DESIGN FOR VERY LARGE-SCALE CONVERSATIONS

Warren Sack

B.A. Computer Science and Psychology, Yale University, 1985
M.S. Media Arts and Sciences, Massachusetts Institute of Technology, 1994

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

Doctor of Philosophy
at the
Massachusetts Institute of Technology

February 2000

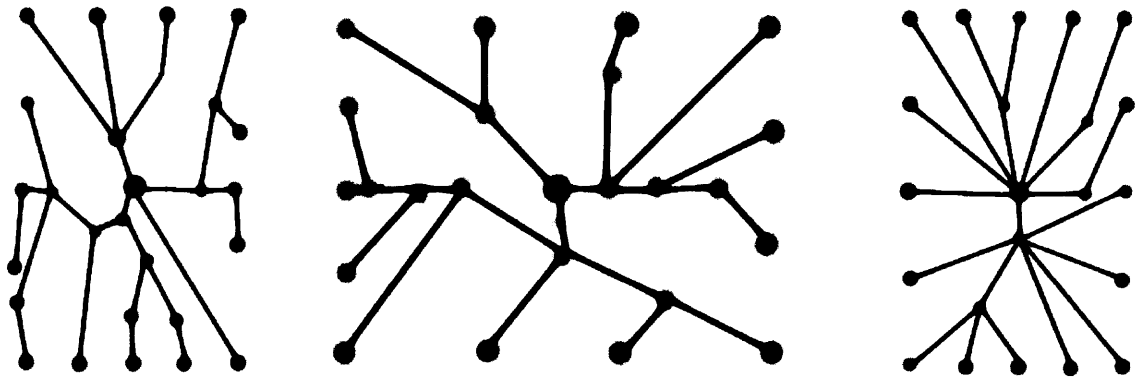

Author .....
Program in Media Arts and Sciences
January 2000

Certified by .....
Kenneth Haase
Visiting Associate Professor of Media Arts and Sciences
Thesis Supervisor

Accepted by .....
Stephen A. Benton
Departmental Committee on Graduate Students
Program in Media Arts and Sciences

ABSTRACT

Design for Very Large-Scale Conversations

Warren Sack

Submitted to the Program in Media Arts and Sciences
School of Architecture and Planning on January 7, 2000
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

On the Internet there are now very large-scale conversations (VLSCs) in which hundreds, even thousands, of people exchange messages across international borders in daily, many-to-many communications. It is my thesis that VLSC is an emergent communication medium that engenders new social and linguistic connections between people. VLSC poses fundamental challenges to the analytic tools and descriptive methodologies of linguistics and sociology previously developed to understand conversations of a much smaller scale. Consequently, the challenge for software design is this: How can the tools of social science be appropriated and improved upon to create better interfaces for participants and interested observers to understand and critically reflect upon conversation? This dissertation accomplishes two pieces of work. Firstly, the design, implementation, and demonstration of a proof-of-concept, VLSC interface is presented. The *Conversation Map* system provides a means to explore and question the social and linguistic structure of very large-scale conversations (e.g., Usenet newsgroups). Secondly, the thinking that went into the design of the Conversation Map system is generalized and articulated as an aesthetics, ethics, and epistemology of design for VLSC. The goal of the second, theoretical portion of the thesis is to provide a means to describe the emergent phenomenon of VLSC and a vocabulary for critiquing software designed for VLSC and computer-mediated conversation in general.

Thesis Supervisor — Kenneth Haase
Title — Visiting Associate Professor of Media Arts and Sciences

COMMITTEE

Advisor .....

Kenneth Haase
Visiting Associate Professor of Media Arts and Sciences
Program in Media Arts and Sciences
Massachusetts Institute of Technology

Reader .....

Walter Bender
Senior Research Scientist
Program in Media Arts and Sciences
Massachusetts Institute of Technology

Reader .....

John Rajchman
Visiting Professor of History, Theory and Criticism
Department of Architecture
Massachusetts Institute of Technology

Reader .....

Krzysztof Wodiczko
Professor of Visual Arts
Department of Architecture
Massachusetts Institute of Technology

CONTENTS

ACKNOWLEDGMENTS

Thank you to my advisor Ken Haase. You gave me the opportunity to come to the Media Lab and welcomed me into the Machine Understanding Group. It has been an honor and a delight to be a part of your research group; a group which has included so many amazing people during my stay at the Lab. Thank you for your advice, your support, and your friendship.

Thank you to my reader Walter Bender. By supporting my work through the News in the Future Consortium and then later giving me the privileges of a virtual membership in your own Electronic Publishing Group you have simultaneously fed me and anchored my work to the fascinating pragmatics that make the Media Lab more than a think tank. Thank you for your grounded sense of what is good and what is right.

Thank you to my reader John Rajchman. Thank you for your careful readings, rigorous but gentle critiques, conversation and hospitality. You have taught me how to think in new ways and introduced me to a world of philosophy. Your encouragement and friendship have provided me with confidence in my work and a motivation to make it much better.

Thank you to my reader Krzysztof Wodiczko. By making me a member of your Interrogative Design Group at the MIT Center for Advanced Visual Studies and then giving me the opportunity to teach with you, you have completely changed my work. By teaching me so much you have made me acutely aware of how little I know. In your Socratic way I know you will insist that you have taught me nothing, but you have been a mentor to me despite what you might sometimes insist.

Thank you Mike Best. Your work has inspired me and your calm advice and sharp humor has saved both my morale and the quality of my work over and over again. I am glad we got to be graduate students together.

Thank you Joe Dumit. You have been both a collaborator and a guide. Even during those times when this project seemed stalled or without direction, you always found something insightful or motivating to tell me. Your belief in me, your advice, your suggestions and your work have provided my work with a coherence and a direction that I could never have found without you.

Thank you Sawad Brooks. Your skills as a designer, programmer, and critical thinker have been inspirational. Thank you also for the many pieces of practical advice and theoretical insights you have given me. Better officemates cannot possibly exist. Every day you shared with me a sense of the absurd and a conversation about art, design, philosophy, and programming.

Edith Ackermann, thank you! You have been my teacher, my friend, and my apartment mate. This last nine months of work towards the degree has been difficult and exhausting, but it would have been infinitely more so without the warm home, cheer, and intellectual exchange you have so generously shared with me.

Thank you Marc Davis. You and Mike Travers were the ones to first invite me to the Lab and your continued interest and support for my work has been one of my lifelines that allowed me to make it through. Thank you for giving me a year to spend with you and your group at Interval Research Corporation. Your friendship, charisma and incredible capacity for work are amazing.

Thank you Mitchel Resnick for inviting me to participate in some of the life and outings of your research group. Your interest in me and my work gave me new strength and insights.

In addition to the people I have named above, over the past several years I have had conversations and email exchanges with a number of people who have given me new insights into this dissertation project: Phil Agre, Hayward Alker, Amit Bagga, Jon Beller, Erik Blankinship, Rick Borovoy, Joost Breuker, Norman Bryson, Janet Cahn, Justine Cassell, Anil Chakravarthy, Joe Chung, Dalton Conley, Sharon Daniel, Sara Elo, Teun van Dijk, Eric Dishman, Kelly Dobson, Cris Dolan, Judith Donath, Jack Driscoll, Ron Eglash, Tom Erickson, Vadim Gerasimov, Jennifer González, Matt Grenby, Dan Gruel, Mark Guzdial, Sanda Harabagiu, Marti Hearst, Graeme Hirst, Camilla Holten, Henry Jenkins, Natalie Jeremijenko, Yasmin Kafai, Karrie Karahalios, Sung Ho Kim, Greg Kimberly, Tammy Lackner, Jason Lewis, Henry Lieberman, Pia Lindman, Tsipora Mankovsky, Patrice Maniglier, Cameron Marlow, Kevin McGee, Marvin Minsky, Erik Mueller, Jon Orwant, Jonah Peretti, Randy Pinkett, Roy Rodenstein, Nitin Sawhney, Sha Xin Wei, Phoebe Sengers, Sylvia Sensiper, Tara Rosenberger Shankar, Brian Smith, Elliot Soloway, Carol Strohecker, Sunil Vemuri, Ann Weinstone, Marc Weiss, Adam Whiton, Terry Winograd, Becca Xiong, Espen Aarseth. Thank you all.

Thank you Felice Gardner and Rebecca Prendergast for your friendship and for making everything work.

Thank you Linda Peterson for shepherding me along the way towards graduation and always having the time to give me good advice when I needed it.

Thank you Nicholas Negroponte for creating such an amazing place to work and study.

Thank you News in the Future Consortium sponsors, AT&T, and Interval Research Corporation for the research fellowships you have given me during my years in graduate school.

Joe Dumit, Sylvia Sensiper, and Andrew, Anne Harley, Mike Best, Judith Donath and Joe Chung, Janet Cahn, Joey Berzowski, Maggie Orth and Josh Smith, Claudio and Monica Pinhanez, Push Singh, Laura Teodosio, Ken and Margaret, Ben Vigoda, Carol Strohecker, Sung Ho Kim, Tim Anderson, David Cavallo, and Edith Ackermann have shown me that there is life in Boston/Cambridge outside of the Lab. Thank you for taking such good care of me and Jennifer especially during the final months of the thesis.

My most recent officemates, Erik Blankinship, Tammy Lackner, and Tara Rosenberger Shankar have made the last nine months in the Lab a fun and friendly place. The place in which our offices have been located, the Cube, has been a great place to work and I want to thank everyone in Mitchel Resnick's group, Brian Smith's group, and Walter Bender's group for that.

I have not named all of my friends at the Lab and I hope everyone not mentioned will forgive me. I love you all. I have not even begun to thank all of my friends outside of the Lab who have patiently indulged me in the seeming endlessness that is called graduate school. Surf's up dudes! *J'ai fini!*

Thank you to my family: my dad, Ron; my mom, Peg; my sister, Leslie; and my grandfather, Lyman. You are my roots and you have nurtured me and given me strength, love and a home.

Thank you to the family of my partner Jennifer González: Crispin González, Kirsten González, Consuelo González, Richard Middleton, Lorraine Bell, and Nelson Bell. You have taken me into your family and given me a new place to also call home.

Thank you to my partner in life, my soulmate Jennifer González. Thank you for your love, your laughter and your wisdom. I would not have survived this without you, but with you it was a fun and enlightening adventure.

INTRODUCTION

The utopic hope that motivates the work of this thesis is a hope for a global, international, cross-cultural, trans-linguistic, all-inclusive conversation. Although truly global conversations (involving billions of people) of this sort do not yet exist, very large-scale conversations (VLSCs) involving hundreds or thousands of people do now occur. It is argued that a number of design challenges must be met to transform existing VLSCs into global conversations. A number of approaches to meet these design challenges could be envisioned. A tactical software design approach is described and two means of developing the approach are identified. Firstly, such an approach can be developed through the production of prototype systems. As part of the work for this thesis, a prototype system for browsing VLSCs has been designed and implemented. Secondly, this thesis formulates an ethics, aesthetics, and epistemology of software design for VLSC in particular and computer-mediated conversation in general. Abstracts for each of the remaining chapters of this thesis are collected at the end of the introduction to provide the reader with an outline of the thesis.

Although there does not yet exist a truly global conversation, on the Internet there are now very large-scale conversations (VLSCs) in which hundreds, even thousands, of people exchange messages. These messages are exchanged daily -- and even more frequently -- across international borders. Unlike older, one-to-many media (for example, television or radio) where a small group of people broadcast to a larger number of people, VLSCs are a many-to-many communications medium. Also, unlike older, one-to-one media (e.g., the telephone), the people engaged in VLSCs do not necessarily know the electronic addresses of the other participants before the start of the conversation. For these reasons, VLSCs can create new connections between people who might otherwise not even have imagined the other's existence. The work described in this thesis is intended to help transform existing VLSCs – e.g., Usenet newsgroups -- into truly global conversations.

This dissertation accomplishes two pieces of work. Firstly, the design, implementation, and demonstration of a proof-of-concept, VLSC browser is presented. The *Conversation Map* system provides a means to explore and navigate the social and semantic structures of VLSCs (e.g., Usenet newsgroups). Secondly, the thinking that went into the design of the Conversation Map system is generalized and articulated as an aesthetics, ethics, and epistemology of design for VLSC. The goal of the second, theoretical portion of the thesis is to provide a means to describe the emergent phenomenon of VLSC and a vocabulary for critiquing software designed for VLSC and computer-mediated conversation (CMC) in general.

VLSC is both a well-known phenomenon but also, simultaneously, something as yet largely unexamined by designers and theorists. On the one hand, VLSC is well known in the form of busy Usenet newsgroups and large, electronic mail lists. For participants and observers alike, VLSC manifests itself as huge lists of messages in a conventional email reader like RN, Eudora, or Netscape Messenger.

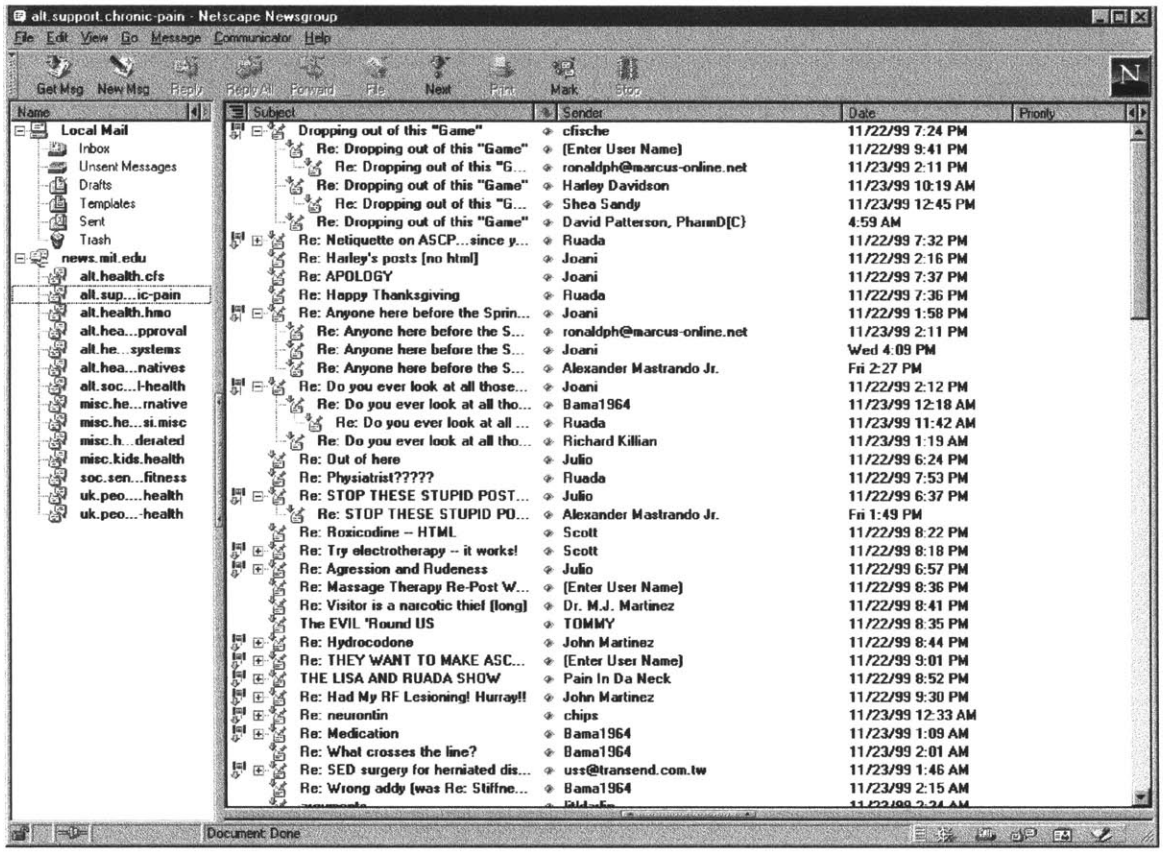

On the other hand, VLSC is largely unexamined. What does it mean to have a conversation that involves hundreds or thousands of people? Existing theories of conversation and discourse do not cover this scale of conversation. Moreover, very little design work for VLSC has been done. For example, why is VLSC usually represented as a long list of email messages? Isn't something better possible? In fact, with a better theory of VLSC, better software for navigating VLSCs can be designed.

