an environment that fosters discussion may be more conducive to disagreement for the purposes of “winning” that disagreement or swaying others.

These correspondences between group characteristics and goals and practices are neither exhaustive, exclusive nor deterministic; in observing how these groups interacted through Webbed Footnotes, it was interesting to observe what similarities and differences are exhibited in their behavior.

4.1 Classroom Study

The scholarly annotation of texts, the initial inspiration for this thesis and Webbed Footnotes, was the goal for the classroom user study. Especially in a seminar, an overwhelming

proportion of classroom time is generally spent discussing the assigned readings. The students are supposed to consider each other's alternative viewpoints, and learn from one another. In many seminars, however, it is often the case that discussion takes a little while to ramp up, for getting through the preliminaries of discussing general themes and "breaking the ice," so to speak. A large proportion of these preliminaries are spent establishing common ground among the students, as they relate their individual experiences with the readings. It often takes some time to get past these preliminaries and into the more complex issues raised by the readings. If the students can begin to have a shared experience outside the classroom, then hopefully once they reach the discussion room they will already have a sense of shared experience with the readings.

Though the small size of the group I studied means that evaluation of their use is far more anecdotal and brief than that of the large *New York Times* group, there is nonetheless ample reason to believe that Webbed Footnotes was at least somewhat successful in getting students to engage with one another through the materials before reaching the classroom. Through Webbed Footnotes, the students pointed out questions that the class ought to address in the class the following day, and also posed questions of their own. The annotations made reference to their homework assignments, as well. The most promising result was that one student reported that annotations made reading online "a less isolating experience" and made her feel "as though you are reading as part of a group."

However promising these outcomes are, these are things that students do in other course-based online environments like discussion boards. Indeed, many of the benefits of Webbed Footnotes I discussed earlier, like Approval-Based Moderation and a heightened importance of document legibility, do not really apply to small, private groups. They are not likely to post inappropriate material for the reasons discussed above, and so a moderation system is far less necessary, and the private nature of their community means that the underlying documents themselves may be the property of the class, thus eliminating the annotator-author tension that might result when a public document is annotated. An educational scenario that might serve Webbed Footnotes well is one in which students could critique one another's work by annotating it. This preserves the separation among the students' individual contributions – whereas a wiki, for example, may not – while allowing commentary to be located in context – which a bulletin board may not.

4.2 New York Times Discussion Study

The *New York Times* discussion study comprised 97 participants recruited on the web. The participants were invited to use Webbed Footnotes daily to read and annotate a pre-cached version of the *New York Times*, which was located on a web server at the Media Lab. The articles were updated very early each morning; new articles were added, but old articles remained in place, so that any conversations going on within that article could continue beyond the day the article was published. The system was running for a total of four weeks between March and April 2005.

Because the subject pool for this study were to read and discuss news, subjects were recruited from places online where such people might congregate. These included news and current events discussion boards on the *New York Times*' own website and on community bulletin board site *craigslist.org*. Additionally, Webbed Footnotes was mentioned by several weblogs, ranging in visibility from highly prominent to likely rather obscure.

In this larger, public study, I was interested in testing the effects of Approval-Based Moderation, and so participants were randomly assigned to two groups, a “test” group in which annotations had a lifespan of 10 views, and a “control” group in which annotations had no expiration. Both groups had the “thumbs up” button, however; the groups were given different explanations as to what effect the button had. Both groups were told the following:

Approving annotations tells the Webbed Footnotes system about the kinds of contributions readers are making. By approving the annotations you think are worthwhile, you help us get a sense of how happy people are with the quality of others' contributions.

Additionally, the “test” group was told this:

More importantly, every time you “approve” an annotation, **it will be shown to more people**. Normally, each annotation is shown only to a few readers before it expires. But if you see one you like, “approve” it, and we'll show it to even more readers. When everyone “approves” the annotations they like, it's more likely you'll see interesting annotations and be spared the uninteresting ones.

Beyond the difference in whether approval had an effect, the two groups' systems were identical.

There is a great deal of competition for the attention of web users and, unfortunately, the test group participants were not active enough over the course of the trial for any expirations to occur. Therefore, there was no difference between the two groups in their experimental conditions and it is not possible to make any definitive claims about the functioning of approval based moderation. This is disappointing, but it is impossible to evaluate how a system deals with an overabundance of negative content when there is not an overabundance of anything at all. This setback highlights what is perhaps the greatest challenge of evaluating a conversation space: it is difficult or impossible to construct community, especially in a short amount of time. A community is a group of people bound together by social ties developed over time through repeated interaction, and must grow as a result of users coming together; it cannot be realistically or meaningfully imposed on an otherwise unconnected group of users.

Part of the problem is a “chicken-and-egg” problem. Namely, people are attracted to other people; they like watching them and being in their presence [70]. They might, therefore be more likely to show interest in a system that is highly populated. This is problematic, because a new system that is unpopulated might face significant difficulties attracting enough initial users to reach a point at which the system is highly populated enough to attract and retain newcomers. Indeed, one of the chief reasons subjects gave for spending little time with Webbed Footnotes is, ironically, its lack of contributions from other users.

