

What Does A Very Large-Scale Conversation Look Like?

Warren Sack
UC Berkeley
SIMS
314 South Hall
Berkeley, CA 94720-4600
+1.510.643.3904
sack@sims.berkeley.edu

Introduction

The new electronic spaces that I am interested in have the following characteristics in common:

- *They are large.* Many server sites now support interchanges between hundreds and thousands of people. Usenet newsgroups and large listservs are the most common of such sites. I call these usually text-based, usually asynchronous interchanges *very large-scale conversations* (Sack, 2000a).
- *They are network-based.* More specifically, they support network-based communities. The boundaries of these spaces and the communities they support are not geographic boundaries. Communities of artists, writers, and scientists are examples of pre-Internet, network-based communities; i.e., communities based upon a social network and some shared interests or needs. Network-based communities are of a different kind than geographically-based communities like neighborhoods, cities, and nations. Network-based communities -- e.g., the scientific community -- have continued to grow with the help of new network technologies, but contemporary technologies have also engendered a variety of new communities; e.g., the open source community.
- *They are public.* As more and more people gain access to the Internet from their home or school rather than from their workplace the Internet increasingly becomes a space for public discussion and exchange. Very large-scale conversations are a common event within the confines of large industry; e.g., the huge number of communications between thousands of people necessary to design and build an airplane or coordinate the production of a film. However, these have a distinctly different character than the very large-scale conversations in which people are participating as individuals rather than as employees. The Internet is engendering the production of new public spaces that may offer the means to reinvigorate public discourse (cf., Tsagarousianou et al., 1998; Hague and Loader, 1999).

From the perspective of the history of media, very large-scale conversation (VLSC) is a new and mostly unexplored phenomenon. At no other point in history have we had a medium that supports many-to-many communications between hundreds or thousands of people. VLSC takes place across international borders, often on a daily or hourly basis. Unlike in older media (e.g., telephones) participants in these very large-scale conversations usually do not know the addresses of the others before the start of a conversation. Current social scientific theories and tools we have for understanding and investigating conversations and discourse include those of discourse analysis (e.g., Schiffrin, 1994) and conversation analysis (e.g., Hutchby and Wooffitt, 1998). These existing theories and techniques can handle the analysis of small-scale conversations; e.g., interactions between thirty or fewer people. But, it is not obvious how the existing methods can be scaled up to handle the huge, many-to-many interactions that have now become commonplace on the Internet. So the challenge is this: What software can be designed to help us navigate the new public spaces of VLSC?

Navigation

Michel Foucault has pointed out that "the comparison between medicine and navigation is a very traditional one" (Foucault, 1983). Medicine, navigation, and government have to do with self-guidance, control, and governance. Etymologically, the verb "navigate" comes from the combination of words *navis* [ship] and *agere* [to guide]. Thus, in the case of the navigation of a large, public, information space the "ship" has been replaced by the self, and so the point of navigation is self-guidance or self-governance. From this perspective, the right way to evaluate or critique a browser -- or any other piece of navigation

software -- is with respect to how well it supports self-governance. In the particular case of a VLSC browser it should help us better understand where we are located (and where we might go) in a wider network of social and semantic relations. It should also help us consider the existence of a collective, self-organization constructed through the text and talk of a VLSC. I am interested in the larger ethical and aesthetic implications of this understanding of navigation.

To better understand the issues of designing software for navigation, I've borrowed a conceptual framework from Paul Dourish and Matthew Chalmers. In 1994 they wrote a paper entitled "Running Out of Space: Models of Information Navigation" (Dourish and Chalmers, 1994). In their paper, Dourish and Chalmers assert that there are at least three ways in which large bodies of information can be navigated: (1) social, (2) semantic, and (3) spatial navigation.