Maps and Navigation

Lucy Suchman has shown how the activity of navigation can be a very different practice in different cultures. To illustrate this, she compares the navigation activities of European sailors, which are more plan-directed, to those of the Trukese whose navigation seems to be more contingent upon the environment.[1] In its current forms, VLSC is usually an inter-cultural phenomenon since it is usually conducted on the Internet between participants from many different countries. To theorize and to design software to browse and navigate VLSCs it is necessary to make explicit the culturally-specific assumptions that go into the design work. Consequently, in addition to implementation-oriented discussions of the Conversation Map VLSC browser, this thesis includes a large amount of

[1] Lucy Suchman, *Plans and Situated Actions: The Problem of Human-Machine Communication* (New York: Cambridge University Press, 1986), vii-x.

philosophical and historical analysis in order to provide the tools necessary to critique software like the Conversation Map system.

Michel Foucault has pointed out that "the comparison between medicine and navigation is a very traditional one in [ancient] Greek culture."[2] Medicine, navigation, and government have to do with guidance, control, and governance. Etymologically, the verb *navigate* comes from the combination of words *navis* [ship] and *agere* [to guide]. Thus, in the current case of the navigation of a large information space -- like a large archive of Usenet newsgroups -- the "ship" has been replaced by a self, and so the point of navigation is self-guidance or self-governance. From this perspective, the right to 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 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. The larger ethical and philosophical implications of this understanding of navigation will be explored in this thesis.

VLSC is a medium; i.e., it mediates people. As such, it connects some people together and separates others. Like previous media it functions as a substrate in which and on which groups of people constitute themselves. There is a long history of the use of media as *technologies of the self*,[3] as reflective and communicative media for the construction of social, psychological, economic, and political self-governing people and peoples. Diaries have been used for millennia by particular people as a medium for self-reflection, for writing down and shaping the person's image of self.[4] Diaries are a medium that functions as a technology of the self where "self" is understood to be the self of one person. The oral storytelling practices of folktales function in an analogous manner for the formation and description of a slightly larger self, a self of a small group of people. Oral storytelling of folktales is a means for articulating the values and identity of small, tightly knit clusters of people. The facilitation of the production of larger selves, of the selves of self-governing nations, for instance, requires a

---

[2] Michel Foucault, "Parrhesia and Community Life," In *Discourse and truth: the problematization of parrhesia*, edited by Joseph Pearson, Six lectures given at the University of California at Berkeley, Oct-Nov. 1983. The text was compiled from tape-recordings made of the lectures delivered, in English, by Michel Foucault at the University of California at Berkeley in the Fall Term of 1983, transcribed and edited in 1985. The lectures were given as part of Foucault's seminar, entitled "Discourse and Truth." Since Foucault did not write, correct, or edit any part of the text transcripts, they lack his imprimatur and do not present his own lecture notes. The transcripts are the notes of one of his auditors and are online here: http://perso.club-internet.fr/kmille/discourse.htm

[3] Michel Foucault, "Technologies of the Self," in *Ethics: Subjectivity and Truth; Essential Works of Foucault 1954-1984, Volume I,* edited by Paul Rabinow (New York: The New Press, 1997).

[4] *Ibid.*

different kind of medium. Scholars have shown how the mass production capabilities of high-speed printing presses made possible the media of novels and newspaper stories that were essential to the formation of the modern nation-state.[5] VLSCs do and can function as the substrate for new kinds of selves, new sorts of groups of people, that are as yet unnamed.[6] These new groups of people can be transnational or international in scope.

Individuals need a map of the group in order to find their current or desired position in the group. Groups need a map to reflect on their limits and internal structure. This map can be either a metaphorical or a literal map. Maps have historically been very important for geographically-based groups. On a country's map, citizens can find their home, their proximity to the capital, their range of travel experience, etc. Maps usually incorporate several kinds of information; e.g., political boundaries, roads, and elevations might be included on one map of a region. No map can incorporate all kinds of information, but -- at least for geographically-based maps -- it seems clear that certain kinds of information are essential to all maps. For example, the map of a country needs to include some representation of the country's borders. However, it is not clear what needs to be included in a map of a non-geographically-based group. What, for instance, does a member of a Usenet newsgroup need to see in order to navigate through the VLSC of the newsgroup?

A VLSC is a "space" created through the electronic exchange of words. A map of a VLSC depicts a sea of information in which participants sail through, make waves in, and sometimes beach themselves to become an island or even a continent. Since the map of a VLSC depends upon the words and deeds of the participants, a participant might use the map to navigate and by so navigating change the map. In this way map construction and navigation of VLSCs can be mutually recursive activities. One might imagine that this relationship between navigation and map-making is specific to information spaces, but, in fact, it takes only a quick glance at a world map of ten years ago -- before the end of the Cold War -- to see that geographically-based map-making is predicated on social and political forces and the boundaries of social and political forces are outlined on maps.

What then are the "boundaries" of a VLSC that can be drawn on a map? In this thesis it will be argued that VLSCs can produce and reproduce social relationships between people and semantic relationships between words and, so, social networks and semantic networks are essential to the mapping of VLSC. Paul Dourish and Matthew Chalmers -- researchers in human-computer interaction (HCI) and computer-supported cooperative work (CSCW) -- claim there are at least three forms of navigation which can be combined in information

---

[5] Benedict Anderson, *Imagined communities: reflections on the origin and spread of nationalism* (London: Verso, 1983).

[6] See Gilles Deleuze and Félix Guartarri, *A thousand plateaus: capitalism and schizophrenia*, translated by Brian Massumi (London: Athlone Press, 1988), 469-471.

systems: (1) spatial; (2) semantic; and, (3) social navigation.[7] Most graphical interfaces make use of spatial layout (and thus facilitate spatial navigation); some use semantic navigation (e.g., hypertexts); and some social navigation (e.g., a variety of work in the sociology of social networks). However, very few interfaces combine all three sorts of navigation. The Conversation Map system provides the means to spatially navigate through social networks, semantic networks and intersections of the two, thus, implementing all three of Dourish and Chalmers' forms of navigation.

To better explain why a combination of social and semantic networks might be a useful description of a VLSC, it is helpful to examine the claims of the linguist Michael Halliday. According to Halliday[8], language has at least three meta-functions: (1) *ideational*: language can represent ideas; (2) *interpersonal*: language functions as a medium of exchange between people; and, (3) *textual*: language functions to organize, structure, and hold itself together; this function allows the various devices of cohesion, including citation, ellipsis, anaphoric reference, etc. to be used. Thus, for example, I can write here, in this present sentence, about the first sentence of the present paragraph and the reader can infer that I am referring to the sentence that begins like this: "To better explain why..." Any interesting map of VLSCs is incomplete if it does not incorporate all three of these meta-functions (ideational, interpersonal, and textual). As will be demonstrated with the Conversation Map system, a suitably improved implementation of social networks can represent, even if only very roughly, the interpersonal and textual aspects of a VLSC; and, semantic networks can be an approximation of the ideational content of a VLSC.

Social Networks

One of the results of a VLSC is a set of *social networks*. After a while, participants in an online discussion get to know one another and exhibit characteristic patterns of interaction with one another. Some participants post messages that often strike others as interesting, evocative, or worthy of a reply and so these people tend to garner more responses to their messages than others do. Some participants take pains to reply to the posts of newcomers and regulars alike and so build a reputation as virtual moderators for the discussion (even in groups with no officially designated moderators). Some people post what others consider to be "spam" and so, even though they may tend to post frequently, they are ignored (i.e., their posts do not earn replies from other participants).

These collections of posting profiles for the participants of a group, when graphically assembled together, constitute what would be recognizable to social

---

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

[8] Michael A. K. Halliday. *An Introduction to Functional Grammar, Second Edition* (London: Edward Arnold, 1994), 179.

scientists as a *social network*.[9] A social network can summarize a large number of interactions between participants. A map of this sort can be more or less sophisticated.

The simplest social networks possible for VLSCs are those that can be computed through an analysis of references between messages; i.e., an analysis of the "threading" of the messages. Threading is easy to compute given the format of email. But a simple reply to someone does not constitute a conversation with that person. At the very least, conversation is a form of two-way, back-and-forth exchange. If participant A responds to a message posted by participant B and, then later, participant B responds to a message posted by participant A some form of *reciprocation* can be assumed to exist between A and B. A social network that maps out reciprocated replies between participants is one of the simplest summarizations of the social structure of a VLSC that can be assembled.

A more sophisticated social network can incorporate citation information. This sort of analysis has been extremely useful in science studies and is generally referred to as *citation indexing*.[10] However, while instances of citation are relatively easy to spot automatically in scientific papers because citations are required to appear in a standard format and must be listed in the references at the end of a paper, this is not the case for email messages. Within email messages, citations often occur without mention of the original author, citations are often nested inside one another, and citations do not have a standard format, even though it is quite common (but not required) to precede each line of a quotation with some punctuation, like this:

```
>>> On 31 February 2001 Warren Sack
>>>                wsack@media.mit.edu                wrote:
>>> Hi guys!  The future is really great!
```

These complicating factors make an automatic citation analysis procedure difficult to implement for email messages. One of the small contributions of this thesis is the description of a procedure to identify quotations in email messages.

Combining information about who has responded to and/or quoted from whom produces a social network that can provide an overview of interpersonal activity. However, such a social network does not incorporate anything about the content of the messages exchanged and/or cited. That is to say, it does not incorporate a representation of the *discussion themes* of a conversation.

---

[9] Stanley Wasserman and Joseph Galaskiewicz (editors) *Advances in Social Network Analysis: Research in the Social and Behavioral Sciences* (Thousand Oaks, CA: Sage Publications, 1994).

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

If participant A mentioned the word "baseball" in a post that also quoted a part of a message from participant B wherein B wrote about the term "football," and then, later in the conversation participant B wrote about basketball in response to a message by A concerning soccer, then the link between A and B in the social network might be labeled with the term "sports" since baseball, football, soccer, and basketball are all sports. An analysis of reciprocated "discussion themes" of this sort is possible to compute given a new algorithm for *inter-message lexical cohesion* presented in this thesis.

Strictly speaking -- i.e., according to the terminology of linguistics -- the sports example outlined above is not an illustration of the identification of discussion themes *per se*, but, rather, is an example of an analysis of *lexical cohesion.* Lexical cohesion is a measure of the repetition of similar terms in one or a series of texts.[11] Performing an analysis of lexical cohesion is only one step of many that would be required if – within linguistics -- it could be claimed that a system can identify discussion themes. However, since an analysis of lexical cohesion is a necessary step in the determination of discussion themes, I will call the analysis an analysis of discussion themes for the sake of simplicity.

A social network constructed and labeled with information about who replies to whom, who quotes from whom, and the discussion themes addressed through these replies and/or citations provides a map of a group's *social cohesion.*[12] With this map it is possible to estimate whether the group is loosely constituted or close knit and which themes of discussion hold the group together in conversation. A contribution of this thesis is a computable definition of social cohesion for VLSCs.

Semantic Networks

Another result of a VLSC is a set of *semantic networks.* Over the course of many exchanges, participants in a VLSC coin new acronyms (e.g., IMHO, "in my humble opinion"), new punctuation (e.g., the (in)famous smiley faces :-), and new or idiosyncratic ways of using existing vocabulary terms (e.g., "a flame" means something rather particular online). Semantic relations between new and existing lexical items (i.e., words and abbreviations) can be represented in a semantic network.

It is possible to understand the semantic networks of a VLSC as a map of the emergent *metaphors* or *definitions* of participants' discussion. To clarify this insight it is worthwhile examining an example of how the cognitive scientists

---

[11] Michael A.K. Halliday and Ruqaiya Hasan *Cohesion in English* (New York: Longman, 1976).

[12] Warren Sack "Diagrams of Social Cohesion," in *Descriptions of Demonstrated Systems, ACL'99* (University of Maryland, College Park: Association for Computational Linguistics, June 1999).

George Lakoff and Mark Johnson[13] demonstrate the common usage of specific metaphors. Lakoff and Johnson offer the following sentences to support their claim that, in "our" (presumably U.S.) culture, the metaphor TIME IS MONEY is a common presupposition:

TIME IS MONEY
You're wasting my time.
I don't have the time to give you.
How do you spend your time these days?
I've invested a lot of time in her.
I don't have enough time to spare for that.
You're running out of time.
You need to budget your time.

In these sentences the word "time" could meaningfully be replaced with the word "money." The hypothesis is that if two words or concepts are thought about in the same way by a group of people, then they will be systematically used in the same way in the conversation of those people. Examples of language use can be used as symptoms for the diagnosis of semantically related words and phrases.

By mapping out which words in a VLSC are used like which other words, a map of some of the semantics -- i.e., the meaning -- of the conversation can be displayed. The resulting semantic network for a given VLSC is a way to begin to investigate whether or not a group is producing new metaphors or definitions, new ways of talking about known or new subjects. Given an archive of messages, the Conversation Map system computes a set of semantic networks for a given VLSC.

Combining Social and Semantic Networks

For the most part, sociologists working on social network computation and visualization[14] and artificial intelligence (AI), collaborative filtering researchers[15] do not build technologies to facilitate semantic navigation of information. On the other hand, the technologies of corpus-based, computational linguistics[16] and

---

[13] George Lakoff and Mark Johnson. *Metaphors We Live By* (Chicago: University of Chicago Press, 1980), 7-8.
[14] E.g., Barry Wellman, "Living networked in a wired world," *IEEE Intelligent Systems*, (January/February 1999), 15-17.
[15] E.g., Paul Resnick and Hal R. Varian, "Introduction: Special Section on Recommender Systems," *Communications of the ACM*, 40 (2) (March 1997).
[16] E.g., Phil Resnik and Judith L. Klavans (editors), *The balancing act: combining symbolic and statistical approaches to language* (Cambridge, MA: MIT Press, 1996)

information retrieval researchers[17] "black box" the social aspects of the information they index.

Corpus-based computational linguistics is most often performed on large corpora described as, for instance, "10 million words from several volumes of the Wall Street Journal," or "1 million words from a wide variety of text genres." How the authors of the texts included in the corpora interact with one another or are related to one another is not factored into the analysis of the corpus. The one exception to this anonymity of authors is the use of corpus-based techniques for author identification purposes. But, even in these cases, the task is usually to determine who is the most likely author of a text given a small set of possible candidate authors. The social network that incorporates (or the fact that no known social network incorporates) the set of candidate authors is not something that is taken into account in the design of the corpus-based, computational linguistic methods of analysis.[18]

The techniques of corpus-based, computational linguistics are oftentimes technically related to the techniques employed by sociologists since both sets of techniques can depend upon similar tools from statistics and information theory (e.g., measures of mutual information and entropy). But the techniques are essentially inverses of one another due to the fact that what the sociologists black-box in their analyses is almost exactly what the corpus-based linguistics and information technology researchers do not black-box in their own research, and vice versa.

Any significantly new technology of information navigation should involve the combination of these two approaches. To see an archive of VLSC messages -- e.g., Usenet newsgroup messages -- as embedded in and (re)productive of both a set of semantic associations *and* as a facilitator or inhibitor of social networks, it is necessary to explore how social and semantic networks overlap. Moreover, by exploring their overlap it is possible to create new kinds of maps. For example, the combination of lexical cohesion and social networks creates a computable definition of social cohesion for online groups. A second example is the use of social network and social cohesion information to sort out the possibly important semantic networks from the probably trivial or uninteresting. These uses of sociological information to sort out the results of linguistic analyses are implemented in the Conversation Map system and constitute some of the novel and interesting software design results of this thesis.

---

[17] E.g, Gerald Salton, James Allan, Chris Buckley, Amit Singhal. "Automatic Analysis, Theme Generation, and Summarization of Machine-Readable Texts," *Science*, Vol. 264, (3) (June 1994), 1421-1426.

[18] See Warren Sack, "On the Computation of Point of View," in *Proceedings of the National Conference of Artificial Intelligence* (AAAI 94) (Seattle, WA: Morgan Kaufmann Publishers, July 31-August 4, 1994).

Strategies and Tactics

Design work to help establish, support, and extend VLSCs can take one of two approaches: (1) a tactical approach; or, (2) a strategic approach.

Tactics are to be distinguished from *strategies.*[19] Strategies are practices performed by those who have the powers of ownership and/or authorship. For example, owning and administering a private forest for the purposes of hunting is a European, medieval strategy of the aristocracy. Poaching is a tactic that the non-aristocrats performed to survive this strategy. Strategists have the powers to independently redraw a map by changing the territory. Tacticians use a map to navigate through a territory.

By analogy, a strategic design approach is an approach for those who own or control a place or network. A tactical approach is for those who might be inhabitants of the place or network, but who do not have the power to autocractically make decisions about the place or network. A strategic technology creates a territory. A tactical technology, like a map, helps one find one's way through a territory.