It may have been possible to set up a “toy” version of Webbed Footnotes, wherein actual conversation was not necessary. It would have been possible to stock articles with annotations, have a group of users vet or not vet annotations, and then have a second group of users rate the quality of the system both with and without the vetting information affecting annotation placement. Though such a test may have been possible, it would not have been naturalistic enough to be able to draw meaningful conclusions about the system as a whole. Despite the fact that there is less data than might have been ideal, there is still a great deal that can be learned from the users and their activity. For example, from the 97 subjects, there are 162 people-days of activity and 131 annotations, which are analyzed statistically as well as

qualitatively in the next section. Since there was no discernable difference in the overall behavior of the two groups, for the remainder of the evaluation, all analyses combine the two groups into a single group of 97 users.

4.2.1 Subject Activity

Like most conversational systems, levels of activity among individual subjects ranged greatly, with a few subjects participating very frequently, and most subjects participating comparatively infrequently. I measure activity two primary ways, the number of distinct pages (articles) a user visited and the number of annotations a user wrote. These measures, of course, describe their reading and writing behavior, respectively. Most users read only a few news articles; 80% of users read 9 or fewer articles, and 4% read more than 30 articles. Writing activity was somewhat higher, as almost half of all users authored annotations. While 57% of users did not annotate at all, 7% of users wrote more than 5 annotations. However, these two activity levels are strongly correlated; more active reading is strongly associated with more active annotating ($p = .000$).

Additionally, more active annotators are also more active vetters ($p < .002$). Though few users vetted annotations – 87% did not – those who did were the same users who wrote the most annotations.²² Consistently, we see that users who are more likely to participate in one way are also more likely to participate in other ways. In this, Webbed Footnotes seems to follow statistically in the pattern observed in other electronic communities anecdotally.

These patterns notwithstanding, a significant proportion of annotators did not vet, and likewise several participants who vetted annotations did not write annotations themselves. Specifically, of the 13 participants who vetted annotations, 8 of them (62%) also wrote annotations; those 8 users make up the 19% of annotators who also vetted others' annotations. That approving of others' annotations and authoring annotations oneself are each activities that can be practiced without also practicing the other is important. It demonstrates that users who may be reluctant to author annotations for all the reasons suggested in Chapter Two may be willing to have an impact on the makeup of the community through vetting. However, it also means that even some of the most active users in terms of writing may yet be unwilling to put time into vetting others' annotations. This is, however, not a reason that jeopardizes Webbed Footnotes' potential; on the contrary, the moderation system functions perfectly well even if very few people vet annotations. The result will be that annotations have shorter life spans than they would if more people were vetting.

Overall, the generally low levels of annotating and vetting demonstrate that, like most electronic conversational environments, a large percentage of participants will simply want to lurk. In Webbed Footnotes, this is important to know because lurkers do have an effect on the community's information resources, as they consume the scarce resource of annotation hits. One way of addressing this issue might be to have two tiers or classes of user, one whose members are interested in vetting and so they contribute to the decrementing of the annotation hits, and one whose members are uninterested in or unwilling to vet annotations and do not affect annotations' hit counts. By default, users could be in the latter group and only by vetting an annotation would they be moved to the former group. If a user did not continue vetting at

²² Users could not vet their own annotations.

least once in a while, that user might be moved back to the latter group. The general idea is that such users should not adversely impact the annotations they read but would never vet.

4.3 An Analysis of Annotating Behavior

Earlier in Chapters One and Four, I discussed several reasons why people might annotate documents, both physical and electronic. In this section, I examine the content of annotations written in the *New York Times* study. This analysis is useful because it allows us to understand more closely how people use annotation in a shared context when reading and annotating news articles. This is a very important set of content, because of the prevalence of current events discussion on the web. Though this analysis cannot reveal with certainty how another group of users might use annotation on a different body of content, it provides a look at what kinds of comments the users were inspired to make within the context of Webbed Footnotes and the features with which it provided them.

I performed a qualitative analysis of each of the 131 annotations in the system, developing a codebook of eighteen codes using a grounded theory approach [22]. The coding was performed by two independent coders, who could assign as many or few of the codes to each annotation as deemed necessary. Disagreements among the two coders were resolved through discussion until a set of codes for each annotation was decided upon. No disagreement lasted beyond brief discussion, demonstrating quite good reliability among the coders. Following grounded theory, the codes were developed inductively from several readings of the text [22]. Through each reading, actions were identified and classified based on the answer to the question, “what is the annotator doing here?”

1. critiquing / questioning / engaging author	10. raising related question
2. inviting response	11. reply to another person
3. blank	12. deixis
4. test	13. personal experience
5. engaging with quotes	14. this system
6. self-reference	15. digressive
7. non-self-reference	16. sharing/discussing pointer
8. presenting alternatives	17. summarizing / paraphrasing author
9. evaluative statement	18. irrelevant

Table 4-1. The codebook used in the content analysis of annotation text.

Each of the codes represents a kind of speech event the annotator can perform. The annotations can vary in a variety of dimensions, from who the intended interlocutor is, to whether it is a declaration or question, and whether the subject is intrinsic or extrinsic to the article. Though the codes are, for the most part, nonexclusive, there are a few notable exceptions. For example, #6 and #7 are opposites, so each annotation is either one or the other. Any annotation that satisfies code #3 by being blank cannot do any of the other things. However, an annotator can do many things in a single annotation, so a single annotation could, for example, both relate a personal experience (#13) and be digressive (#15) or paraphrase the article (#17) at another point in the annotation.