- *Social Navigation* : Firstly, Dourish and Chalmers claim that software can be designed to support the social navigation of information (see Munro, Hook and Benyon, 1999). By social navigation I understand them to mean people helping other people to find information. Examples of social navigation software include the mechanisms employed in recommender systems and collaborative filtering (cf., Resnick and Varian, 1997). Work done in organizing texts through citation analysis -- as is done in the field of science studies -- can also be counted as support for social navigation (e.g., Garfield, 1979).
- *Semantic Navigation* : Semantic navigation requires, for instance, the sorts of computation we have available to us when we use a search engine on the web. Using techniques from information retrieval and computational linguistics, semantic navigation can be supported through the calculation of some approximation to the meaning of a set of documents.
- *Spatial Navigation* : Spatial navigation entails the kind of manipulations often performed in the area of information visualization to convert a large body of data into a two- or three dimensional image. The image then can function as an interface to the information that it incorporates (see, for example, Card, Mackinlay and Shneiderman, 1999).

To support all three of these types of information navigation, I use some techniques and tools from sociology to support social navigation; some ideas from linguistics to support semantic navigation; and, some aspects of graphical interface design to support spatial navigation of VLSCs. A more complete description of my approach can be found elsewhere (e.g., Sack, 2000b). I have designed and implemented a piece of software to embody this approach. The Conversation Map system is a prototype, VLSC browser.

Conversation Map

The Conversation Map system can analyze several thousand messages at a time. It employs a set of computational linguistics and sociology techniques in order to generate a graphical summary of the messages. The graphical summary includes

- a set of *social networks* that illustrates who is corresponding with whom;
- a menu of *themes* of discussion that are important to the conversation embodied in the messages; and,
- a *semantic network* that articulates some of the emergent synonyms or metaphors of the discussion.

One can use the Conversation Map like Netscape Messenger, Outlook, Eudora, or any other conventional news or mail reader. However, right now, the text analysis procedures are too slow. An analysis of several thousand messages currently takes the system several hours. I am re-engineering the system (and redesigning the interface) to allow one to use the Conversation Map as an everyday email reader or news browser.

Social Networks: The upper left quadrant of the interface depicts a set of social networks that record who is corresponding with whom. By "corresponding" I mean who is mutually responding to and/or quoting from whom. According to my definition, two participants -- say "Sally" and "Spot" -- correspond with one another if Sally posts to the newsgroup, Spot responds to (or cites) Sally's message and then, later in the discussion, Spot posts to the group and Sally responds to (or quotes from) Spot's message. In the social network, Sally and Spot will be represented as two nodes with a line connecting them. If they correspond

frequently, then the line between them will be short. In contrast, those pairs of participants who correspond only once will be plotted relatively far apart. Note that posters who spam the group with many messages, but who receive no replies, do not even show up on the graph. Those participants who show up closely connected are pushed to the middle of the graph and can be understood as virtual mediators of the newsgroup. They are virtual moderators because most of the analyses I have done have been of unmoderated, public discussion spaces on the Net. To end up in such a position one needs -- not only to post many messages -- but also to have others in the group reply to or quote from many of one's messages. So, the social network display acts both as a filter for spammers and a means to identify some of the main players in a discussion.

Themes: The menu in the upper middle of the interface lists the themes of the conversation. Imagine that Sally posts a message about football, and then Spot responds with a message that includes some reference to baseball. Then, perhaps later in the discussion, Spot posts a message about skiing and Sally responds with one concerning skating. This correspondence will be represented in the social network, but some approximation to the theme of their exchange will also be listed in the middle menu. In this case, since football, baseball, skiing, and skating are all sports, the term "sports" might be listed on the menu of themes. Calculating that these four terms are all sports requires, of course, a machine-readable thesaurus. The thesaurus employed in the Conversation Map system is WordNet, a lexical resource created by George Miller, his colleagues, and students at Princeton University (see Fellbaum, 1998). The algorithm for calculating the multi-authored themes is akin to (but not exactly the same as) a set of procedures from computational linguistics designed to analyze the lexical cohesion of single-authored texts (cf., Hirst and St-Onge, 1998).