The strategic approach to design for computer-mediated communication (CMC), community, or conversation presupposes a set of enormous powers and resources. Strategically designed technologies take the form of an architecture or framework within which all communications or coordinated actions must take place. A system is built around a strategic design and all participants are encouraged to abandon any old media of exchange or expression and to

---

[19] The distinction between tactics and strategies comes from the French, philosopher Michel de Certeau who, in his book, *The Practice of Everyday Life* wrote this: *I call a strategy the calculation (or manipulation) of power relationships that becomes possible as soon as a subject with will and power (a business, an army, a city, a scientific institution) can be isolated. It postulates a 'place' that can be delimited as its own and serve as the base from which relations with an 'exteriority' composed of targets or threats (customers or competitors, enemies, the country surrounding the city, objectives and objects of research, etc.) can be managed. As in management, every 'strategic' rationalization seeks first of all to distinguish its 'own' place, that is, the place of its own power and will, from an 'environment.' A Cartesian attitude, if you wish: it is an effort to delimit one's own place in a world bewitched by the invisible powers of the Other. It is also the typical attitude of modern science, politics, and military strategy. ... By contrast with a strategy, a 'tactic' is a calculated action determined by the absence of a proper locus. No delimitation of an exteriority, then, provides it with the condition necessary for autonomy. The space of a tactic is the space of the other. Thus it must play on and with a terrain imposed on it and organized by the law of a foreign power. ... It operates in isolated actions, blow by blow. It takes advantage of 'opportunities' and depends on them, being without any base where it could stockpile its winnings, build up its own position, and plan raids. It poaches ... It creates surprises ... It can be where it is least expected.* Michel De Certeau, *The Practice of Everyday Life*, translated by Steven Rendall (Berkeley, CA: University of California Press, 1984), 35-37.

conform to the governing influences of the new protocol and/or language. Many designers of CSCW systems have had the powers to pursue a strategic approach because they have built software for private or institutionally localized discussion; e.g., the communications exchanged within a company. But, to acquire the powers necessary for a strategic approach to work on a global scale would be tantamount to gaining power over how everyone communicates. This is unusual but not unprecedented in the "public" areas of the Internet. When Tim Berners-Lee invented the hypertext transfer protocol (HTTP) and hypertext markup language (HTML), he effectively fathered the World Wide Web and so transformed electronic communications.[20] Strategic design is more or less a winner-takes-all game. The nation, company, group, or individual who controls the strategic protocol or language, sets the terms of exchange and, thereby, wields enormous powers. Witness the ongoing power struggles to (re)define the standards for the World Wide Web.

Unlike a strategic, "architectural" approach, a tactical approach is an approach of *bricolage*,[21] a means to patch or incrementally improve an existing "house," i.e., an existing architecture. Strategic approaches are only useful for huge, powerful players, while tactical approaches can be undertaken without gaining the resources and forces of a real or metaphorical army. Instead of attempting to convince or coerce everyone to use a new protocol or idealized language, the tactical approach is to provide participants with the technological means to improve their situation. This thesis describes a tactical approach to the design of computational tools for computer-mediated conversation. A tactical approach to improve the quality of online, public discourse is an attempt to work within existing, strategically established standards. The tactical approach discussed in this thesis is an attempt to produce tools that can help existing VLSCs grow into global conversations. In other words, the idea is to design tools so that the archives of conversations can be explored and navigated by participants. The aim is to play with what we already have and to produce tools that can help us improvise something better as we continue our conversations.

Outline of the Thesis

The thesis is divided in two. The first half is a description and demonstration of the prototype Conversation Map system. The second half is a theory of design for computer-mediated conversation with special attention to VLSC. The theory chapters expand and generalize the design issues touched on in the description of the Conversation Map system. The theory can be read without the system demonstration and vice versa, but it is hoped that the demo system will help make the theoretical discussion more concrete.

---

[20] Tim Berners-Lee. *Weaving the Web: The original design and ultimate destiny of the World Wide Web by its inventor* (San Francisco: HarperSanFrancisco, 1999).

[21] Claude Lévi-Strauss. *The Savage Mind* (Chicago: University of Chicago Press, 1966), 17.

The first series of three chapters, concerning the Conversation Map system, deals with (1) the interface; (2) a set of example uses of the system that show what good the system provides; and, (3) the analysis methods developed to process and cross-index the messages from VLSCs. The three chapters of theory are entitled: (1) ethics; (2) aesthetics; and, (3) epistemology. Related work is discussed throughout the text.

A final, seventh, chapter outlines the specific claims, contributions, and conclusions of the thesis work. Plans for future work are outlined at the end of this chapter.

Following is a list of six abstracts. Each summarizes one chapter of the thesis. The abstracts are repeated at the beginning of each chapter. They have been gathered together here to provide the reader with a roadmap through the text.

Part One: System

(1) **Interface**: This chapter explains the design rationale used to borrow graphical elements and analyses procedures from the social sciences and rework and extend them for use as interface devices. In particular, two tools from the social sciences -- social networks and, semantic networks -- are examined as possible interface components for a VLSC browser. The design challenge lies in the attempt to stretch the tools of the social sciences both (a) to meet the demands of social scientists interested in studying the phenomenon of VLSC; and, (b) to render the tools in a form that makes it possible for the layperson – the non-specialist participant in these VLSCs – to achieve some of the insights that these tools make possible for the social scientist.

(2) **Examples**: The online conversations of several large groups are examined through the graphics generated by the Conversation Map system. By considering the reflections of the VLSCs of (1) audiences, (2) social movements, (3) user groups, and (4) a group I will name "citizen diplomats," a clearer picture can be had of the structure of these new, network-based groups. For the social scientist, this is useful because it provides a place to begin to explore the social and semantic significance of these new groups. For an actual or potential participant, gaining a feel for the social and semantic structure of a VLSC is a means to decide whether the VLSC is worth joining or whether particular contributions to the group have been pivotal or marginal to the structure of the group.

(3) **Analysis**: The analysis procedures of the Conversation Map system use, extend, and combine in novel ways techniques from computational linguistics and quantitative sociology. The goal of the analysis procedures is to generate a graphical summary of archives of online, email-based VLSCs. Three new analysis techniques of potential interest for computational linguistics and sociology are developed: (1) Inter-message lexical cohesion:

An algorithm has been created for the computation of lexical ties between email messages that takes into account information about both threading and quotation in the messages (thereby producing an approximation of the themes of conversation). (2) Social cohesion: A definition of social cohesion has been developed. Social cohesion is the overlay of lexical cohesion information on top of a social network (thus creating a rough description of what the conversations between participants are about). And, (3) Social network-based sorting: A means has been designed and implemented to use the computed social networks to weight the importance of word entries in automatically computed rough draft thesauri (which is useful for selecting the terms which are important as discussion themes and thus important to the social structure of the group). A step-by-step description of the analysis procedures is presented. Improvements over related work in CSCW, sociology, and computational linguistics are discussed.

Part Two: Theory

(4) **Ethics**: The good of new computer network-based technologies for conversation is that they can potentially provide us with new means for free speech and therefore new powers of self-governance. If these new forms of conversation work effectively as new technologies of self they should make it possible for us to better understand how and where we are located in a variety of social and linguistic networks. Furthermore, by making these networks visible they should give us the tools to plan how we want our interconnections to be between us, other people, and non-human entities, like the machines we depend upon everyday. I argue for a means to visualize the emergent social and semantic networks that are now hard to perceive. By making them visible for everyone to see, it is hoped that it will be possible for them to be used as resources for public discussion. Who is excluded from the linked social networks of public view? Who is central? Why? What are the predominant associations connected with the signs of public life and public discourse? Who and what are accorded positive even glorious associations? Who are given more modest or even negative associations? Why? Why not? A great amount of effort is currently devoted to gathering online information about people for the work of marketing, advertising, etc. At the same time, many people are building home pages on the World Wide Web and creating other electronic projections of themselves. The ethics of software design must examine how these two activities coincide and sometimes clash.

(5) **Aesthetics**: Previous software design approaches (especially those of artificial intelligence) are closely tied to a *commonsense aesthetics*, i.e., an aesthetics that presumes a commonsense, a predictably emergent commonsense, or the uncanny, interference of the commonsense world. An alternative to these approaches must be found if we are to design for VLSCs where a potential, or virtual, commonsense is contingent upon the possible

(but not necessarily probably) emergence of a community of people who create their own stable semantic and social structure through continued interaction on the Internet. This new aesthetics, therefore, must be useful for the practices of design for emergent communities.

(6) **Epistemology**: Detailed, micro-analyses of face-to-face conversation usually involve a very different kind of work and produce a very different type of research result – i.e., a very different type of knowledge -- than do macro-scale analyses of discourses involving thousands or millions of people. This micro/macro divide is a recurrent one in many of the social sciences and has been widely discussed in, for example, economics and sociology. Bridging this divide for the analysis of VLSCs is necessary because, on the one hand, the phenomenon under examination is macro-scale by definition; but, on the other hand, one of the most important motivations for analyzing VLSCs is to give participants a means to find their way and locate their position in a VLSC. Consequently, standard social scientific methods of dealing with macro-scale phenomenon by working with norms and averages are unworkable because they risk effacing the contributions of particular individuals. It is argued that a bridge can be found between micro and macro scale analyses of online conversations. This bridge is the lexicon, or what might be called the "thesaurus," of a group conversation. On the micro-scale, contributions to a conversation are judged to be coherent and cohesive partially according to whether or not they are taken to be "on topic" by the participants. Knowledge of deviation or convergence with a given topic is based on knowledge of a lexicon; i.e., according to the relationships between and the definitions of words. But, over the course of the lifetime of a group, new (e.g., slang) words are coined, some words gain new meanings and others lose their currency, connotations, or the controversy that surrounds them. Thus, conversation both depends upon and changes the lexicon or "thesaurus" of a group.

PART ONE: SYSTEM

INTERFACE

This chapter explains the design rationale used to borrow graphical elements and analyses procedures from the social sciences and rework and extend them for use as interface devices. In particular, two tools from the social sciences -- social networks and, semantic networks -- are examined as possible interface components for a VLSC browser. The design challenge lies in the attempt to stretch the tools of the social sciences both (a) to meet the demands of social scientists interested in studying the phenomenon of VLSC; and, (b) to render the tools in a form that makes it possible for the layperson – the non-specialist participant in these VLSCs – to achieve some of the insights that these tools make possible for the social scientist.[22]

[22] Parts of this chapter appear in 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: Association for Computing Machinery, January 2000); 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 2000).

Most contemporary VLSCs are conducted as email exchanges in Usenet newsgroups and in large, locally hosted email newsgroups. This chapter describes the interface design of a prototype newsgroup browser, Conversation Map. The Conversation Map system employs a set of text analysis procedures to produce a graphical interface. With the graphical interface one can browse a set of Usenet newsgroup articles according to who is "talking" to whom, what participants are "talking" about, and the central terms and possible emergent metaphors of the conversation. To allow this combination of social and semantic navigation[23] the Conversation Map system computes a social network[24] corresponding to who is replying to (or citing) whose messages. The Conversation Map system also parses and analyzes the contents of the newsgroup articles to calculate a semantic network[25] that highlights frequently used terms that are similar to one another in the Usenet newsgroup discussion. For example, if the discussion includes messages concerning "time" and other messages concerning "money" and these two terms ("time" and "money") are used in similar ways by the discussants (e.g., "You're wasting my time," "You're wasting my money," "You need to budget your time," "You need to budget your money") then the two terms will show up close to one another in the graphically displayed semantic network and so indicate the presence of a literal or metaphorical similarity between the terms (e.g., "Time is money"). In addition, the Conversation Map system analyzes connections between messages to extract an approximation of the discussion themes shared between newsgroup participants.

The output of the text analysis procedures are automatically translated into interface devices that allow one to browse the Usenet newsgroup articles in ways that would be impossible with a conventional, "format-based" news reader (e.g., RN, Eudora, or Netscape). One of the purposes of this research is to produce a better Usenet newsgroup browser for newsgroup participants and others who might like a quick way of discovering the terms and conversational rules of a newsgroup (e.g., sociologists and anthropologists of on-line text and social activity). Examination of the calculated semantic networks provides a means to begin to understand how the shared common terms of a VLSC have emerged over the course of a discussion. The computed themes and social networks are not an explicit statement of the possible rules of conversation of a VLSC, but an examination of them will provide a means to consider who is responding to whom, about which subjects, under what circumstances.

---

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

[24] Stanley Wasserman and Joseph Galaskiewicz (editors) *Advances in Social Network Analysis: Research in the Social and Behavioral Sciences* (Thousand Oaks, CA: Sage Publications, 1994).

[25] Cf., M.R. Quillian "Semantic Memory" In M. Minsky (editor) *Semantic Information Processing* (Cambridge, MA: MIT Press, 1968).

To some, compared to a conventional email reader, the Conversation Map system might look like a fancy appliance overloaded with a bunch of spurious, new features. However, these features are derived from standard tools of the social sciences; tools which have been useful for scientists attempting to understand the social and semantic structures of individuals and groups.

Thus, the obvious first "users' group" for the Conversation Map system is to be found in those scientists for whom these tools are already familiar. In fact, a handful of projects are now underway in which social scientists are using the system to gain a better understanding of various VLSCs.[26] The system is being adapted and extended to meet their needs in a sort of participatory design process.[27]

Simultaneously, several laypeople are experimenting with the system. Why might a layperson want to use tools from social science? Imagine wanting to join in on a VLSC that has been in progress for several months or even years with an archive of thousands of messages. Having a set of machine-generated summaries of the archive provides one with a relatively quick way to decide if the group is focussed on what is expected or desired. The summaries provided also give some indication of how various subjects are being discussed and who is central to the discussion. In addition, there are ways of using the interface to steer around "spam" and focus in on threads of discussion that address specific themes. In short, the social science-derived tools of the interface give one a means of seeing the "forest" of the conversation before diving into the "trees."

So, what a layperson wants to know about an archive of messages can overlap with the goals of social science. Consequently, it seems possible to start at the same point (i.e., with the same prototype system) and then evolve the system in two – perhaps only slightly – different ways: one way for the scientists, another way for the non-scientists. The difficulty of the interface design challenge is proportional to the size of the divergence between these two groups of people. The design challenge lies in the attempt to stretch the tools of the social sciences both (a) to meet the demands of social scientists interested in studying the phenomenon of VLSC; and, (b) to render the tools in a form that makes it possible for the layperson – the non-specialist participant in these VLSCs – to achieve some of the insights that these tools make possible for the social scientist.

This chapter explains the design rationale used to borrow graphical elements and analysis procedures from the social sciences and rework and extend them for use as interface devices. In particular, two tools from the social sciences --

---

[26] One example of collaborations with social scientists is Warren Sack and Joseph Dumit, "Very Large-Scale Conversations and Illness-based Social Movements," presented at the conference *Media in Transition* (Cambridge, MA: MIT, October 1999).

[27] Morten Kyng, "Designing for Cooperation: Cooperating in Design," *Communications of the ACM*, 34(12), 1991, 65-73.

social networks and, *semantic networks* -- are examined as possible interface components for a VLSC browser.

What do social and semantic networks look like?

Rendering a semantic network generally entails sketching edges between a set of nodes labeled with words or concepts. The edges are understood to be a representation of the semantic relations between the words or concepts.[28]

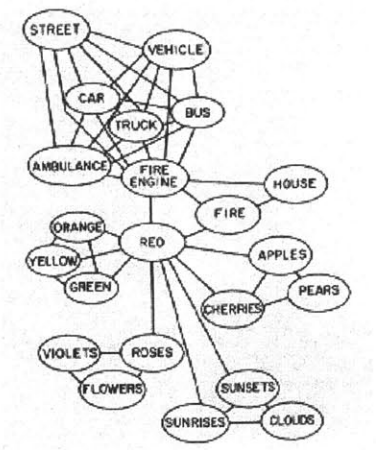

Social networks can be drawn in a similar manner although, obviously, for very different representational purposes. The nodes in a social network are often labeled with the names of people. The edges between nodes represent, for instance, social interactions between the people.[29]

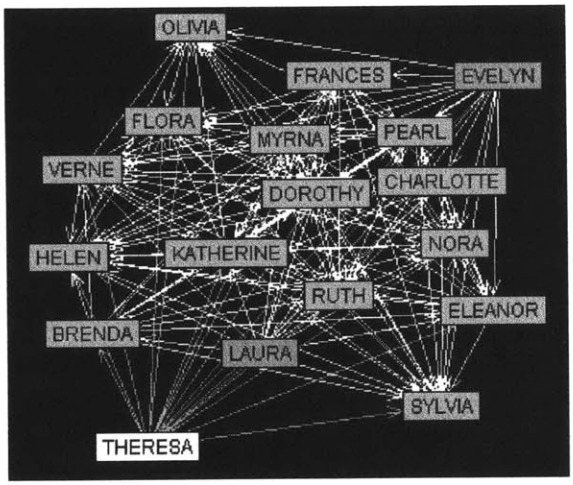

[28] Semantic network from A.M. Collins and E.F. Loftus "A Spreading Activation Theory of Semantic Processing," *Psychological Review, 82*, 1975: 407-428.