4.3.1 Examples and Analysis

This section consists of a more detailed discussion of each code and its prevalence, as well as several illustrative examples of many of them. The codes are grouped into several categories; annotations unrelated to the underlying text, annotations that describe or refer to the text,

annotations that approach the text critically, and finally annotations that supported or engaged in conversation.

4.3.1.1 Unrelated Annotations

Several annotations' content were completely unrelated to the text, but interesting nonetheless. For example, #3 identified blank annotations. There were only three blank annotations recorded, by three separate users. These can be construed only as mistakes, or as experimenting with the interface. There were several clear attempts to test out the interface (#4) or messages pertaining to the system (#14). For example, ten annotations either contained the word "test" or were clearly for testing out how to post an annotation. It is very common for users to want to experiment with a new system. In wikis, for example, many wikis will have a special page called a "Wiki Sandbox" for new users to become familiar with the functionality; Wikipedia and C2 both have this feature. In addition, four annotations were about Webbed Footnotes itself or other annotation systems. One user wrote: "Isn't this system toooo much like Third Voice?" Another posted about a project he or she was working on and believed was similar. A third user complained that not enough other people were annotating ("no way there is no one here"), and a fourth posted about a bug he found.

4.3.1.2 Reference and Description

The majority of the annotations were in fact directly related to the content of the articles. Summarization, one of the most common reasons for annotation [62,31], was practiced in Webbed Footnotes as well. A good summary distills and often clarifies what an author is intending to express:

```
(1) He is revealing the conflict between what some users want (simple phones) vs.
    what other users want (featureful phones) vs. what the phone manufacturers
    want (customers that buy a new phone every year).
```

Other times, a summary captures not just what the author was saying, but also the reader's reaction to it:

```
(2) All I will remember from this article is that getting high makes you live
    longer. Sweet!
```

Interestingly, summarization and paraphrasing (#17) took place far less than I expected; Only 3 annotations were purely summarization in which the annotator ostensibly attempted to be value-neutral. Many annotations that could not, strictly speaking, be considered summaries would necessarily summarize part of an article's argument, if only to challenge it or offer some other commentary on it. I suspect two reasons that simple summarization did not occur more frequently. The first pertains to the subject matter; these news articles are rather short, and are not documents that the readers likely intend to review again except for conversation purposes. Therefore, the summarization that might be useful in a longer, more complicated text would not be necessary here. This likely extends to the rest of the web, where it has been shown that readers employ skim reading of a text as a strategy to combat information overload [51]. This strategy may be a reaction to the difficulty of reading long texts on a screen [45], which is one reason that usability experts suggest it is good "web writing" style to use fewer words, shorter paragraphs, and the journalistic "inverted pyramid" style [45]. Another reason that is perhaps behind summarization not appearing frequently is that summaries tend to be written for

oneself, for later reference. Since Webbed Footnotes places all annotation in a shared space, people may be less willing to write purely for themselves in a space that is shared with others.

One of the benefits of being able to place annotations arbitrarily on a page is the ability to use deictic references (#12). A deictic reference is one “whose interpretation is relative to the extralinguistic context of the utterance;”²³ that is, it refers – usually with a pronoun – to something that is outside the context of the speaker, but within the context of the discourse the speakers are engaged in. Common examples include the (referential) use of the words “this,” “that” and “it,” for example, “put it on top of that.” The reason deixis is possible in Webbed Footnotes is that, by placing an annotation geographically near other text, the annotation is placed in the linguistic context of that text as well. One of the main reasons deixis is useful is that it saves a great deal of redundancy:

(3) This seems very similar to Judaism’s view.

(4) Actually, it is exactly what is being proposed.

(5) This is a strange type of apology to make.

Comments such as these may be more opportunistic and less deliberate than comments users are forced to post as entire messages at the end of a document, the way they are in weblogs and other environments. They can be more opportunistic because the person writing the comment need to worry about reestablishing the context for the comment. For example, in (3), instead of “this” the user may have had to write something like, “The view the author expresses in the third paragraph about the permissibility of capital punishment. . .” This is much longer, but is necessary if the context relating the comment to its referent is broken. The importance of deixis cannot be overstated; one of the most prevalent kinds of speech coded for, it was seen in just under a quarter of the annotations (23%).

4.3.1.3 Critical Annotations

As I discussed very early in Chapter one, annotating a text requires reading it actively. Annotating while reading helps one learn the material better because it encourages more deliberate reading [62]. I am therefore pleased that many of the annotations in Webbed Footnotes do appear to have approached the content critically, as annotators offered their own opinions about what was discussed in the text. 40% of the annotations contained statements evaluating the actors or actions discussed by the article (#9). Statements that were coded as critical were those that expressed a viewpoint, positive or negative, about the very things the journalists were discussing, but not about the journalists themselves or their writing. It is almost as if annotators were paralleling the journalist, by reading the journalist’s first-hand account and then, like the journalist, offering their own perspective on the account. This is important because it shows that readers are internalizing what they are reading, thinking about it, and forming their own opinions about it.