Semantic Network: The calculations performed to create the semantic network shown in the upper right-hand corner do not use a thesaurus, but, rather, *automatically generate* a rough-draft thesaurus. To create a rough-draft thesaurus the Conversation Map system does the following: First, the content of all of the messages exchanged during the conversation is parsed. In other words, the subjects, verbs, objects and some of the other modifying relations are identified between the words of each sentence in the texts of the messages. Next, for each unique noun mentioned in the corpus of messages a profile is built. By "profile" I mean that, for each noun, a vector is created that records all of the verbs for which the noun functioned as a subject; all of the verbs for which the noun functioned as an object; all of the adjectives which modified the noun; etc. Once a profile has been calculated for each noun, the nouns' profiles are compared to one another and each noun's nearest neighbor is identified. An algorithm described in Grefenstette (1994) is used to calculate and compare the noun profiles. If two nouns have similar profiles, then they can be said to have been "talked about" in similar ways by the participants in the discussion. Therefore they may be considered synonyms or possibly metaphors for one another. In the semantic network, if two nouns are nearest neighbors, then they are plotted as two nodes connected to one another.

Why is this sort of analysis of interest for the navigation of very large-scale conversations? To answer this question, I compare this sort of analysis with some work done by the cognitive scientists George Lakoff and Mark Johnson. Lakoff and Johnson wrote a book entitled *Metaphors We Live By* (Lakoff and Johnson, 1980). The book is filled with a set of metaphors that Lakoff and Johnson claim are central to our (presumably North American, English-speaking) culture. In their book, for instance, they claim that one emergent metaphor of our culture is that *arguments are buildings*. As part of their method to argue for the validity of insights like this, they show how two nouns, which might *a priori* be considered to be completely unlike one another, show up in very similar contexts. For example, one can say "The building is shaky" but one can also say "The argument is shaky." One can say "The building collapsed" but also "The argument collapsed." Similarly, both buildings and arguments can be said to have "foundations," "to stand," and "to fall," "to be constructed," "to be supported," "to be buttressed," etc. A set of similar sentences of this sort provides an empirical means for thinking about and discovering how synonyms and metaphors are produced over the course of a large amount of discussion.

Thus, this tool for automatic, rough-draft thesaurus generation gives one the means to begin to generate the sorts of hypotheses that Lakoff and Johnson explore in their book. Alternatively, one can understand the noun profiles and semantic networks in Michel Foucault's terms; namely as "statements" and

"diagrams" respectively. Gilles Deleuze explains Foucault's terms (Deleuze, 1988). So, the Conversation Map gives one some data exploration/navigation tools to start to understand how conversations differ from one another according to the metaphors, synonyms and "statements" that are produced by the collective efforts of their participants.

Message Archive: The lower half of the interface is a graphical representation of all of the messages that have been analyzed by the Conversation Map system. Messages are organized into threads. A thread is defined as an initial post, all of the responses to the initial post, all of the responses to responses, etc. The threads are plotted like spider webs. The first message posted is represented as a large node, the responses, responses to responses, etc. are plotted as radiating out from the center. Double-clicking on a message thread in the lower half of the interface, will cause a larger picture of the thread to be displayed.

The lower half of the screen is divided into a grid and the threads are organized in chronological order from upper left to lower right. If a thread contains many messages, it shows up as an almost completely green square on this display. If a thread contains few messages, then it shows up as an almost completely black square. So, scanning across from upper-left to lower-right the lower-half of the screen can be seen as a rough guide to the posting activity in the newsgroup.

Examples

In the following section I show twelve example Conversation Maps that were generated from a wide variety of online, public discussions. With these examples, I hope the semiotics of how to read these maps will become understandable. Also I hope that these one page, graphical summary of hundreds or thousands of email messages will be seen to be a useful thing for gaining a quick glimpse into a very large-scale conversation.

Politics

The map on the left and the map on the right were created about a week apart using messages from the newsgroup alt.politics elections. The one on the left was generated immediately before the presidential election. Notice how the main themes of discussion center around the candidates: Gore, Bush, and Nader. A week after the election the conversation has moved away from a discussion of the candidates. Now it is a discussion of the technicalities of elections: votes, counts, ballots, laws and courts are the newly prominent themes of discussion. This can be seen in the themes and semantic network of the map on the right.

Media

This pair of maps show the same newsgroup (a discussion about the television show X-Files) at two different times. Notice how many themes of discussion there are in the map on the left. Now, notice how very few themes of discussion are listed in the map on the right. Because the Conversation Map uses a very generous means of counting the themes of discussion, it usually lists too many, not too few. What the map on the right tells us is that no one is following up on what other people are saying. The two snapshots in time represented by these two maps demonstrate how an online discussion can change from being one where people talk to one another into one where they just talk at one another. This fact is also represented in the very scattered appearance of the social network. For a detailed analysis of these types of maps as a new practice of media audience studies see Sack (1999).