[29] Social network is produced from data obtained in the 1930s in a study by Davis, Gardner and Gardner and plotted by software written by Michael Chan at the University of Illinois. The software can be found here: http://carnap.ss.uci.edu/java/illinois/davis.html

Social and Semantic Networks as Analytic and Generative Devices

In the social sciences, social and semantic networks can be used and discussed as *analytic devices*, as representations of observed or hypothesized phenomenon, i.e., as scientific models. They are, in short, proposed as answers to outstanding scientific questions. For instance, in cognitive psychology, linguistics, and artificial intelligence, semantic networks have been understood to be an answer to this question: "What constitutes a reasonable view of how semantic information is organized within a person's memory?"[30] Within sociology, social networks are used to model and summarize empirical studies of interactions and relationships between people, groups, and institutions.

In architecture the diagram is historically understood in two ways: as an explanatory or analytic device and as a generative device.
Peter Eisenman[31]

In addition to their use as analytic, scientific representational devices, I am proposing that social and semantic networks be designed as *generative devices*; i.e., as interfaces with which one can embark upon an exploration of VLSCs. This use of maps as generative -- as well as analytic -- devices is relatively well understood in other design disciplines as is illustrated by the quote from Peter Eisenman on the role of diagrams in architecture. Generative devices, as they are to be discussed in this chapter, serve as possible conversational resources: they provide a set of possible conversational foci and a means to reflect upon and navigate a conversation. Such devices are generative insofar as they are evocative objects meant to engender discussion or reflection. One can think of this generative use of social and semantic networks as an interface into the archive of a discussion instead of (or in addition to) analytic summaries of a discussion.

Here is an outline of the remainder of this chapter: First, several examples of semantic and social networks used as interface elements are shown. These examples are computed and mapped by an implemented Usenet newsgroup browser – the *Conversation Map* system -- that allows one to explore large archives of newsgroup messages. Second, after describing how social and semantic networks are automatically generated, mapped, and incorporated into the Conversation Map system, several design criteria for conversation maps will be sketched out.

---

[30] Quillian, 80.
[31] Peter Eisenman "Diagram: An Original Scene of Writing" in *Architecture New York* (23:27, 1998).

Main Screen

The image shown below was produced by the Conversation Map system after an analysis of over 1200 messages from about 260 participants in 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. The following explanations of the interface will use images from the analysis of this newsgroup as an example. However, it should be clear that this is only one example. The Conversation Map system can be run on the message archive of any newsgroup concerning any topic and will produce a unique interface image for each and every newsgroup archive.

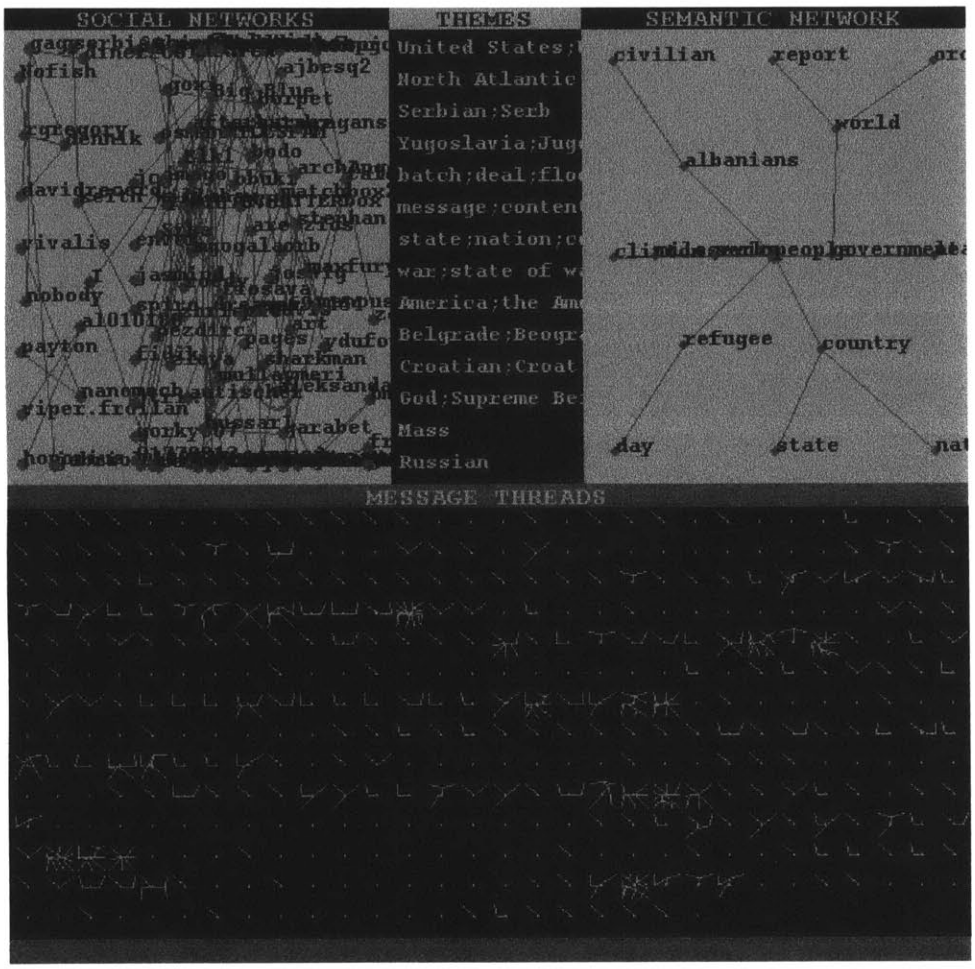

The interface is divided in two pieces:

(1) below the gray line labeled "message threads" is a graphical representation of the newsgroup messages that have been analyzed by the Conversation Map system;

(2) above the gray line is a display of the three-part analysis:
   (a) social networks;
   (b) discussion themes; and,
   (c) a semantic network.

All parts of the interface are interconnected with the other parts of the interface, so clicking on one part will highlight parts of it, but will also highlight other parts of the interface too. To explain the parts and their interconnections, each part of the image shown above will be explained.

Social Networks

The following figure shows an enlarged view of the social networks drawn in the upper left-hand corner of the main screen rendered in the previous figure.

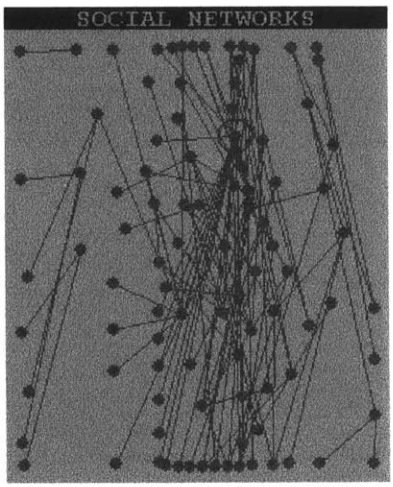

By positioning the mouse over the social networks panel and then pushing the right mouse button, the names labeling the nodes of the social network can be turned off (as has been done in the figure above). With the names off, it becomes easier to see that some participants are central to the newsgroup discussion and others are more marginal. The nodes with many connections represent participants who are both responding to and being responded to by many other participants. In other words, *reciprocity* is highlighted in the computed social networks.

Note that there are certain "hubs" in the social network. These hubs represent participants who post many messages but who also receive many responses to their messages. They are virtual moderators of a sort for the newsgroup even though the group depicted has no official moderators. The layout algorithm used tends to push the central participants to the center.

By simultaneously holding down the Shift key and the mouse button one can drag the nodes of the social networks around and get a better feel for the connectivity of various portions of the networks. To turn the names labeling the

nodes of the networks on again, press the right mouse button again (or, simultaneously the Meta key and mouse button if you have a one-button mouse).

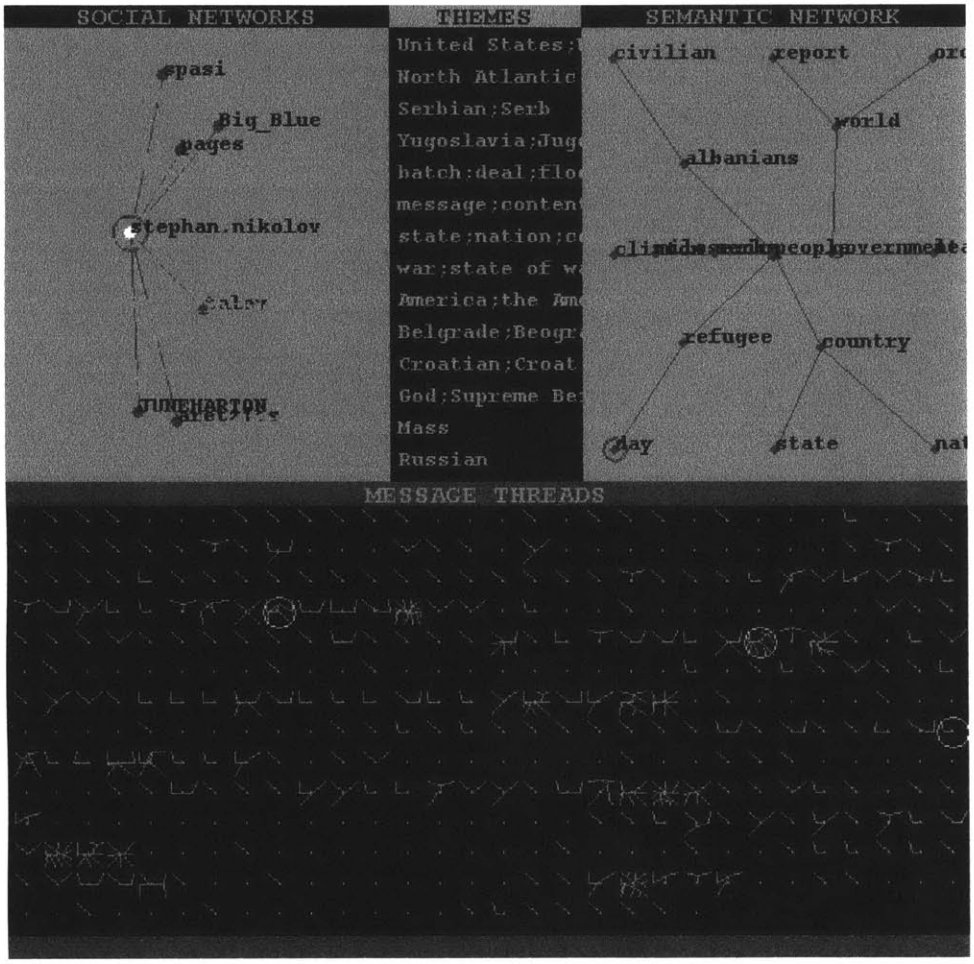

If one clicks the mouse button over one of the nodes in the networks, a small portion of a network is highlighted and the rest of the social networks disappear. The node selected (representing one participant in the newsgroup) and all the nodes linked to it are highlighted. At the same time, all of the threads in the archive are highlighted (with a light gray oval) in which the selected participant posted one or more messages.

By holding down the Control key and simultaneously clicking the mouse button, a second participant in the social networks can be selected. The edge between the two selected participants is highlighted, the threads where the two exchanged messages (and/or citations) are highlighted (in the case shown below, only one thread is highlighted), and, also, the discussion themes apropos of the messages exchanged by the pair are highlighted in the themes menu (in this case, two themes are visibly highlighted: the posters sent messages and/or quoted one another on the subject of the North Atlantic Treaty Organization (NATO) and the subject of war).

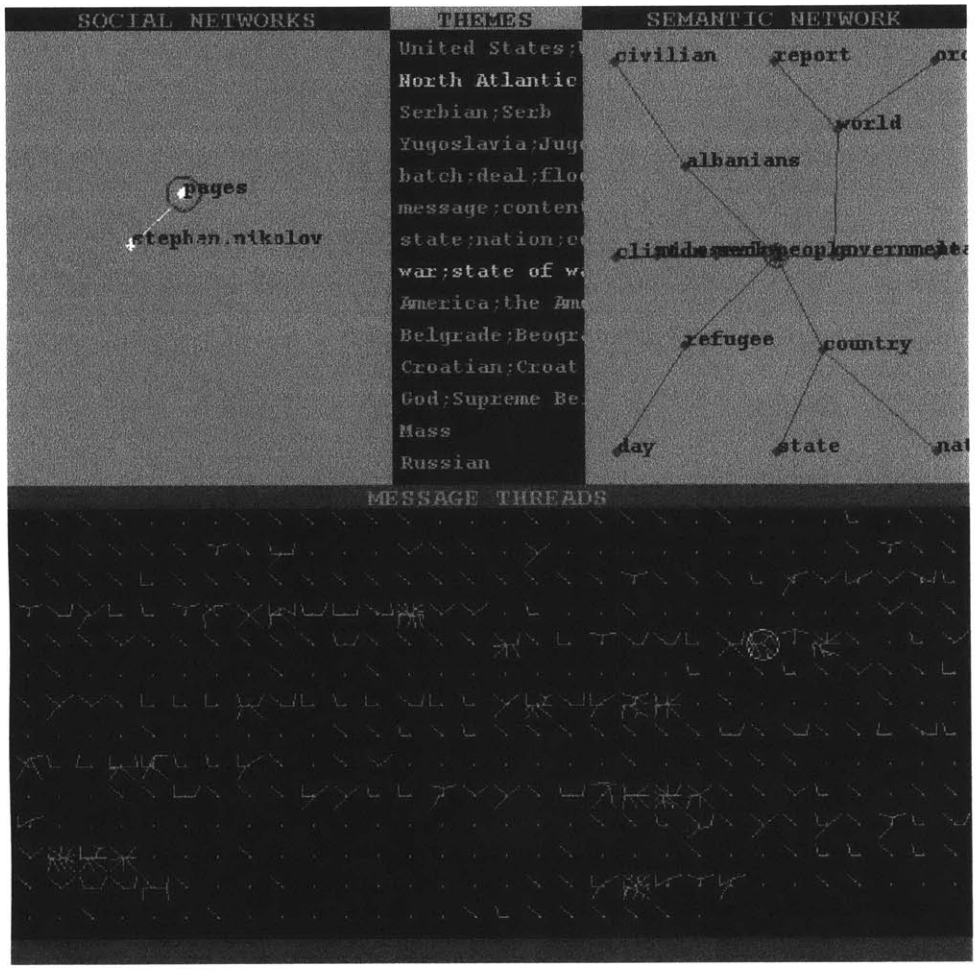

To make all of the social networks reappear, hold down the Alt key and click the mouse button.

Discussion Themes

Following is a picture of the same social network shown in the previous figure along with the menu of discussion themes that link messages, and thus, people together in conversation about the larger topic of Kosovo and Albanian culture in general. The "NATO" item in the themes menu has been highlighted by clicking on it with the mouse. The figure shows which pairs of posters have exchanged messages concerning NATO. Again, the unhighlighted portions of the social networks disappear from view and the portions of the archive where NATO connects two or more messages together are highlighted in the lower portion of the interface.