²³ This definition comes from the linguistics glossary of SIL International, a non-profit linguistics and education association. <http://www.sil.org/linguistics/GlossaryOfLinguisticTerms/WhatIsDeixis.htm>

Evaluative annotations ranged from the very short – even one and two word annotations like “terrible” and “please... ridiculous.” – to somewhat longer, more reflective ones about the qualities of a movie in a movie review:

(6) I loved Best in Show and its parody of the way pet owners identify and are gratified by the “accomplishments” of their pets, really a reflection of themselves. “Mighty Wind” was enjoyable but a bit long.

Annotations like these represent users’ confidence in their opinions and their willingness to express them. Such confidence and the desire to express one’s opinions are the foundations for engaging in conversation and discussion with others. Relatedly, the expression of alternative solutions or explanations to problems discussed in an article (#8) also demonstrates a belief in the value of one’s own ideas. Such ideas were commonly (10%) offered in annotations in the trial:

(7) I love that this stuff gets classified as though it were a disease. WE’RE JUST HAPPY

Many annotations were directly critical of the journalist; 12 annotations (10%) addressed what the users who wrote them considered biases or omissions in the journalist’s writing (#1). This ability to turn a mass communication channel into one that can carry the voices of the consumers as well as the those of the producers, is one of the most important aspects of Webbed Footnotes. As discussed in Chapter Two, this is what the Third Voice founders envisioned their product doing: keeping the media honest:

(8) Why is this put in quotation marks? Not to designate a quote but to doubt it’s [sic] content. If there is reason for doubt: explain, don’t just drop a “hint” like this.

This annotator is responding to the journalist’s use of quotation marks to set apart a phrase about which the journalist is implicitly expressing doubt. The annotator seems to believe that this is a subtle bias and wants to make this clear to other readers. By taking this step to inform other readers, this annotator is enacting the old practice of correcting a text so that future readers will not be, from this annotator’s perspective, misled. If more savvy and thoughtful readers can make points like these clearer to more naïve readers, then the latter can more easily be exposed to new perspectives, making reading a more educational experience for them.

4.3.1.4 Annotation for Conversation

As I have discussed throughout this thesis, annotations are quite often written for the benefit of others and, in shared electronic spaces, can form the basis for conversation. Indeed, in Webbed Footnotes, there is a great deal of evidence that users used annotation to engage one another in conversation. Some annotations were explicitly conversational, in which the annotator invited others to respond (#2):

(9) How is this related to Columbia? Can you guys explain a bit more?

(10) I have never seen one of these games... Anyone ever played?

These invitations to respond can be direct requests for information, as in (9), or questions that the annotator asks the other participants to answer, as in (10). Only four annotations were

coded as containing explicit invitations to respond. However, more annotations (8%) were replies to a specific person (#11). Replies demonstrate interactivity [54], the condition in which messages refer to previous messages. Interactivity is necessary for a group to build a history, for the participants to develop causal relationships between their own behavior and that of others, and for a community and a social structure to form. Therefore, interactivity is a highly desirable trait for an electronic environment to have. In the Webbed Footnotes trial, the replies to other participants show the very beginnings of a group becoming interactive. Though there is a technical function for replying, messages were coded as “reply” if they could be reasonably construed to be addressing a previous, nearby annotation, regardless of whether they hit the “reply” or “post” button to write their annotation.

In being conversational, users most often related a personal experience (#13; 13%) or raised a question related to one raised in the article (#10; 14%). Personal experiences not only help people contextualize and relate personally to the stories they are reading, they have the potential to become a source for additional conversation:

(11) My country has no handguns. Period. We aren't afraid of being shot, and we aren't afraid of not being able to defend ourselves. Our murder rate is an infinitesimal fraction of North America's.

(12) I've always been interested in the social dynamics of digital photography... In my experience, 'Moms' have a terrible time trying to do anything more than the absolute minimum on a camera.

Raising related questions is another way of stimulating additional conversation, and can also be an attempt to fill in gaps in the story where the annotator may believe the journalist neglected to provide enough information to satisfy him:

(13) This is so scary. This is the man that made the war happen. What damage could he do at World Bank?

(14) Did they really try to speed through the check point? If not, why were they shot?

By asking such questions and discussing them amongst themselves, readers pick up where journalists leave off, putting into practice the principle that, though authors have the first word, readers can always have the last word [31] through annotation. Further, the conversations that result from such questions add new information to the article, which enriches its value for subsequent readers.

When readers have the power to annotate, they also have the opportunity to diverge significantly from the topic at hand. Though few (2%) of the annotations were coded as being completely irrelevant (#18), several annotations (7%) were coded as digressive (#15), having begun discussing concepts from the underlying article but diverging to a related topic that itself could not reasonably be construed as being related to the original topics in the article:

(15) I think this is yet another indication of the extent to which the rest of the world perceives American consciouss [sic] to have sunk into nationalist self obsession.

Digression from the original topic is not inherently problematic, however. In Usenet, for example, “thread drift” – the slow shift from one topic of conversation to another in a single

thread – is common and in many cases acceptable. Indeed, users may find a digressive annotation interesting and have an extended conversation about the new topic that it brought up.