Environment

The map on the left represents about a month's worth of messages posted to the group sci.environment. The map on the right represents the same newsgroup one month later. By comparing the two maps you can get some idea of how the group has changed over time. One thing that has remained stable between the two maps is the connection in the semantic networks between the terms "people" and "problem." This is a clue that perhaps, in this newsgroup, people are seen to be one of the causes of environmental problems. But a hypothesis like this that one can come up with by looking at the maps needs further investigation to be confirmed or discarded.

Education

On the left is a map of about 300 messages from the Usenet newsgroup misc.education. Note the themes of discussion and compare them to the map on the right. Both maps summarize discussions about education and learning. The map on the right summarizes a semester's worth of messages posted by a distance learning course taught by Prof. Linda Polin of Pepperdine University. In comparison with the first map, note how much more tightly knit the social network is here: people are responding to one another. Note also the elaborate threads containing many messages as compared to the sparse threads in the first map. These elaborate thread structures show that the participants are repeatedly elaborating on one another's posting. This sort of an exchange is perhaps much deeper than, for example, the quick question-and-answer format of the technology discussions depicted below and the curt exchanges that one can note in the threads of the political discussions above.

Technology

The conversation map on the left was created from about a month's worth of messages posted to a public listserv devoted to the construction of Lego robots. Note how the social network shows that there are a multiple hubs: these correspond to an expert in mechanical systems, an expert in programming, and expert in electronics. The second map is an analysis of about 2500 messages from the newsgroup devoted to the Perl programming language: comp.lang.perl.misc. Note the dense social network and also compare the thread pattern here with the deep discussion of education analyzed above. The pattern here is indicative of a series of brief question-answer clusters. In contrast, the elaborate threads in the deep education conversation indicate that participants are repeatedly elaborating on one another's responses.

Health

The conversation depicted here, on the left, took place in a public newsgroup devoted to attention deficit disorder. However, one can see from the map that the discussion was not just about the illness, but about family members as well. The map on the right is summary of several hundred messages sent to a newsgroup on chronic fatigue syndrome. As can be seen here too, the discussion focuses – not just on the illness – but on a more general discussion of people and citizenry. The anthropologist of science, Prof. Joseph Dumit of MIT, argues that illness like these (ADD and CFS) are illnesses one has to "fight to get" because they are often not recognized by doctors and insurance companies. Consequently, online discussions can become places where sufferers can meet and illness-based social movements can emerge (see Sack and Dumit, 1999).

Conclusions: Mapping Common Ground

In a recent essay the writer and frequent contributor to *Artforum*, Frances Richard, discusses the Conversation Map in juxtaposition with other work on mapping conversations, specifically the work of artists Janet Cohen, Keith Frank, Jon Ippolito, and Mark Lombardi:

Sack's project unfolds a tiered grid on which this collective polemic can be tracked. Electronic communication is often theorized in terms of a return to epistolary or conversational consciousness, and the opportunity of discussing, say, the Kosovo situation with political scientists, Balkan historians, NATO-watchers, Albanian teachers, and Serb journalists represents a previously unimaginable crucible for spontaneous intercultural and interdisciplinary debate. The existence of such a collective is so fascinating that the interface seems transparently beneficial, a labor-saving device without which important knowledge would smear into static. Sack's high-tech browser and Mark Lombardi's painstaking low-tech works on paper thus perform similar procedures on the information glut, but their interventions point to opposite feelings about that information. The group Lombardi examines is a suspect elite, and conversations are presumed to be exploitative and self-serving, ripe for the whistle blower. Newsgroup and chat-room speech, in contrast, is imagined as vox populi in action. The VLSC map does not expose a closed coterie; it expands an egalitarian fellowship. (Richard, 2001, p. 81)