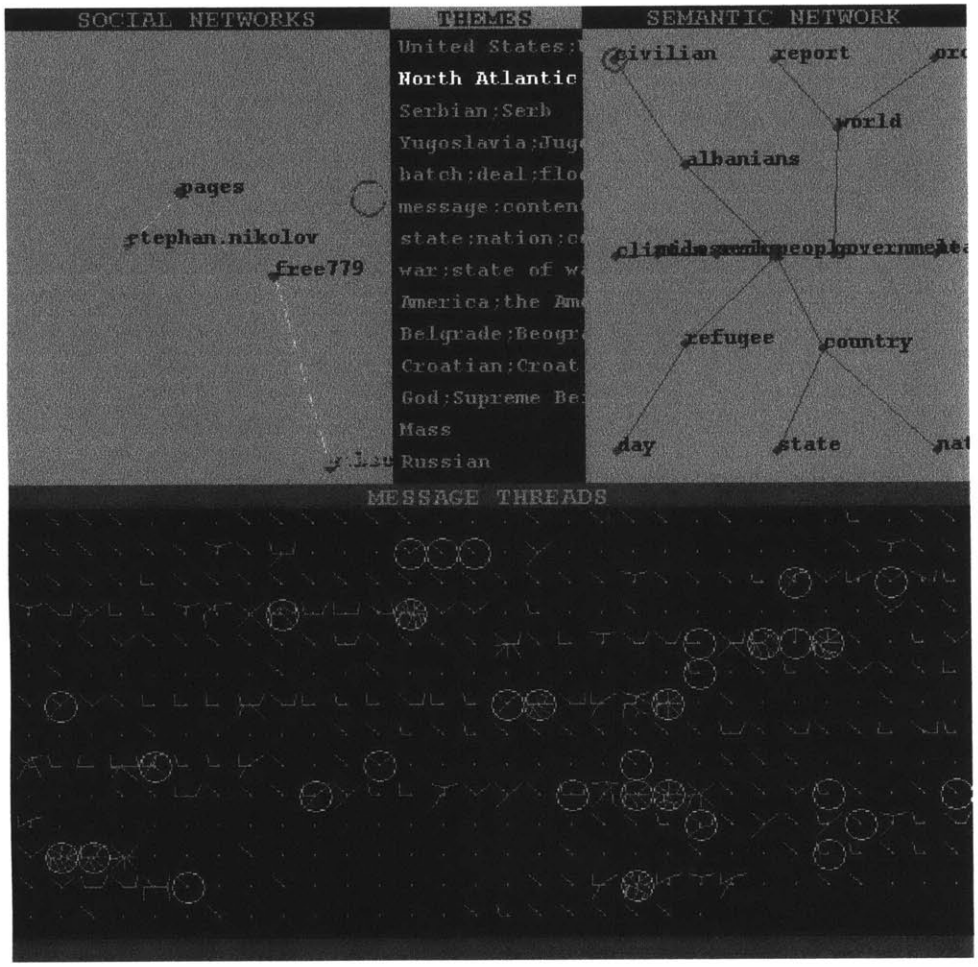

Note that only two pairs of posters seem to have exchanged messages about NATO, but many threads in the archive use NATO as a lexical tie between messages. It is probably not the case that the four participants highlighted in the social networks are responsible for all of the threads concerning NATO. Rather, it must be kept in mind that a pair of posters is highlighted if and only if they have a two-way, back-and-forth exchange involving a given theme while, in contrast, the criteria for highlighting a thread in the archive is less rigorous: a thread in the archive is highlighted for a given theme if the theme connects even one pair of messages in the thread.

Themes in the menu are listed according to the number of pairs of participants they connect in the social network. Thus "United States" is listed above "NATO" because "NATO" links only two pairs of posters while "United States" links three pairs. All of the themes down to "war; state of war; warfare" link two pairs; "America; the Americas" links one pair as do the rest of the following items in the menu.

Clicking on a theme is equivalent to searching the message threads, but the search performed differs from a conventional keyword search. A keyword search

would find, for instance, every mention of the term "NATO." In contrast, the theme search criteria are more rigorous. The theme search criteria are only fulfilled if, for instance, "NATO" is mentioned in one message of the thread and then again in a response or quoting message later in the thread.

Multiple themes can be selected by holding down the Control key while pushing the mouse button. The menu can be scrolled down by simultaneously holding down the Shift key and dragging the mouse. All highlighted themes can be simultaneously unhighlighted by holding down the Alt key and clicking the mouse button.

The Messages

Threads in a newsgroup discussion consist of an initial message concerning some subject, a set of responses to the initial message, a set of responses to the responses, and so forth. Therefore, conceptually, a thread is a "tree" in which the initial message is the "root" and links between responses are the "branches" of the "tree." Graphically, a thread tree can be plotted as a "spider web" in which the initial post is placed in the middle, the responses to the initial post are plotted in a circle around the initial post, the responses to the responses are plotted in a circle around the responses, etc. One of the nice features of plotting the thread trees as "spider webs" is that, at least in theory, any size tree can be plotted within a given amount of space and more-or-less uniformly fill that space.

In the bottom half of the figure below, over 400 threads are plotted as spider webs constrained into rectangular (rather than circular) spaces. The threads are arranged chronologically from upper-left to lower-right. By passing the mouse over each thread, the start and end dates and the subject lines of each thread can be read in turn in light gray text written into the dark gray strip at the bottom of the interface.

Since each thread is allotted the same amount of screen space, a rough guide to newsgroup activity can be read off of the panel in which all of the threads are plotted. If a thread without many messages is plotted, the rectangle containing it in the panel appears as mostly black. Threads containing many messages, and thus a lot of activity, appear very green.

In the figure immediately above the threads in the archive where "NATO" is a theme of discussion are highlighted. It can be seen that the "NATO" theme is discussed in some of the busiest -- i.e., largest -- threads of the newsgroup archive.

In the figure below, one thread from the archive has been selected with a mouse click. The thread selected has a white oval drawn around it. Note also that the dates when the messages of the thread were posted (27 April 1999 - 1 May 1999) and the subject line of the first message in the thread is printed in the dark

gray strip at the bottom of the interface: "Re: Response to: European trouble from a bird eye." In addition, parts of the social network, the themes menu, and the semantic network have also been highlighted. In the social network, those participants who are part of the social networks and who also have posted to the selected thread are highlighted. In the themes menu, those themes which connect two or more messages in the selected thread are highlighted. In the semantic network, those terms which correspond to the highlighted themes are also highlighted. The connection between the themes and the terms in the semantic network will be more fully explained in the section below devoted to the semantic network.

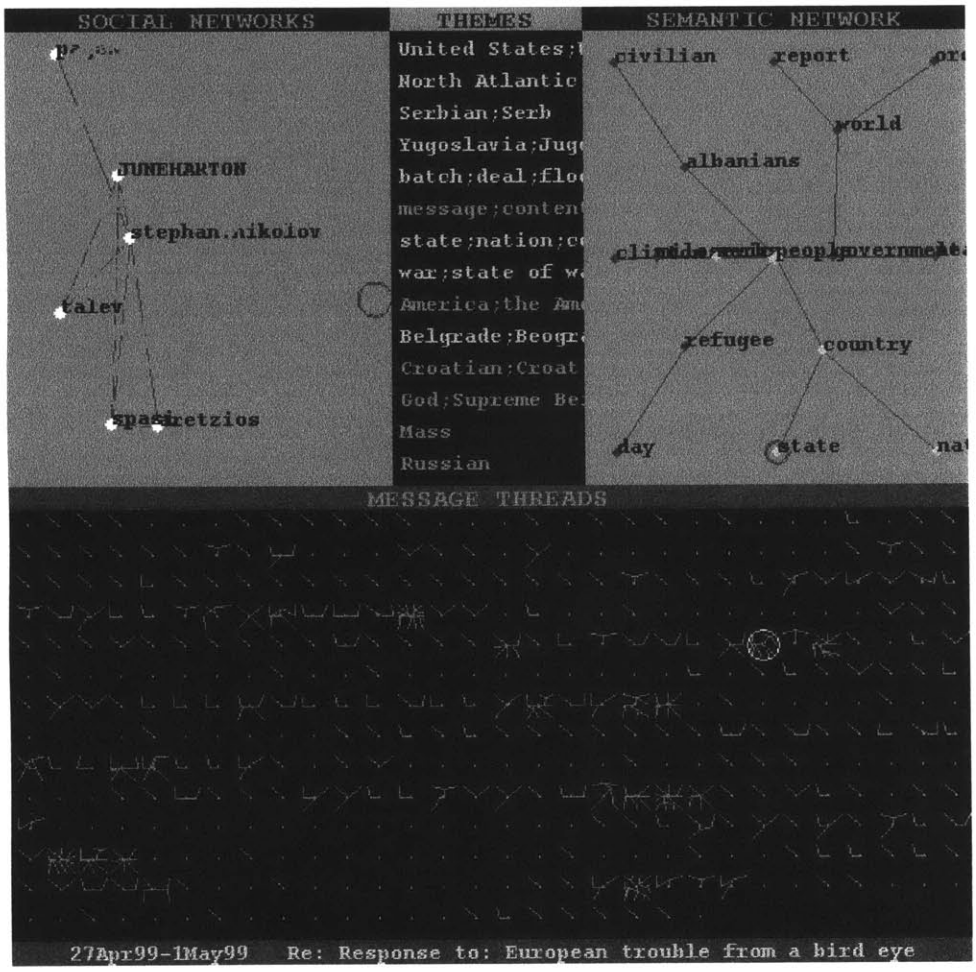

Message Threads

A thread can be opened and explored by double-clicking on it with the mouse. A double-click opens a separate window containing a larger version of the graphical display of the thread. The following figure illustrates an opened thread.

Normally, the nodes of a thread (representing messages in the thread) would be labeled with the names of the participants who posted them. In the figure above, however, the names have been turned off (using the right mouse button or Meta-click combination). In addition, some of the nodes of the thread have been moved around (by holding down the Shift key and dragging the mouse).

The spider web shape of the thread tree can be seen. If the thread was perfectly balanced (i.e., if each message had exactly the same number of responses as every other message), then the graphical plot of thread would more closely resemble a symmetrical web. However, a symmetrical shape is more the exception than the rule. The initial message of the thread is plotted as the largest green node in the center. In the thread shown above, the discussion theme "Croatia" has been highlighted. The menu of discussion themes can be scrolled by holding down the Shift key and dragging the mouse. By clicking on a discussion theme in the menu of themes, it is highlighted in white and the portion of the thread in which it is used as a theme is also highlighted in white. In this case, it can be seen that three of the messages of the thread are connected together by the theme "Croatia."

Message Display

In the thread shown above, a white circle around one of the nodes shows the position of the mouse. If the mouse is clicked, the text of the message (represented by the circled node) is displayed in a separate window.

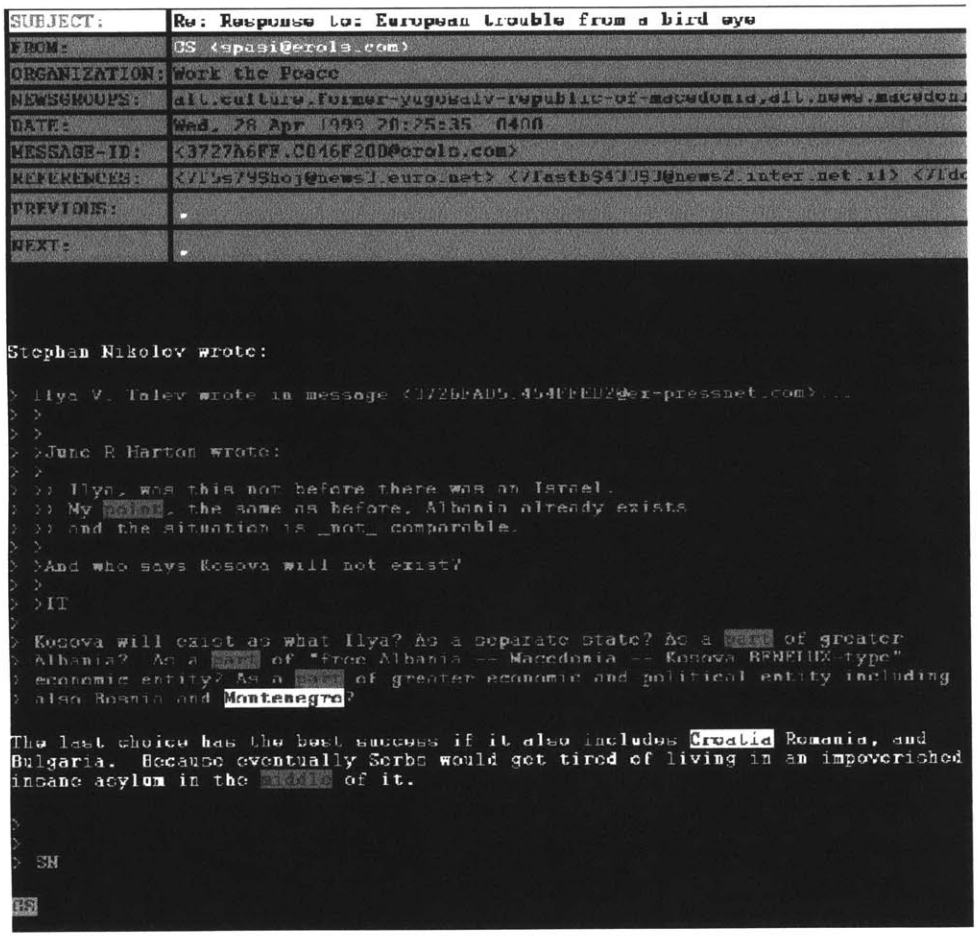

The use of "Croatia" as a discussion theme that links two of the messages of the thread is visible in the display of the message shown above. "Montenegro" is mentioned in a quote from a previous message and "Croatia" is discussed in the present message. The discussion themes analysis procedure of the Conversation Map system connected these two terms together because, in the thesaurus used in the Conversation Map system (i.e., Wordnet version 1.6[32]), Montenegro is listed as a part of Croatia. The text of the message displayed above also illustrates two other features of the Conversation Map system as a Usenet newsgroup browser: (1) Since quotations within messages are identified as a part of the analysis procedure for building the social networks, quotations within a given message are automatically highlighted as hypertext within the

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

display of the text of the message. Clicking on the text of a quotation will open a new window containing the full text of the quoted message. (2) Near the top of every message is a PREVIOUS and a NEXT label. If there is a • symbol listed next to the PREVIOUS label, clicking on the • will open a window containing the text of the message that precedes the current message. A message, A, is said to precede another message, B, if B is sent in reply to A. Since several messages might be sent in reply to a message, one or more •s might appear after a NEXT label. Click on each of the •s listed after the NEXT label to see all of the messages sent in response to the current message.

Semantic Network

The upper right-hand corner of the main screen of the interface displays a semantic network. In the semantic network, if two terms are connected together, then they have been calculated to have been "talked about" in similar ways in the archive of newsgroup messages.

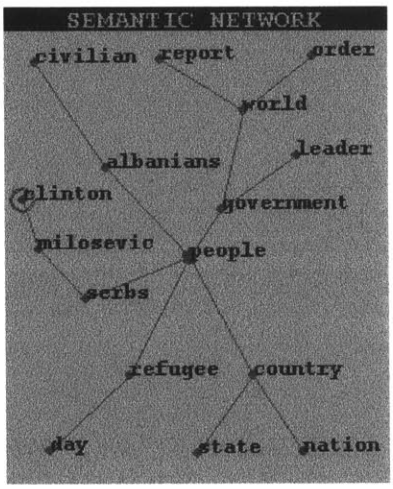

The central terms of a discussion are often connected to two or more other terms. Thus, in the soc.culture.albanian archive "people" is computed to be a central, perhaps neutral, term is the vicious argumentation that characterizes the content of many of the messages in the example archive. In this archive Albanians are "talked about" as people, Serbs are talked about as people, refugees are talked about as people, as are governments and countries. In other words, it appears to be the case that all sides of the argument (which is predominantly an argument pitting the Albanian view of the Kosovo situation against the Serbian view) can agree that the more general term "people" is applicable to both Serbs and Albanians.

The graphical interface uses the same spider web algorithm to lay out the semantic network as it uses to display the thread trees. Note that the algorithm sometimes overlaps nodes of the graph. In the figure above, the nodes of the

semantic network have been rearranged for legibility by holding down the Shift key and dragging the mouse.

Nodes of the semantic network can be selected by clicking the mouse. For example, if the term "country" is selected, all of the themes synonymous with country are highlighted in the themes. Simultaneously, all of the participants in the social network connected by the highlighted themes are also highlighted, and all of the threads wherein "country," or a synonym of country is used as a discussion theme are also highlighted.

To better understand why a given term appears where it does in the semantic network, double click on the term to see all of the associations it has in the archive of messages. Double-clicking on the term "country" produces a web page containing the following information.