One of the most interesting and powerful kinds of alternative or related pieces of content is the pointer to information elsewhere (#16). A few times in the trial (4%), users added references to other sources, either in the form of hyperlinks or in descriptions:

(16) Map of the area: <http://www.topozone.com/map.asp?z=19>

(17) I read a (French) book about this. "Absurd Decisions". I recall the author called that "cognitive bricolage"

Such behavior is also very similar to what Nelson envisioned for his Xanadu system [44], whereby text from anywhere in his information repository could be brought together. Early in Chapter One, I suggested that the hyperlink is powerful only if everyone can use it to bring together disparate pieces of information in support of a larger point. My reasoning is that, when any author composes a text, the author necessarily chooses from among many sources of information those that most strongly support the point he or she is advocating. In that process, the author may link to limitless material to support that viewpoint. Readers, however, cannot generally provide counterarguments or contextualizing information in a persistent way to other readers. When any reader can add a hyperlink to any document, within that document's immediate context there is now the potential for the existence of a link to information that will refute the author's points. Though potentially disconcerting to the author, this check on the author's ability to have complete control of the floor serves readers well because there is greater likelihood that either the author will be pressured to be more honest and balanced in his presentation, or other savvy readers will bring their peers' attention to the fact that the author is not doing so.

4.4 Public and Private Audiences

Given the opportunity, people love expressing their opinions. The most clear outcome from examining users' contributions is that people like sharing their beliefs and ideas and speaking from personal experience. However, they like doing so only when there are others listening. Interestingly, several subjects in the trial were disappointed there were not more people participating. They did not have any knowledge about how many fellow readers were in the system, only about how many annotations they saw; from the number of annotations they generalized about the system's activity level overall. I suspect that if annotation were an intrinsic part of the web, users might feel that there are more eyes on each page, thus increasing their motivation to annotate a page. Though users are generally cautious about what they say and to whom they say it because different groups value and interpret contributions differently [26], they also do not want to waste their time talking into a void. If there is a perception that readership is large, then I suspect annotation, especially of the personal, conversational kind, will grow.

A readership as large as the web itself requires that support for annotation be woven directly into the fabric of the web. Without resorting to massively overhauling the HTTP protocol or addressing the complicated concerns about network trust and security in distributed annotation storage, it is nonetheless clearly possible to make public annotation widespread by integrating

the functionality into the browser – either directly or by a plugin/extension. However, though most public annotation systems like uTok and ThirdVoice did integrate with the browser, they were privately owned by companies whose profit motives were based on promoting the annotation system, leading to the use of advertising, which resulted in a battle for control for screen real estate with other web properties who sought to profit from their own advertising.

Another, more promising alternative is to separate the content storage and delivery from the software that makes it possible, so that many parallel annotation systems can run on a common framework. Users ought to be able to switch among several public or private annotation servers and perhaps use several simultaneously. Annotations need not always be public, and a small and active group may find private annotation extremely useful, as the seminar group demonstrated. Some of the most beneficial interactions people can have electronically are with others with whom they already share close relationships; by analogy, though many wikis are public and truly allow contributions from anyone on the web, many people and groups use private wikis as integral tools in their private collaboration. Likewise, a private annotation server can free a group from the concerns of putting on a public face and allow them to share private information amongst themselves freely, by privately annotating public content.

4.5 Future Directions

As I said at the end of Chapter Three, Webbed Footnotes has many areas in which it could be improved or further developed. At that point, I discussed technical details. Having the benefit of the knowledge gained from observing Webbed Footnotes' use, here I discuss ideas for future changes to the design and their potential impact.

As I discussed toward the end of Chapter One, time and place play an important role in setting the context for an annotation. In light of this, Webbed Footnotes' annotations have a fixed location and lifespans that expire if users no longer find them relevant. However, annotations need not necessarily disappear without a trace.

Once upon a time, Usenet faced an analogous situation; posts would expire after a given time period and were expected to simply disappear forever. However, the keeping of many small archives eventually led to the creation of Deja News and then Google Groups, a persistent, ever-growing repository where all Usenet posts are archived and publicly accessible, perhaps forever. Being permanently archived changed Usenet forever; conversations that once took place in a particular time and context could now be experienced outside of that temporal context and participants had to be concerned with not only how their actions would be interpreted in the present, but also any potential future interpretations by an uncountable number of actors. Likewise, the web itself, which changes every moment, is archived both by Google's search engine and by the Internet Archive. A document's evolution, captured by these archivists, becomes frozen in time at each stage, tracing out the document's history.

If Webbed Footnotes' annotations could be experienced out of context, I believe it is vital that they always maintain their links back to the documents in which they originally lived. Unlike Usenet conversations, for example, annotation conversations cannot stand on their own. For this reason, integration with a perpetual web archive like the Internet Archive makes sense. So long as the timestamps of the annotations and of the web documents are maintained, it ought to be possible to situate an annotation in its original context in the archive. Though Webbed

Footnotes participants may, like Usenet participants, have to face the unfortunate reality of their words persisting forever, there may be some comfort in the knowledge that that conversation will be experienced in the way it was intended, so that the social and cultural milieu of the participants can be appreciated by future observers.

Besides Usenet, the other environments that influenced Webbed Footnotes were the weblog and the wiki. Personal and collective publishing with these tools is receiving a great deal of attention; Webbed Footnotes can learn from them, and vice-versa.