The specific conversation map Richard considers is the one that appears at the beginning of this paper. This image was produced by the Conversation Map system from an analysis of over 1200 messages from the Usenet newsgroup soc.culture.albanian, a group devoted to the discussion of Albanian culture in general, but at this period in time (16 April 1999 - 4 May 1999) especially the war in Kosovo. One can see from the social network that the discussion was rather cohesive and dominated by a few central

voices. But, it is the automatically generated semantic network that illustrates the optimistic politics that motivates this project: the hope for a truly global conversation (cf., UNDP, 1999, pp. 62-63). The upper right-hand corner of the semantic network connects the terms "nation," "state," and "country." These are associations that one might find in a conventional thesaurus and simply show that the empirical procedure for automatically compiling a rough-draft thesaurus from a corpus of messages is working. A closer examination of the automatically compiled thesaurus reveals many of these conventional associations.

The lower left-hand corner of the semantic network contains a cluster of entities that all represent a category one might label as "political or military entity." These include the KLA, the US, Macedonia, Russia, and Croatia. Closely connected to these is an association that is at first surprising but, upon reflection, not so surprising: Clinton is connected to Milosevic. Why? Because, as described in the contents of the messages, both Clinton and Milosevic are acting as (and thus associated with the verbs and adjectives descriptive of) a president of a country.

But, it is neither the upper right-hand corner's reproduction of conventional associations, nor the lower left-hand corner cluster of political actors that is of especial interest here. Rather it is the central portion of the semantic network -- connecting "Albanians" to "Serbs" through "people" -- that instantiates what might be seen as an implicit hope or goal of the conversation: namely, to understand Albanians and Serbs as comparable and equal. One optimistic way of reading the semantic network computed by the Conversation Map system for the soc.culture.albanian group is this: "people" is a neutral term: Serbs, Albanians, refugees, countries, and governments are all "talked about" like people. This is, perhaps, a sort of thin humanism: "after all we are all people."

However, it must be kept in mind that no one in the newsgroup necessarily wrote "we're all people." In fact, the comparison is much more subtle. The neutral term is not necessarily "people" *per se*, but rather attributes (e.g., adjectives and verbs) that may be applicable to everyone (Serbs, Albanians, or people in general) on all sides of the argument. These overlaps, these neutral attributes, can be seen by examining the profiles for the terms in the semantic network. This is done using a part of the Conversation Map interface not discussed in this paper.

If we focus on only those verbs for which both "people" and "Serbs" appeared as a subject of the verb, then the resultant overlapping list looks like this:
SERBS ARE PEOPLE (terms appear as subject for each of the verbs one or more times) allow, be, destroy, die, do, drive, exist, flee, get, give, have, keep, know, lay, leave, live, make, need, pay, remember, tell, think, turn

In other words, by looking at the archive of messages one can find many places where, for instance, both "people" and "Serbs" appear as subjects of the same verb. From the intersected lists of verbs one can see that, in the archive of soc.culture.albanian messages "Serbs" and "people" are discussed in similar terms because there exist one or more statements in the archive for both "Serbs" and "people" where they are describe separately as agents which allow, destroy, die, do, drive, exist, etc.

The verb "to need" is one of these shared verbs found in the intersection of the "Serbs" and "people" profiles. Clicking on a verb in the intersected profiles (not shown here, but displayed by the Conversation Map interface when two terms in the semantic network are selected) reveals the following two example sentences that partially underpin the link between "Serbs" and "people" in the semantic network: (1) "You have to realize that Greeks and Serbs need a just solution, and not just Serbia has a solution" (2) "It is not enough to be alive, people need normal life."

Similar word association lists are computed on-demand by the Conversation Map system for any other pair of terms in the semantic network and, if desired, example sentences of the terms in use can also be viewed.

ALBANIANS ARE PEOPLE (terms appear as subject for each of the verbs one or more times) cross, displace, do, flee, have, hate, hide, leave, lose, say, suffer, think, walk

SERBS ARE ALBANIANS (terms appear as subject for each of the verbs one or more times) do, flee, found, have, insist, leave, shoot, think, want

These sorts of verbal overlaps designate possible common ground and thus potential insights into where and how to start a discussion that all sides of the argument might listen to or participate in. In other words, the associations shown in the semantic network do not document an accomplished humanism, but rather empirically point back into the discussion to show places in the conversation where one might return to build a common ground because Serbs and Albanians are all actants who flee, think, and want. Thus, we can see these associations as both an empirical fact documented in the archive of the conversation and, simultaneously, as a set of potential goals for future discussion.