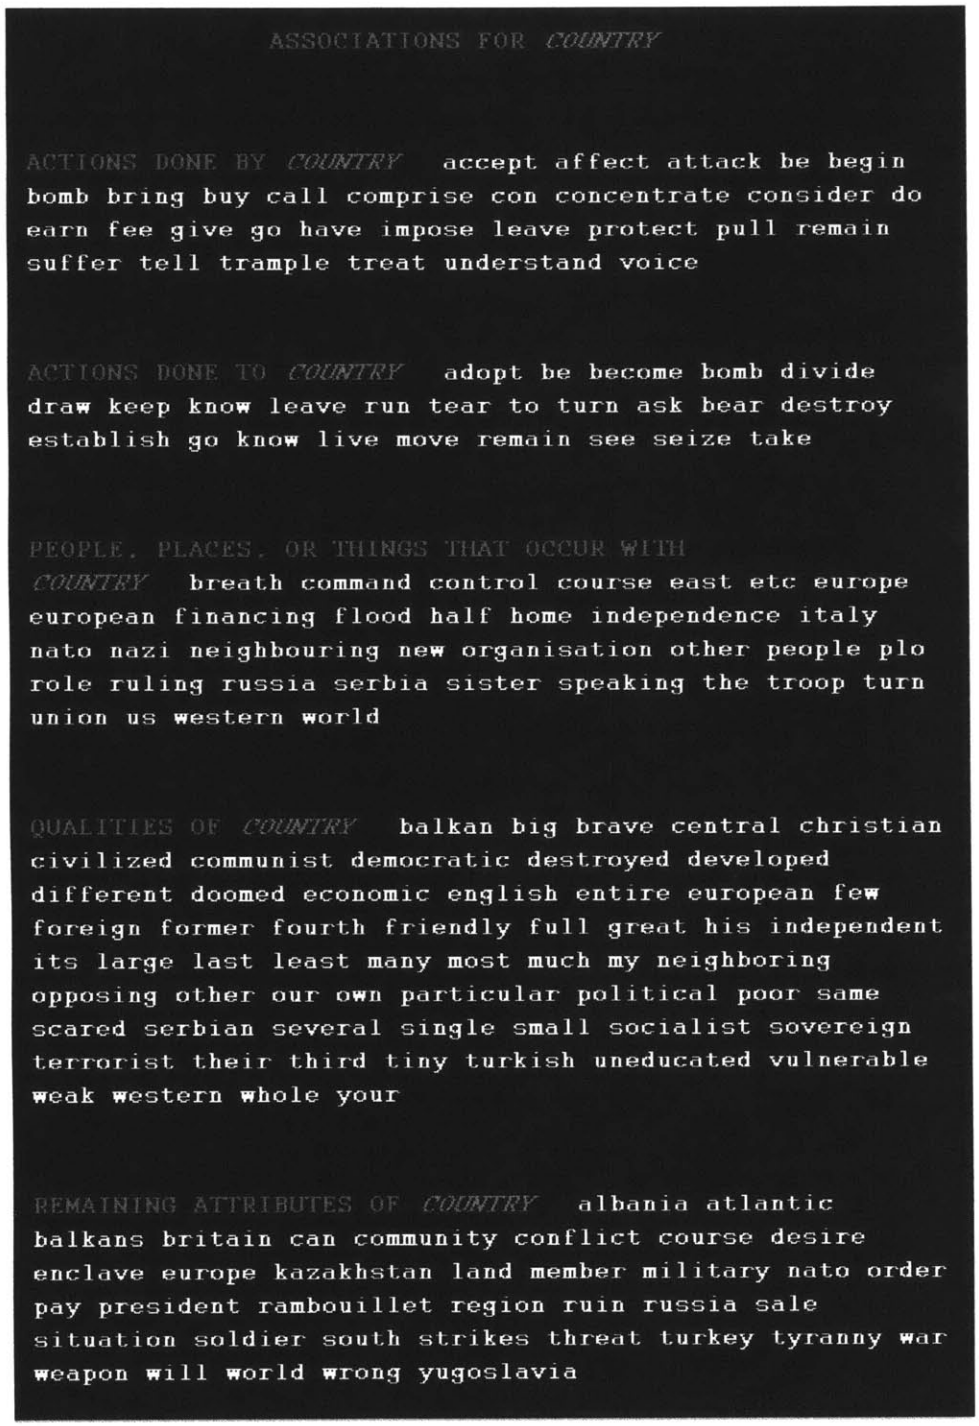

The associations displayed in the image above were calculated by the Conversation Map system.

The word associations that can be viewed by double-clicking on a term in the semantic network is a complete list of the verbs, adjectives, and nouns that are used with the given term. Each of the word associations can be "opened" with a single click. If the verb "consider" is clicked on from the display shown above, a

web browser window containing the following table appears. This table shows all of sentences in the archive of messages where the term "country" has appeared as the subject of the verb "consider." To see the message that contains an example sentence, click on the sentence and a new web browser window will be opened containing the text of the message.

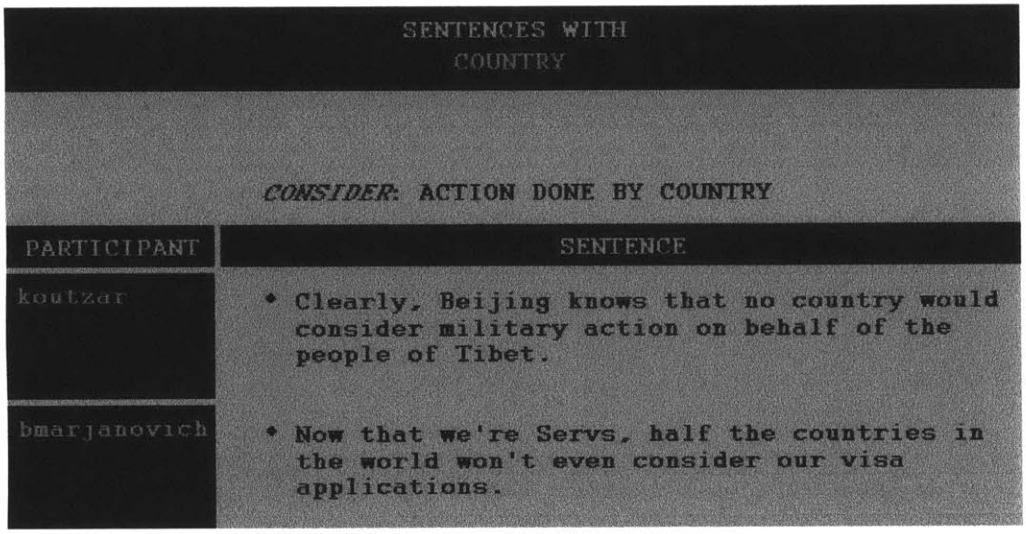

SENTENCES WITH COUNTRY

CONSIDER: ACTION DONE BY COUNTRY

PARTICIPANT	SENTENCE
koutzar	• Clearly, Beijing knows that no country would consider military action on behalf of the people of Tibet.
bmarjanovich	• Now that we're Servs, half the countries in the world won't even consider our visa applications.

It is also possible to compare the associations of one term with the associations of another term. Return to the main window displaying the semantic network. In the semantic network, hold down the Control key and click the mouse twice, once over the term "country" and then over "nation." Now, hold down the Control key again and move the mouse over one of the two selected terms, and double click the mouse.

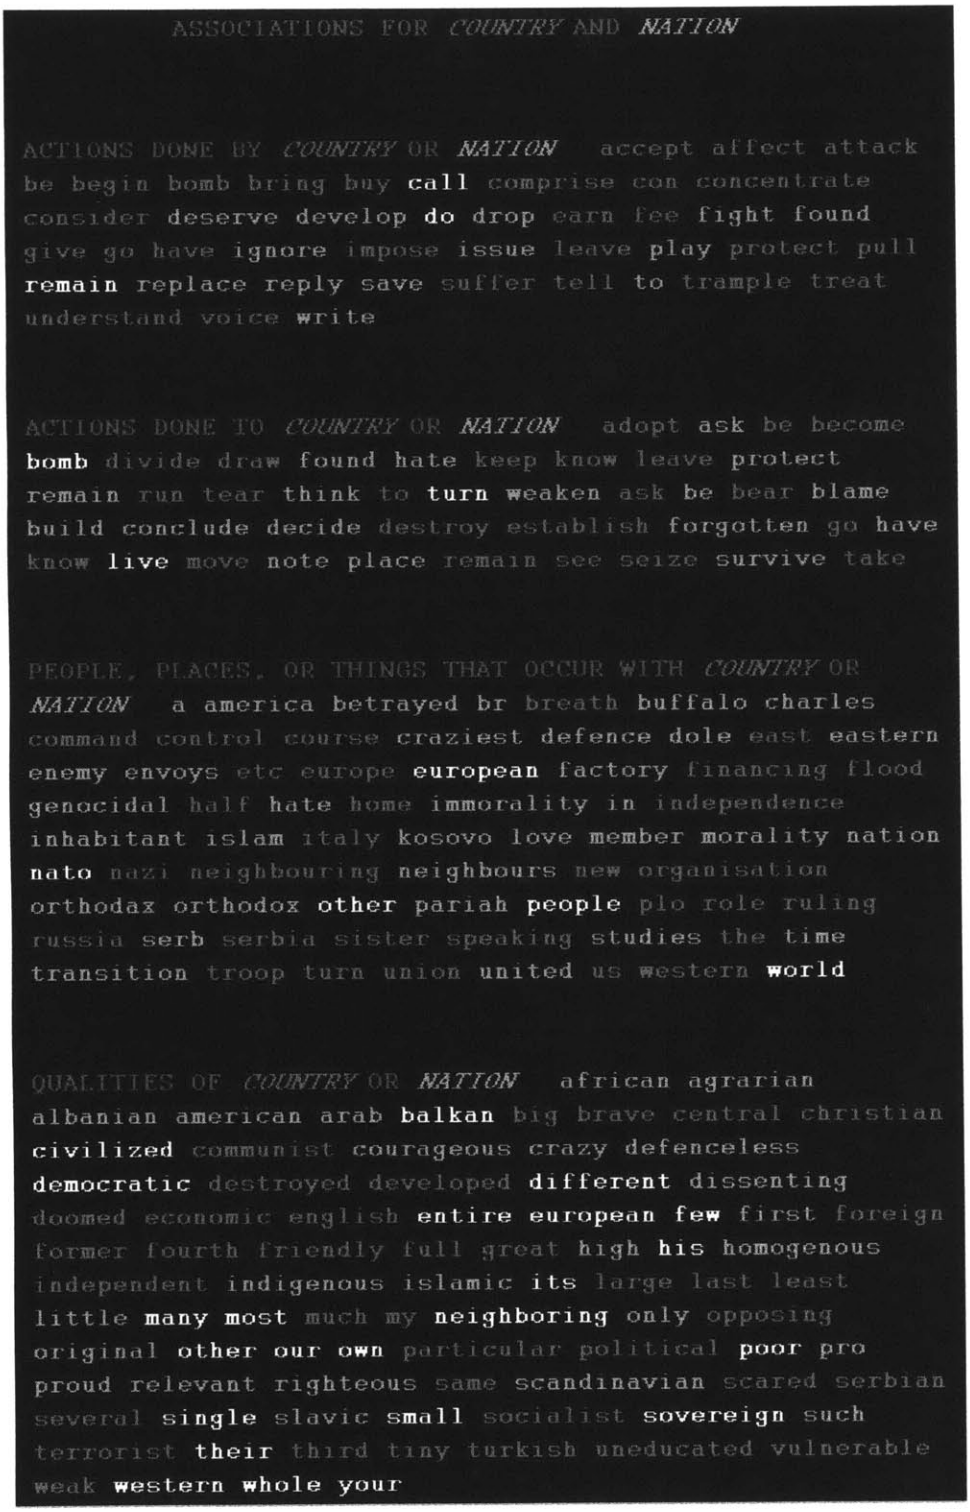

A new window is created. It displays the difference and union of the associations for "country" and "nation." Associations unique to "country" are displayed in green. Associations unique to "nation" are shown in silver. And, associations common to both "country" and "nation" are written in white. Clicking on any of the terms listed in green or silver will create a window of example sentences like

the window shown above for the examples of "country" used as the subject of the verb "consider." If any of the white terms are clicked on, a similar window of examples will be created containing sentences using the term "country" and other sentences using the term "nation."

By comparing terms' intersecting associations it is possible to begin to explore questions like these: In this conversation, how are countries like nations, people like countries, or Serbs like Albanians?

Conclusions: Design Criteria

There are several design criteria that have been used to construct and display the social and semantic networks of VLSCs with the Conversation Map system. These criteria negotiate a divergence between a social scientific use of the social and semantic networks of the Conversation Map and a possible, popular, non-scientific use of the same.

Ideally, for the sake of science, the system-generated, social and semantic networks would be constructed as carefully and rigorously as the "hand built" analyses of, for instance, ethnographic or sociolinguistic studies of online discussions.[33] However, clearly, the system-generated networks will never be as precise as analyses accomplished by hand. Thus, while, from a scientist's viewpoint, it might at first appear to be a good idea to attempt to automate much of the process of online conversation analysis, a closer look at the pragmatics of such a design goal shows it to not be such a good idea.

However, since the system-generated results are quicker and easier to attain than comparable results compiled by hand, the results of the Conversation Map system could serve different needs for scientists and non-scientists. For the scientist – e.g., for someone who is trained as a discourse or conversation analyst or an ethnographer – the maps produced by the Conversation Map system could be understood as a rough sketch of where one might begin to explore an archive of messages. With such a "sketch" in hand, the scientist could begin a set of more rigorous close readings of the archive of messages.

For the non-scientist, the conversation maps produced probably represent a much more detailed analysis of a VLSC than anything the non-scientist would ever produce on their own. However, there are pitfalls associated with non-scientific usages of scientific-looking images. Scientific images have always been open for interpretation and put to alternative uses by non-scientists; e.g., journalists, lawyers, politicians, interested laypersons, and non-specialists (e.g., scientists or doctors who are not from the discipline directly responsible for the

[33] E.g., Susan Herring, Deborah A. Johnson, Tamra DiBenedetto. "'This discussion is going too far!': Male resistance to female participation on the Internet," in *Gender Articulated: Language and the Socially Constructed Self* edited by K. Hall and M. Bucholtz (New York: Routledge, 1995).

production of a set of scientific images). For example, the anthropologist of science and technology, Joseph Dumit, has examined how PET scans of the brain are used and understood both within science and "outside" in medicine, law, journalism, and popular culture.[34]

The common pitfall associated with many vernacular presentations of scientific images concerns the manner in which the images are often "untethered" from the data used to produced them. Untethered scientific images -- i.e., images that have been unlinked from supporting data – sometimes become too easy to manipulate because they are no longer manipulated within the rigorous constraints of science; e.g., when a popular magazine recolors an image received from a biologist to make the image easier to print, or more colorful. At other times, these untethered images become too hard to manipulate because a layperson has no access to the phenomenon pictured; e.g., how would a non-scientist redraw the traces of subatomic particle collisions produced by physicists with an instrumented particle accelerator?

One crucial design question is therefore this: How can the images of science be used for interface devices without untethering the images from the supporting data? Furthermore, how can the social scientific images of semantic and social networks be rendered as *generative maps* for use as interface devices that cross-index and provide access to thousands of email messages from the archives of very large-scale conversations? The approach to this design problem taken in the construction of the Conversation Map system has been this: make all of the original data (i.e., the newsgroup messages) accessible through the maps (i.e., through the act of double-clicking on the maps).

A second tension that exists between possible scientific versus possible popular uses of the Conversation Map system is this: to automatically generate the conversation maps – the social and semantic networks – the Conversation Map produces an enormous amount of quantitative data on the messages and the participants of the newsgroup analyzed. Social scientists often want to see the numbers produced. For example, the Conversation Map system generates a set of statistics which could be useful to get at least preliminary answers to the following questions: What was the mean number of responses sent to a message? How many other participants did a given participant respond to? How many times did a given theme show up as a theme of discussion? What proportion of the population of participants contributed towards the discussion on a given topic? Is there a specific subset of participants who started most of the threads concerning a given theme of discussion?

As interfaces for VLSCs come to be more and more representative of the social structure of the conversing group, many in the group may feel that an interface display impinges on their privacy by rendering transparent the history of their

[34] Joseph Dumit *Whose Brain Is This? PET Scans and Personhood in Biomedical America* (forthcoming; see also, http://symptom.mit.edu).

interactions with the group. Even today, the poster profiles computed and indexed at sites like www.dejanews.com incite these worries for some posters. Preferably then the construction of conversation maps will follow an aesthetics of social translucence[35] and encourage an ethics of social reflection, rather than aiming at an aesthetics of realism and transparency that would make all of us feel we are under a microscope of surveillance. Displaying all of the statistics calculable by the Conversation Map system would probably render the participants' profiles too transparent for comfort.

To negotiate this tension between the need for numbers and the larger scientific and non-scientific need to better understand the linguistic and social structures of VLSCs, the *quantitative* results calculated by the Conversation Map system are displayed as *qualitative*, flexible maps. These maps can be moved and redrawn by the participant or interested observer, but they do not directly yield summary statistics on the group or individual participants of the group.

Finally, while it is often the case that analyses of conversation and discourse have been done by scientists for other scientists, it was a specific design choice to create an interface for the Conversation Map system that can be, at least in principle, accessed by everyone who might be participating in a public, Internet-mediated VLSC. Several Conversation Maps and a user's manual can currently be found on the web: http://www.media.mit.edu/~wsack/CM/index.html. To use the Conversation Map system, a newer, Java 1.2 enabled browser is necessary (e.g., Netscape 4.6 on Windows or Linux machines; Internet Explorer 4.5 on Macs).