The wiki is very good at allowing many people to contribute to a document, but makes negotiation of those contributions challenging. The problem is so acute in Wikipedia that it necessitates having a “talk” page that parallels each entry page for discussing the content of that entry. In contrast, Webbed Footnotes is very good at supporting discussion on top of documents, but currently provides no functionality for incorporating the discussion into the body of the document, so that if many people were to use Webbed Footnotes to mark up a document they are revising, they would find it challenging to finalize the changes they would like to make. Bridging this gap is an important future step for wikis as well as for Webbed Footnotes.

Weblogs, in that their user community has developed a tradition of heavily quoting from outside sources, exhibit some traits of Nelson’s “xanalogical” structure [44]. However, their links are still one-directional; the weblog points to the web content it is discussing, yet the content cannot point back to weblogs that are discussing it. Within the world of weblogs, posts can contain *trackback* links, or links to other weblogs that discuss that post; trackbacks function as a bidirectional link between content and other locations that cite that content. In contrast, the web at large does not allow pointed-to content to refer back to that which cites it. There are some web services that aggregate weblog content by topic of discussion, but one must visit a specific aggregator.²⁴

Annotation could potentially serve as a technique for filling this gap and provide a web-wide trackback-like feature. Consider Webbed Footnotes’ User Page, discussed in Chapter Three. This page consists of a chronologically-ordered list of all the annotations a user authored; in other words, a weblog. In the previous section I discussed the difficulty of achieving pervasiveness. However, a pervasive annotation system that makes the individual user’s page somewhat more prominent would effectively constitute a bidirectional weblog. Though weblogs are hardly starved for attention,²⁵ the bidirectional links of annotations would alert readers to the fact that, elsewhere on the web, they could be joining in a discussion of the content they are presently reading.

4.6 Conclusion

In this thesis I have presented Webbed Footnotes, a tool for annotation-based conversation within web documents. Like other annotation systems that have come before, Webbed Footnotes’ foundations lay in traditional paper document annotation, a practice hundreds of

²⁴ For example, Blogdex (<http://blogdex.media.mit.edu>) and Technorati (<http://www.technorati.com>).

²⁵ In the aggregate, that is. Clearly, most webbloggers “toil in obscurity” and many of them would be grateful for increased readership.

years old in which readers actively mark up their texts to distill and augment knowledge both for their own use and for the use of others. At the same time, Webbed Footnotes was inspired by other modern conversational and collaborative tools, ranging from the conversational ur-space of Usenet, to more recent web tools like the weblog and the wiki, to hypothetical tools like Bush's memex or Nelson's Xanadu.

The design of Webbed Footnotes was planned so as to achieve the often-competing goals of making annotations prominent yet ensuring that web documents remain legible once annotated. To do this, Webbed Footnotes makes use of a two part annotation consisting of a small annotation marker and a large, dynamically-displayed footnote window. Webbed Footnotes sought to employ a moderation system that did not burden users with a managerial chore, yet functioned as intended even when users did not moderate. Approval-Based Moderation, in which users could vet or approve others' annotations, would allow more highly approved annotations to "live" longer, and "expire" annotations that users did not find worthy of vetting.

In two user trials, one large and one small, participants annotated web documents and engaged each other in conversation while critically examining the texts they were reading together. In the small trial, users reported feeling as though they were reading as part of a group, and in the large trial participants often were critical of the journalists whose articles they were reading, offering their own critiques, questions and explanations, and occasional references to other web material. These results demonstrate that, given the opportunity, users may become more active and savvy readers, potentially transforming the mass media content of the web into a shared space for conversation, with the possibility of shared learning.

References

- 1 Adler, Annette, Anuj Gujar, Beverly L. Harrison, Kenton O'Hara and Abigail Sellen. (1998) "A Diary Study of Work-Related Reading: Design Implications for Digital Reading Devices." *Proceedings of the ACM Conference on Human Factors in Computer Systems (CHI)*.
- 2 Avery, Chris, Paul Resnick and Richard Zeckhauser. (1999) "The Market for Evaluations." *American Economic Review*. American Economic Association. 89(3). 564-584.
- 3 Bandura, Albert. (1976) *Social Learning Theory*. Prentice Hall.
- 4 Barabasi, Albert-Laszlo. (2002) *Linked*. New York: Penguin.
- 5 Bechar-Israeli, Haya. (1999) "From Bonehead to Clonehead." *Journal of Computer-Mediated Communication*. 1(2).
- 6 Bikhchandani, Sushil and Sunil Sharma. (2001) "Herd Behavior in Financial Markets." *IMF Staff Papers*. Vol. 47. No. 3.
- 7 Bouvin, Neils Olof. (1999). "Unifying Strategies for Web Augmentation." *Proceedings of the 10th ACM Hypertext 1999 Conference*. Darmstadt, Germany.
- 8 Bray, Hiawatha. (2004) "One Great Source – If You Can Trust It." *Boston Globe*. July 12, 2004.
- 9 Brush, A.J. Bernheim, David Barger, Anoop Gupta and JJ Cadiz. (2001) "Robust Annotation Positioning in Digital Documents." *Proceedings of the ACM Conference on Human Factors in Computer Systems (CHI)*.
- 10 Bush, Vannevar. (1945) "As We May Think." *The Atlantic Monthly*. July 1945.
- 11 Cadiz, JJ, Anoop Gupta and Jonathan Grudin. (2000) "Using Web Annotations for Asynchronous Collaboration Around Documents." *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*.
- 12 Castleman, Craig. (1982) *Getting Up: Subway Graffiti in New York*. Cambridge: MIT Press.
- 13 Daft, Richard L. and Robert H. Lengel. (1986) "Organizational Information Requirements, Media Richness and Structural Design." *Management Science*. Vol. 32, No. 5.
- 14 Deleuze, Gilles. (1992) "Postscript on the Societies of Control." *OCTOBER* 59. Cambridge: MIT Press. pp. 3-7.
- 15 Donath, Judith. (1998) "Identity and Deception in the Virtual Community." *Communities in Cyberspace*, Marc Smith and Peter Kollock, eds. Routledge.
- 16 Donath, Judith, Karrie Karahalios and Fernanda Viégas. (1999) "Visualizing Conversations." *Proceedings of HICSS-32*.
- 17 Donath, Judith and Niel Robertson. (1994) "The Sociable Web." *Proceedings of the Second International WWW Conference*. Chicago, IL. October 1994.
- 18 Dourish, Paul and Victoria Bellotti. (1992) "Awareness and Coordination in Shared Workspaces." *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*.
- 19 Erickson, Thomas. (2003) "Designing Visualizations of Social Activity: Six Claims." *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*.
- 20 Fiske, John. (1989) *Understanding Popular Culture*. New York: Routledge.