Closer examination of this particular conversation reveals that the participants used a variety of languages (English, but also the languages of the region including Croatian, Macedonian, etc.) and pursued the discussion in a highly combative, argumentative style. From a philosophical perspective it is extremely hard to understand this exchange as a *dialectic* in which the many sides might eventually reach a compromise or synthesis (e.g., the classical pro + con ==> compromise; or hypothesis + antithesis ==> synthesis).

This conversation illustrates a non-dialectical exchange in which, potentially, no common ground might ever be accomplished. In the words of the philosopher Jean-Francois Lyotard, this conversation may very well illustrate a *differend*, a difference so vast between participants that it can never be bridged (Lyotard, 1988). But, the machinery of the Conversation Map (i.e., those functions which automatically compile a rough-draft thesaurus for a set of messages) works in a strictly mechanical manner that sums and then averages together the language of the group. The Conversation Map is doggedly dialectical. Because of the way it is built, it *cannot not* find a common ground. Consequently, even for an argument so vicious or incoherent that a skilled, human negotiator might find no place to start building common ground, the Conversation Map will diagram -- through its mechanical operations -- a potential synthesis. In a recent interview, I explain how this role of the software (to articulate the synthesis and limits of common ground or common sense) is akin to the role played by performance artists and philosophers of a Socratic persuasion (Dumit and Sack, 2000). Also, in an entry written for the *Oxford Encyclopedia of Aesthetics*, I explain how this constitutes a new sort of artistic, software design aesthetic (Sack, 1998).

The Conversation Map is unlike other electronic art and software design work that has been done to map out the written exchanges of email and online chat. Previous work in this area (e.g., Donath, Karahalios and Viegas, 1999; Cannon and Szeto, 1998) has tended to concentrate on how messages are threaded and/or how social networks of response patterns are constructed without building any sophisticated linguistic analysis into the software. Finding the limits and dialectical syntheses of contentious language from an email archive of an argument is thus difficult, if not impossible, using the work of these other artists and designers.

There is a sort of -- perhaps ridiculous -- optimism built into the machinery of the Conversation Map. The output of the Conversation Map is therefore not simply a description of the *status quo*. Rather, the output can be interpreted as a set of possible goals, a set of landmarks that can be used to navigate, to steer, the conversation forward into the future. This is quite unlike much recent artistic mapping work. For example, Laura Kurgan, in a piece for the show *World Views: Maps & Art* (Kurgan, 2000) entitled "Spot 083-264: Kosovo, June 3, 1999," used the images from French SPOT satellites to map out the burning villages, the mass graves and hidden refugees ("seen" in forests and hills with heat sensors) of the Drenica valley of Kosovo for approximately the same time as the Conversation Map image discussed above. Kurgan's images are invaluable as memorials to horrific events which must neither be forgotten nor trivialized. However, they lead us out of discussion, conversation and dialectics and into an *aporia*: What can possibly follow these events? In contrast, the images of the Conversation Map are naive and ridiculous: they are invented, hopeful landmarks with which to navigate through the conversation. They play the role of the Socratic jester who voices the unthinkable: perhaps compromise and common ground is possible, perhaps healing can be accomplished after these unforgivable acts have taken place?

References

Steve Cannon and Gong Szeto, *Parasite*, <http://parasite.io360.com/index.html> and <http://www.cybergeography.org/atlas/topology.html>, 1998.

Stuart K. Card, Jock D. Mackinlay and Ben Shneiderman (editors), *Readings in information visualization: using vision to think* (San Francisco, CA: Morgan Kaufmann Publishers, 1999).

Gilles Deleuze, *Foucault*, translated and edited by Sean Hand (Minneapolis: University of Minnesota Press, 1988).

Judith Donath, Karrie Karahalios, and Fernanda Viegas "Visualizing Conversations," in *Proceedings of HICSS-32* (Maui, HI: IEEE Computer Society, January 5-8, 1999).