[35] Thomas Erickson, David N. Smith, Wendy A. Kellogg, Mark Laff, John T. Richards, Erin Bradner "Socially Translucent Systems: Social Proxies, Persistent Conversation, and the Design of 'Babble'" in *Proceedings of CHI'99* (Pittsburg, PA, May 1999).

EXAMPLES

The online conversations of several large groups are examined through the graphics generated by the Conversation Map system. By considering the reflections of the VLSCs of (1) audiences, (2) social movements, (3) user groups, and (4) a group I will name "citizen diplomats," a clearer picture can be had of the structure of these new, network-based groups. For the social scientist, this is useful because it provides a place to begin to explore the social and semantic significance of these new groups. For an actual or potential participant, gaining a feel for the social and semantic structure of a VLSC is a means to decide whether the VLSC is worth joining or whether particular contributions to the group have been pivotal or marginal to the structure of the group.[36]

[36] Parts of this chapter appear in Warren Sack "Stories and Social Networks" in the *Proceedings of the American Association of Artificial Intelligence Workshop on Narrative Intelligence* (Cape Cod, MA: AAAI, November 1999).

To get a better feel for the navigational powers of the Conversation Map system it is helpful to compare the outputs of the system to the results obtained by social scientists and discourse analysts. To facilitate this sort of rough comparison, four group types have been selected: audiences, social movements, user groups, and citizen diplomats. Some of these groups – e.g., audiences – have been studied by social scientists and others. By looking at specific instances of online group interchange through the "lens" of the Conversation Map, we can get a feel for what the system misses that the social scientists (working "by hand") manage to see, and vice versa. Certain aspects of the Conversation Maps for these groups will be highlighted to show how these groups sometimes function (and at other times do not function) as very large-scale conversations. This simultaneously demonstrates the existence, limits, and productiveness of the concept of VLSC.

Audiences

Two archives of audience messages will be discussed in this section. Both archives contain messages posted to the Usenet newsgroup alt.tv.x-files, a group devoted to discussion of the internationally broadcast television show entitled *The X-files*. The Usenet newsgroup discussion is archived and publicly available at a variety of websites including, www.deja.com. The staff at Deja.com was kind enough to provide us with the two archives discussed here.

The X-files is a weekly show produced by Twentieth Century Television in association with Fox Broadcasting Company. The show has two main characters, FBI Agents Dana Scully and Fox Mulder (played by actors Gillian Anderson and David Duchovny, respectively), who investigate cases reported to involve extraterrestrials, paranormal phenomena, and government conspiracy. It is an award winning television show now in its sixth season. More information about the show and short descriptions of the episodes can be found at the official *X-files* website: http://www.thex-files.com/.

Audience Message Archive 1: These messages were exchanged during the week following the airing of the episode entitled "Quagmire" (4 May 1996 - 10 May 1996). In the "Quagmire" episode a Loch Ness monster-like creature is suspected of killing several people. About 700 participants posted over 1900 messages to the Usenet newsgroup alt.tv.x-files during this week after this episode was shown. A sketch of the analyzed messages from this archive can be seen in Figure 1.

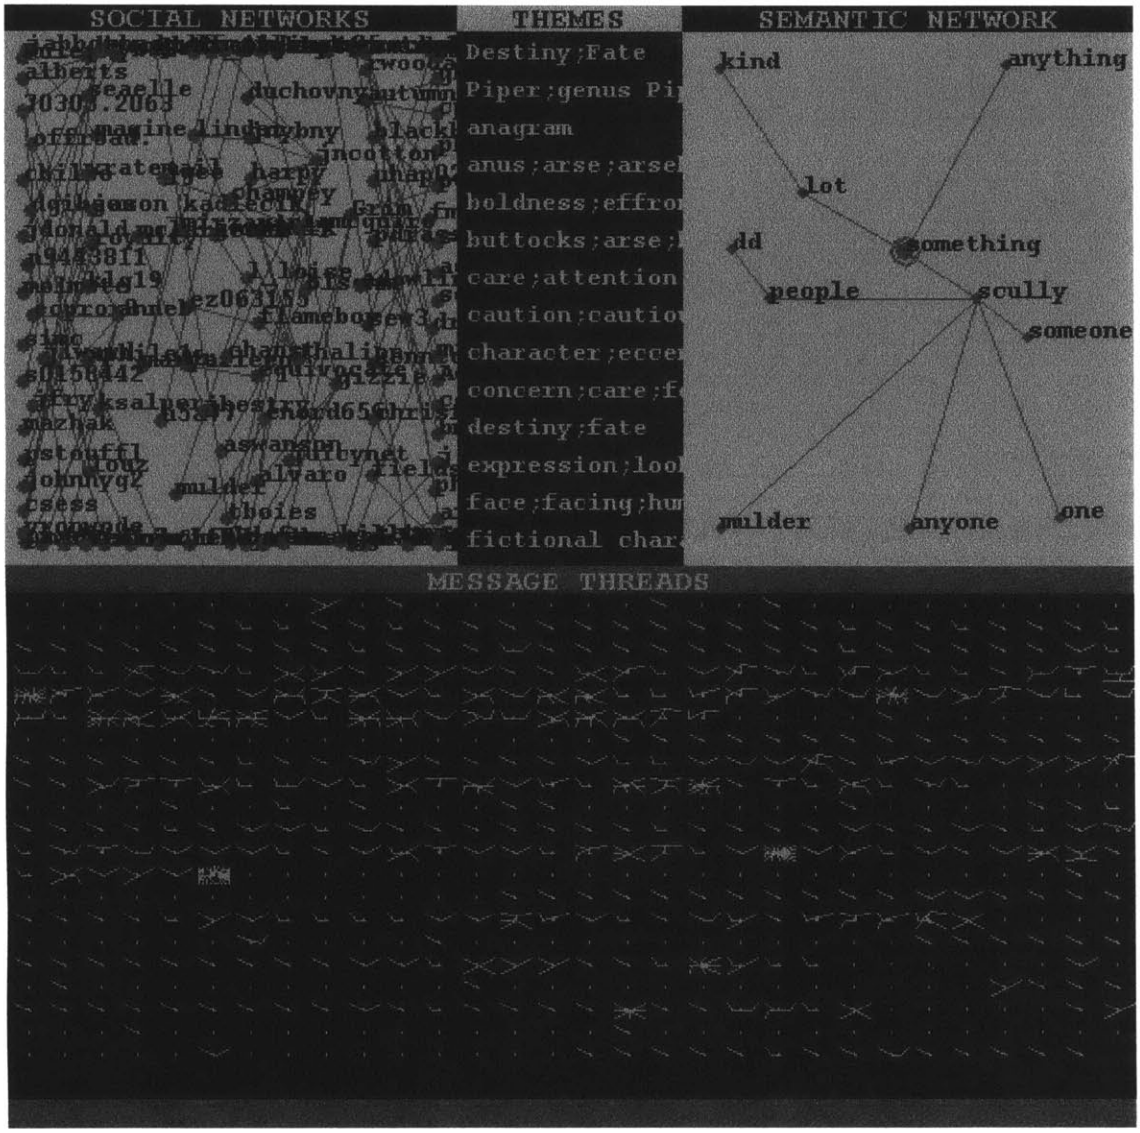

Figure 1: Conversation Map interface for Audience Archive 1

Audience Message Archive 2: These messages were exchanged during the week following the airing of the episode entitled "Hell Money" (30 March 1996 – 5 April 1996). The "Hell Money" episode concerns a high-stakes gambling game in which the players risk their own organs (e.g., their eyes and kidneys). Approximately 900 participants posted 2400 messages to the Usenet newsgroup after this episode. Figure 2 below shows the Conversation Map automatically generated from the analysis of messages posted that week.

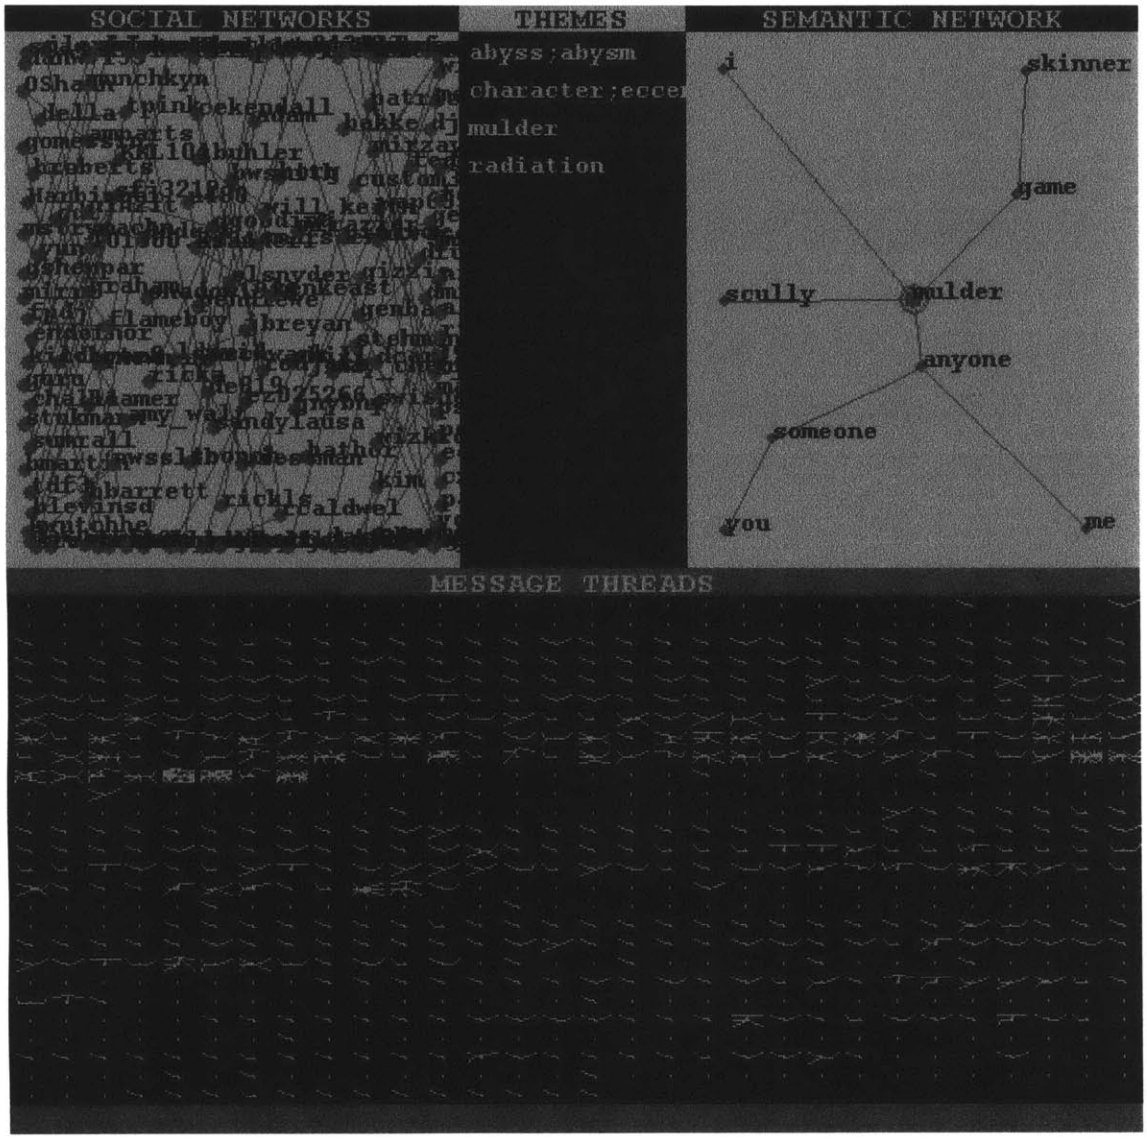

Figure 2: Conversation Map interface for Audience Archive 2

Before proceeding to a closer examination of the Conversation Maps, two points need to be made. Firstly, in many structuralist, formalist, and/or older Marxist-inspired analyses of media audiences, the audience member is often assumed to be a "cultural dupe." That is to say, it is assumed that a story delivered through the media (e.g., radio, television, the Internet, etc.) is not really open to interpretation and/or appropriation and means, more or less, one -- and only one -- thing. Moreover, the one and only meaning of the story is exactly what the audience member receives and, in this reception, is seen to be "programmed" by the story to behave or think in a certain manner by the story. This description is an over simplification, but it underlies the heat generated in arguments over which literatures should or should not be taught in schools (i.e., the debate over the so-called "canon") and also the anger over permissiveness versus censorship of the content of television, movies, and the Internet.

On the other end of the realist-to-relativist spectrum are many post-structuralist and cultural studies-inspired media scholars who have tended to emphasize the extraordinary creativity of audience members. Stories, and media productions in general, are seen as raw materials for audience members to rewrite, reinterpret, and recreate in novel and undetermined ways.

By spending some time with the Conversation Maps of audiences' online conversations, it should become clear that neither of these idealisms is empirically supported. On the one hand, the range of responses to the television stories is very diverse both in content and in genre. The "genres" of response include these: some responses are close intertextual analyses of the plot and characters of the episode, others are simple questions (e.g., "What's your favorite X-files episode?"), others are wildly tangential (e.g., "I have two kittens, one named Mulder, the other Scully, and I'm looking for someone to adopt them..."). On the other hand, only someone who is very easily amused will be likely to see the messages contained in these archives as wildly creative.

Thus, as a first point, I maintain that a machine-assisted, empirical examination of audience conversation makes it quite easy to resolve an issue that is often a point of debate in narrative theory and media studies: audience members are not "cultural dupes," but, neither are they more likely than any of the rest of us to be wildly creative with the "raw material" of the stories seen, heard, or read.

The second point also concerns the computational form of the analyses presented here. It has often been the case that audience studies have been formulated and written in a specialist's language (e.g., the vocabulary of academic media studies) and presented in a medium unlike the medium of the story and unlike the media used by the audience members to communicate amongst themselves (e.g., studies of television audiences are oftentimes written up as academic books). For Internet-based audiences, it is now possible to build technologies that are designed to be accessible to the audience members and specialists alike. The Conversation Map system has been designed to be available online accessible to anyone with a web browser.

My second preliminary point is this: audience-accessible, networked, media studies cannot – as previous work repeatedly has – treat audiences as commodities or scientific objects because the network provides a means for the audience members to dispute the interpretations offered by the specialists. Consequently, what is presented below can best be understood as one place to begin an examination of the audiences' understandings of the two X-files episodes, and not as a definite, final discovery of those understandings.

Two Audience Analyses

In what follows, the social networks, themes, and semantic networks displayed in the Conversation Maps of the two message archives will be more closely

examined.

Audience Social Networks

Figures 3 and 4 are enlargements of the social networks visible in Figures 1 and 2 respectively. In Figures 3 and 4 the names of the newsgroup participants have been turned off to allow one to see the topology of the networks more clearly.

What should be clear in Figures 3 and 4 is that participants are grouped into many small networks. The small networks are not connected to one another although it can be seen that the social networks shown in Figure 3 are more highly connected than the networks shown in Figure 4. In Figure 3, for example, the circled participant is a "linchpin" of sorts holding together several smaller networks.

The lack of connections in the social networks is interesting because a quick glance at them makes it immediately apparent that the newsgroup is a space in which many different, probably unrelated, conversations are happening. Obviously the "effects" of a television story do not automatically include the straightforward production of a cohesive social order.

Figure 4 especially illustrates the possible use of the Conversation Map system to spot non-conversations. Participants in this group (i.e., the discussion after the X-files' "Hell Money" episode") are "speaking" in the same space (i.e., the alt.tv.x-files Usenet newsgroup), but they are mostly "talking at" rather than "talking to" one another. The highly fractured nature of the social networks in Figure 4 (but also, to a lesser degree, in Figure 3) show how few of the participants are commenting on or quoting one another's messages.