- 21 Fogg, B.J. (2002) *Persuasive Technology: Using Computers to Change What We Think and Do*. Morgan Kaufman.
- 22 Glaser, Barney G. and Anselm L. Strauss. (1967) *Discovery of Grounded Theory*. Aldine.
- 23 Goffman, Erving. (1959) *The Presentation of Self in Everyday Life*. Doubleday.
- 24 Goffman, Erving. (1969) "Strategic Interaction." *Strategic Interaction*. Philadelphia: University of Pennsylvania Press.
- 25 Golder, Scott A. and Judith Donath. (2004) "Hiding and Revealing in Online Poker Games." *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*.
- 26 Golder, Scott A. and Judith Donath. (2004) "Social Roles in Electronic Communities." Presented at *Internet Research 5.0*. Sussex, England. September 19-22, 2004.
- 27 Hall, Kira. (1996) "Cyberfeminism." In *Computer Mediated Communication: Linguistic, Social and Cross-Cultural Perspectives*. Ed. Susan Herring. Amsterdam: John Benjamins.
- 28 Hauben, Michael and Ronda Hauben. (1997) *Netizens: On the History and Impact of Usenet and the Internet*. Los Alamitos, CA: The IEEE Computer Society. Also available online at <http://www.columbia.edu/~hauben/netbook>
- 29 Herring, Susan. (1999) "Interactional Coherence in CMC." *Journal of Computer Mediated Communication*. 4(4).
- 30 Hill, William C. and James D. Hollan. (1992) "Edit Wear and Read Wear." *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.
- 31 Jackson, H.J. (2001) *Marginalia*. Yale University Press.
- 32 Jacobson, David. (1999) "Impression Formation in Cyberspace: Online Expectations and Offline Experiences in Text-Based Virtual Communities." *Journal of Computer Mediated Communication*. 5(1).
- 33 Kristeva, Julia. (1980) *Desire in Language: A Semiotic Approach to Literature and Art*. Columbia University Press.
- 34 Knowledge@Wharton. "Can Firefox outfox IE?" Special to Cnet News.com. March 14, 2005.
- 35 Kollock, Peter. (1999) "The Production of Trust in Online Markets." In *Advances in Group Processes*. Vol. 16. Eds. E.J. Lawler, M. Macy, S. Thyne and H.A. Walker. Greenwich CT: JAI Press.
- 36 Kollock, Peter and Marc Smith. (1996) "Managing the Virtual Commons." In *Computer Mediated Communication: Linguistic, Social and Cross-Cultural Perspectives*. Ed. Susan Herring. Amsterdam: John Benjamins.
- 37 Kumar, Aparna. (2001) "Third Voice Trails Off." *Wired News*. April 4, 2001.
- 38 Lampe, Cliff and Paul Resnick. (2004) "Slash(dot) and Burn: Distributed Moderation in a Large Online Conversation Space." In *Proceedings of Conference on Human Factors in Computing Systems (CHI)*.
- 39 Lenhart, Amanda, John Horrigan and Deborah Fallows. (2004) "Content Creation Online." Pew Internet and American Life Project.
- 40 McHenry, Robert. (2004) "The Faith Based Encyclopedia." Tech Central Station. November 15, 2004. <http://www.techcentralstation.com/111504A.html>