Paul Dourish and Matthew Chalmers. "Running Out of Space: Models of Information Navigation." Short paper presented at *HCI'94* (Glasgow, UK, 1994).

Joseph Dumit, "Artificial Participation: An Interview with Warren Sack," in *Zeroing in on the Year 2000: Late Editions 8, Cultural Studies for the End of the Century*, edited by George E. Marcus (Chicago: University of Chicago Press, 2000).

Christiane Fellbaum (editor) *WordNet: An Electronic Lexical Database* (Cambridge, MA: MIT Press, 1998).

Michel Foucault, "Parrhesia and Community Life," In *Discourse and truth: the problematization of parrhesia, six lectures given at the University of California at Berkeley, Oct-Nov 1983*, edited by Joseph Pearson. The text was compiled from tape-recordings made of the lectures delivered in English. The transcripts are the notes of one of his auditors and are online here: www.parrhesiast.com

Eugene Garfield, *Citation Indexing: Its Theory and Applications in Science, Technology and Humanities* (New York: John Wiley, 1979).

Gregory Grefenstette, *Explorations in Automatic Thesaurus Discovery* (Kluwer Academic Publishers: Boston, 1994).

Barry N. Hague and Brian D. Loader, *Digital democracy: discourse and decision making in the Information Age* (New York: Routledge, 1999)

Graeme Hirst and David St-Onge. "Lexical Chains as Representations of Context for the Detection and Correction of Malapropisms" in *WordNet: An Electronic Lexical Database*, edited by Christiane Fellbaum (Cambridge, MA: MIT Press, 1998).

Ian Hutchby and Robin Wooffitt, *Conversation analysis: principles, practices, and applications* (Malden, MA: Polity Press, 1998).

George Lakoff and Mark Johnson. *Metaphors We Live By* (Chicago: University of Chicago Press, 1980).
Jean Francois Lyotard, *The differend: phrases in dispute*, translation by Georges Van Den Abbeele (Minneapolis: University of Minnesota Press, 1988)

Alan J. Munro, Kristina Hook and David Benyon (editors), *Social navigation of information space* (New York: Springer Verlag, 1999).

Paul Resnick and Hal R. Varian, "Introduction: Special Section on Recommender Systems," *Communications of the ACM*, 40 (2) (March 1997).

Frances Richard, "Utterance is place enough: mapping conversation," *Cabinet: A quarterly magazine of art and culture*, Issue 2, Spring 2001, pp. 76-81.

Warren Sack, "Artificial Intelligence and Aesthetics," in *The Encyclopedia of Aesthetics, Volume 1*,

Michael Kelly, editor-in-chief (New York: Oxford University Press, 1998)

Warren Sack "Stories and Social Networks" in the *Proceedings of the American Association of Artificial Intelligence Workshop on Narrative Intelligence* edited by Phoebe Sengers and Michael Mateas (Cape Cod, MA: AAAI, November 1999).

Warren Sack, "Discourse Diagrams: Interface Design for Very Large-Scale Conversations" in the *Proceedings of the Hawaii International Conference on System Sciences, Persistent Conversations Track* (Maui, HI: IEEE Computer Society, January 2000a).

Warren Sack "Conversation Map: A Content-based Usenet Newsgroup Browser," in *Proceedings of the International Conference on Intelligent User Interfaces* (New Orleans, LA: Association for Computing Machinery, January 2000b).

Warren Sack and Joseph Dumit, "Very Large-Scale Conversations and Illness-based Social Movements" presented at *Media in Transition* (Cambridge, MA: MIT, October, 1999).

Deborah Schiffrin, *Approaches to discourse* (Cambridge, MA: Blackwell, 1994).

Laura Kurgan, "Spot 083-264: Kosovo, June 3, 1999," in *World Views: Maps & Art* curated by Robert Silberman (Minneapolis: University of Minnesota Press, 2000).

Roza Tsagarousianou, Damian Tambini, and Cathy Bryan (editors) *Cyberdemocracy: technology, cities, and civic networks* (New York: Routledge, 1998).

UNDP, "New technologies and the global race for knowledge" in *The Human Development Report* (UN Development Programme, 1999)