- 41 Minerals Management Service, U.S. Department of the Interior. "Rigs-to-Reefs Information." <http://www.gomr.mms.gov/homepg/regulate/environ/rigs-to-reefs/information.html>
- 42 Negroponte, Nicholas. (1995) *Being Digital*. Knopf.
- 43 Nelson, Theodor Holm. (1974) *Computer Lib / Dream Machines*.
- 44 Nelson, Theodor Holm. (1999) "Xanalogical structure, needed now more than ever: parallel documents, deep links to content, deep versioning and deep reuse." *ACM Computing Surveys*. Issue 4es. December 1999.
- 45 Nielsen, Jakob. "How Users Read on the Web." Alertbox. October 1, 1997. <http://www.useit.com/alertbox/9710a.html>
- 46 Nonnecke, Blair, and Jenny Preece. (2001) "Why Lurkers Lurk." Americas Conference on Information Systems.
- 47 Nonnecke, Blair and Jenny Preece. (2000) "Lurker Demographics: Counting the Silent." *Proceedings of the ACM Conference on Human Factors in Computer Systems (CHI)*.
- 48 Norman, Don. (1988) *The Design of Everyday Things*. Basic Books.
- 49 Nussbaum, Emily. (2004) "My So-Called Blog." *New York Times Magazine*. January 11, 2004.
- 50 Oakes, Chris. "Third Voice Rips Holes in Web." *Wired News*. July 9, 1999.
- 51 Payne, Stephen J. (2003) "Cognitive Strategies and Design Heuristics for Dealing with Information Overload on the Web." IGR Report GR/M43302/01.
- 52 Petroski, Henry. (1999) *The Book on the Bookshelf*. Alfred A. Knopf.
- 53 Popkin, Samuel L. and Michael P. McDonald. (2000) *Washington Post*. November 5, 2000.
- 54 Rafaeli, Sheizaf. (1988) "Interactivity: From new media to communication." *Sage Annual Review of Communication Research: Advancing Communication Science* Vol. 16. Beverly Hills, CA: Sage.
- 55 Resnick, Paul. Unpublished wiki conversation. April 3, 2005. <http://www.socialcomputingresearch.net/workshop/Golder2005/>
- 56 Resnick, Paul, Neophytos Iacovou, Mitesh Suchak, Peter Bergstrom and John Reidl. (1994) "GroupLens: An Open Architecture for Collaborative Filtering of Netnews." *Proceedings of the ACM Conference on Computer Supported Cooperative Work*.
- 57 Roscheisen, Martin, Christian Mogensen and Terry Winograd. "Shared Web Annotations As A Platform for Third-Party Value-Added Information Providers: Architecture, Protocols, and Usage Examples." *Technical Report CSDTR/DLTR*. Stanford University Department of Computer Science.
- 58 Sellen, Abigail and Richard H. R. Harper. (2002) *The Myth of the Paperless Office*. Cambridge: MIT Press.
- 59 Smith, Marc. (2001) *Netscan: A Tool for Measuring and Mapping Social Cyberspaces*. Available online at <http://netscan.research.microsoft.com>.
- 60 Spiegel, Dana. (2001) "Coterie: A Visualization of the Conversational Dynamics Within IRC." Unpublished MIT Media Laboratory Master's Thesis.

- 61 Sproull, Lee and Sara Kiesler. (1986) "Reducing Social Context Cues: Electronic Mail in Organizational Communication." *Management Science*. Vol. 32, No. 11.
- 62 Starr, Raymond. (1983) "Teaching Study Skills in History Courses." *The History Teacher*. Vol. 16, No. 4.
- 63 Takahashi, Masamichi. (2005) "The Roles and Effects of Active Lurking in In-House Online Communities." *Proceedings of HCII 2005*.
- 64 Turner, Teresa. "The Coral Reef Ecology Homepage." <http://www.uvi.edu/coral.reefer/>
- 65 Viégas, Fernanda, Martin Wattenberg and Kushal Dave. (2004) "Studying Cooperation and Conflict between Authors with history flow Visualizations" *Proceedings of the ACM Conference on Human Factors in Computer Systems*.
- 66 Walther, Joseph B. "Computer Mediated Communication: Impersonal, Interpersonal and Hyperpersonal Interaction. *Communication Research*. Vol. 23, No. 1.
- 67 Walther, Joseph B. and Malcolm R. Parks. (2002) "Cues Filtered Out, Cues Filtered In: Computer-Mediated Communication and Relationships" In M.L. Knapp and J.A. Daly (eds) *Handbook of Interpersonal Communication*. 3rd ed. Thousand Oaks, CA: Sage.
- 68 Wexelblat, Alan and Pattie Maes. (1999) "Footprints: History-Rich Tools for Information Foraging." *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.
- 69 Whittaker, Steve, Loren Terveen, Will Hill, Lynn Cherny. (1998) "The Dynamics of Mass Interaction." *Proceedings of the ACM Conference on Computer Supported Cooperative Work*. 257-264.
- 70 Whyte, William H. (1988) *City: Rediscovering the Center*. New York: Doubleday.
- 71 Witmer, Diane F. (1997) "Risky Business: Why People Feel Safe in Sexually Explicit On-Line Communications." *Journal of Computer-Mediated Communication*. Vol. 2 Issue 4.
- 72 Wolf, Gary. (1995) "The Curse of Xanadu." *Wired*. 3.06. June 1995.
- 73 Yee, Ka-Ping. (2002) "Zest: Discussion Mapping for Mailing Lists." *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*.
- 74 Yee, Ka-Ping (2002) "CritLink: Advanced Hyperlinks Enable Public Annotation on the Web." *Demonstration. Proceedings of the ACM Conference on Computer-Supported Cooperative Work*.
- 75 Zuckerman, Orit. "uTok". (2005) Orit Zuckerman's personal homepage. <http://web.media.mit.edu/~orit/utok.html>


```
ERROR: undefined
OFFENDING COMMAND:

STACK:
```