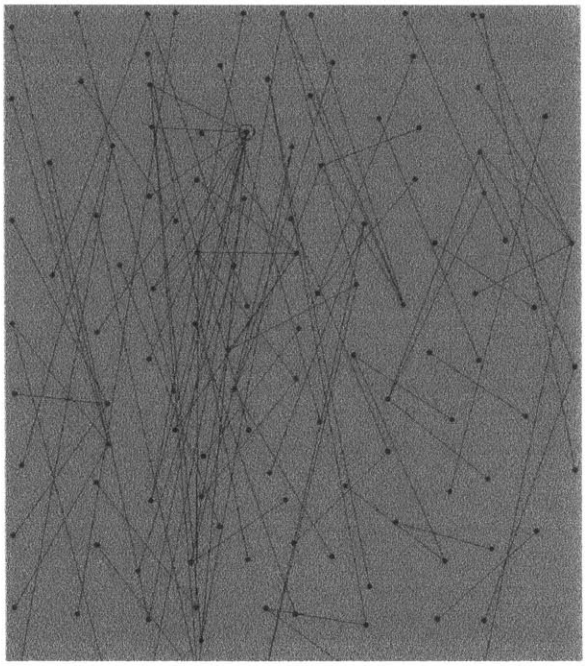

Figure 3: Social Network for Audience Archive 1

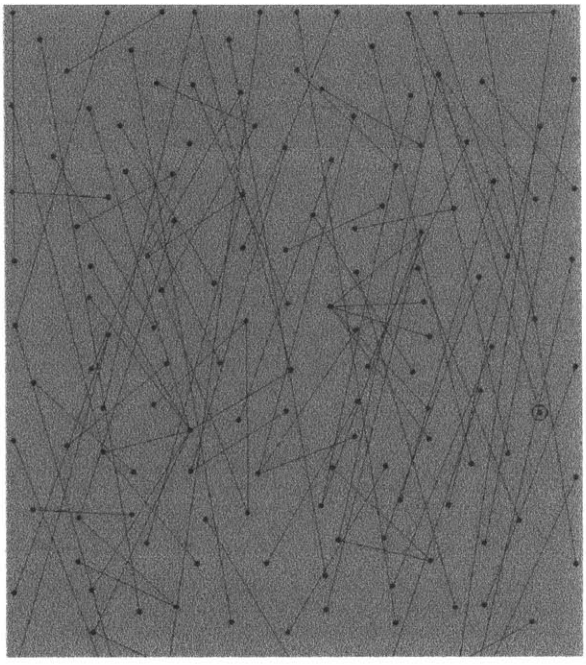

Figure 4: Social Network for Audience Archive 2

Audience Themes

Another measure of the diversity of conversation in a newsgroup is provided by the menu of computed "discussion themes" (i.e., what in linguistics would more properly be described as the *lexical ties* between messages). Figures 5 and 6 list the tops of the theme menus for message archives 1 and 2 respectively.

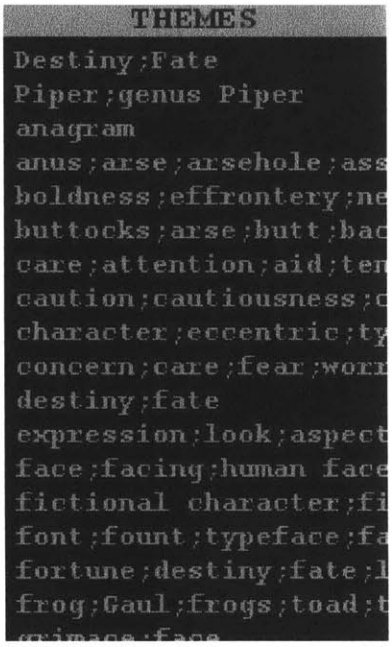

Figure 5: Themes Menu for Audience Archive 1

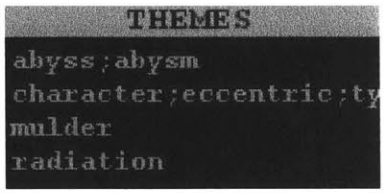

Figure 6: Themes Menu for Audience Archive 2

Themes in the menus of themes are ordered according to the number of pairs of posters in the social network that they connect.

Remember that an arc in the social network connects two newsgroup participants if and only if those two participants have replied to each other or cited from one another's' messages. Thus, for example, A and B are connected in the social network and the arc between A and B is labeled with a theme – e.g., "sports" – if and only if A and B have had at least one interchange like the following: A posts a message about baseball, B replies with a post about football, B posts a message about swimming, and A cites or responds to B's message with one about skiing. Since baseball, football, swimming, and skiing are all sports, the

link between A and B might be labeled with the more abstract term "sports" (computed by the Conversation Map system using the WordNet version 1.6 thesaurus). So, the themes listed in the menus are only there if there has been one or more reciprocated responses in which the theme (or a semantically similar) term was mentioned in each of the exchanged messages.

Figure 6, showing the reciprocated discussion themes in the messages of archive 2, is a rather surprisingly short list. Usually the menu of themes lists many items. Clicking on the items to highlight the parts of the social network that they label shows even more clearly how fragmented the discussion of archive 2 is. All of the themes listed connect only one pair of posters. In short, only a small handful of the interchanges concerning the "Hell Money" episode are focused around a specific theme of discussion. This is more evidence that the group, at this time, is not in conversation with one another but is, rather, just "talking at" one another.

Figure 5, showing the reciprocated discussion themes in the messages of archive 2, shows again that the social interchange visible in the message archives is more cohesive in the first archive than it is in the second archive. This can be interpreted from the longer list of reciprocated themes for archive 1.

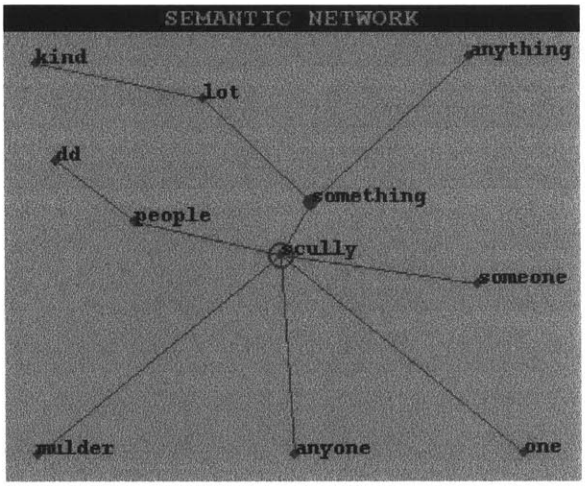

Figure 7: Semantic Network from Audience Archive 1

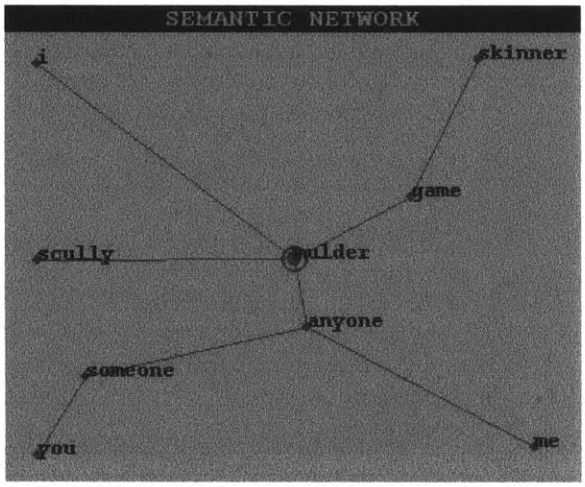

Figure 8: Semantic Network from Audience Archive 2

Audience Semantic Networks

The semantic networks shown in Figures 7 and 8 show that the conversations after both episodes are concerned with the main characters (Scully and Mulder) and, moreover, it is interesting to see the computed similarities between the main characters and the more generic terms of "you," "me," "someone," "anyone" etc. These calculations provide a way of seeing how the audience members talk about themselves in ways comparable to the way they talk about the main characters. This calculation might be compared to analyses of character "identification" discussed in the literatures of film theory and other media studies.

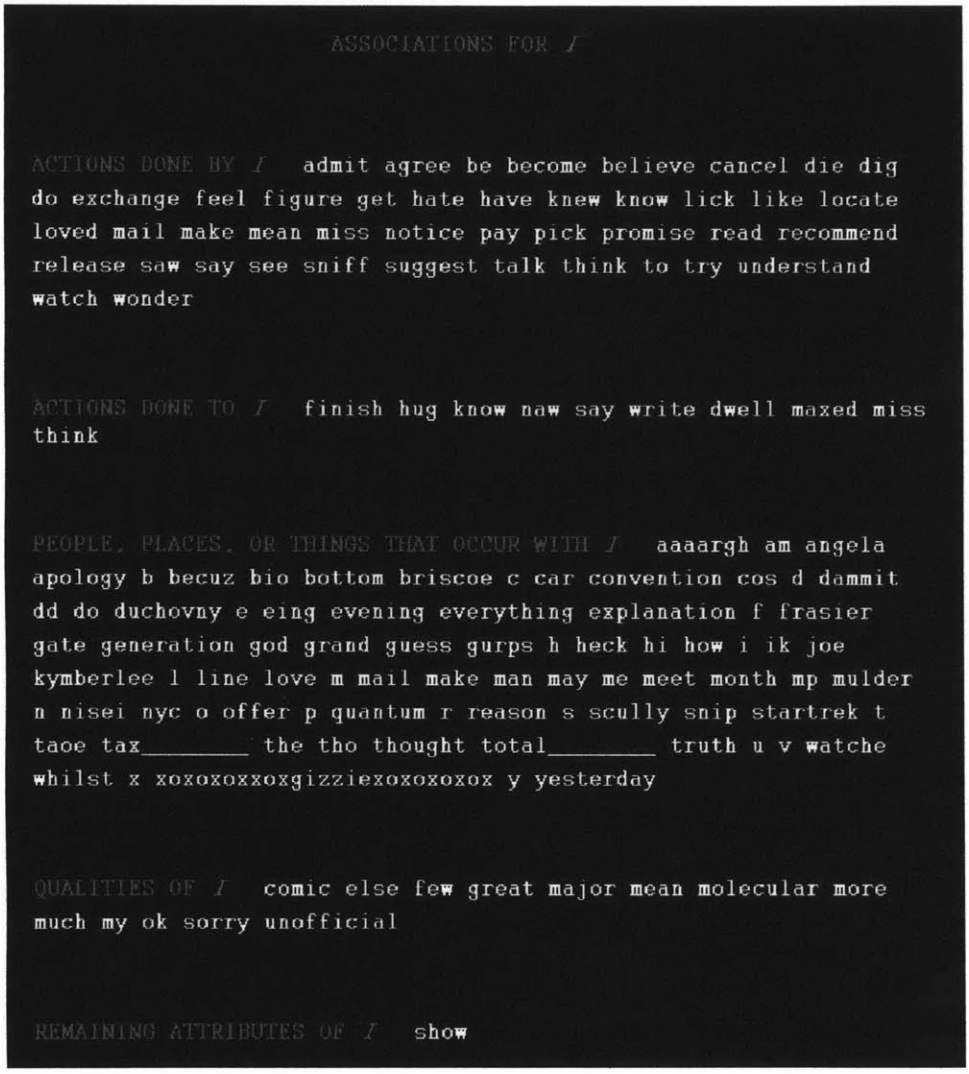

Figure 9: Associations for the pronoun "I" in Audience Archive 2

Audience Identification

Clicking on the pronoun "I" in the semantic network for Audience Archive 2, the information shown in Figure 10 is displayed in a web page. This web page displays all of the verbs, adjectives, etc. that are associated with "I" in the messages of Audience Archive 2.[37] These associations can be compared to the associations for main characters of the show to see if the newsgroup participants are writing about themselves in way that is similar to the way they write about Scully and Mulder.

To find the overlapping associations using the Conversation Map system, hold down the Control key, click on the term "I" in the semantic network, move the mouse over the term "Scully" in the semantic network, and then double-click. A page like the one shown in Figure 10 will appear. (See the previous chapter on the Conversation Map interface for more detailed instructions.)

[37] Obviously, this list incorporates some parsing and/or authoring errors since the pronoun "I" should, grammatically speaking, never occur as the object of a sentence. If the parser was always correct – and the authors of the messages always wrote grammatically correct sentences -- then the list "Actions done to I" should be empty.

Figure 10: Intersecting and unique associations for “Scully” and “I”

If we restrict our attention to just the overlapping subject verbs (i.e., those verbs for which both the pronoun "I" and the name "Scully" appeared as a subject) we can pull out a list like the following from the first block of text in Figure 10:

"Scully" and "I" both are, believe, die, do, feel, get, have, know, lick, make, notice, pay, pick, read, say, see, talk, and think

This list of shared verbs seems rather meager. Likewise, if the associations of "I" and the character "Mulder" are intersected, a similarly thin list of shared verbs emerges:

"Mulder" and "I" both are, become, believe, do, figure, get, have, know, miss, pay, say, see, think, understand

Nevertheless, the point is that the Conversation Map provides a means to begin to explore the issue of audience identification with the characters in an empirical manner.

Thematic Roles

More interestingly for this corpus of messages (Audience Archive 2) is an exploration of what, in narrative theory, are called *thematic roles*. Thematic role is an idea coined by the narrative theorist A.J. Greimas.[38] Thematic roles are part of larger discourses which connect together many stories. For instance, a fisherman is an example of a thematic role. Associated with fishermen are a set of attributes and activities that are carried over from story to story and which remain the same regardless of how a fisherman interacts with others in a story. For example, fishermen usually have a boat and fishing gear, like nets, lines, and hooks, and they often can be found baiting, catching, and cleaning fish. Using this idea we can pose this question: do Scully or Mulder occupy any identifiable thematic roles in the audience discussions of the of X-files show?

First look again at the actions that are attributable to Scully in the audience discussions. In the messages of the newsgroup Scully can be said to accuse, act, adjust, address, answer, appraise, arrive, ask, assume, awaken, be, believe, bother, bring, call, come, complain, convince, cry, decide, deny, die, discover, do, drag, drive, drop, explain, face, fall, feel, find, get, give, go, handle, have, impress, investigate, involve, irradiate, keep, kill, know, launch, learn, leave, lick, liken, look, lose, make, name, need, notice, pay, pick, plead, point, protect, put, rationalize, react, read, realize, refuse, remember, rescue, reserve, roll, satisfy, say, scream, search, see, seem, shoot, show, sit, smile, spend, start, startle, state, stop, stumble, suffer, take, talk, tell, think, throw, tire, use, walk, want, wear, witness, write

---

[38] A.J. Greimas, *On Meaning: Selected Writings in Semiotic Theory*, translated by P.J. Perron and F.H. Collins (Minneapolis: University of Minnesota Press, 1987).

Some verb subsets seem indicative of Scully's thematic roles. For example, the subset {accuse, appraise, discover, drive, investigate, kill, launch, look, notice, protect, rescue, shoot, witness} seems indicative of her role as a detective. Portions of this same verb subset is present in Mulder's profile:

What Mulder does: accuse, appall, appear, appreciate, ask, awaken, be, become, believe, call, commit, cry, describe, do, drag, dye, express, fall, fee, fell, figure, find, fly, found, get, give, grab, guess, happen, have, hear, hurt, insist, investigate, keep, kill, know, lay, let, look, lose, masturbate, miss, pay, pull, pursue, run, saw, say, see, send, set, show, slick, spend, take, tell, think, threaten, throw, twit, understand, walk, witness, work

However, when we look for overlaps in Scully's profile in comparison with Mulder's profile, no clear-cut "detective" thematic role emerges from just the verbs even though some of the tell-tale verbs are in both profiles (namely, {investigate, kill, witness}):

Scully and Mulder both accuse, are, ask, believe, call, cry, do, drag, fall, find, get, give, have, investigate, keep, kill, know, look, lose, pay, say, see, spend, take, tell, think, throw, walk, and witness

What is different about these sorts of comparisons facilitated by the Conversation Map system and the comparisons of thematic roles made by narrative theorists is principally this: while narrative theorists have primarily been concerned with the text of the story itself (e.g., the script of the television show), the comparisons made with Conversation Map are examinations of the texts of the audience, i.e., they are examinations of the audience discussions. This turn towards an effort to understand audiences' interactions with one another (rather than, primarily the text of the story or the relationship between the audience and the story) is part of a larger movement within media studies. This new sort of audience studies is concerned with the writings (and other productions) of the audience.[39]

Idealizations (Fetishizations) of the Audiences

Another aspect of audience studies that can be pursued empirically with the Conversation Map system is, what in media studies, has sometimes been called the audience's idealization of characters; or, in more psychoanalytically inflected media studies, the audience's fetishization[40] of the characters. This too can be studied using the verbs associated with the characters in the texts of audience members' messages. However, the respective lists of adjectives are more interesting in this case for the differences they reveal.

---

[39] A recent book does an admirable job in outline the general strengths of this new movement in media studies research: Virginia Nightingale, *Studying Audiences: The Shock of the Real* (New York, Routledge, 1996).

[40] Cf., Christian Metz, *Psychoanalysis and Cinema* (London: Macmillan, 1983).

Qualities of Scully: big, distraught, least, lethal, lovely, my, our, special, strong, sure, tearful, tense

Qualities of Mulder: competent, dazed, far-fetched, greedy, hard, huge, impending, least, likely, megalomaniacal, married, more, most, much, obvious, our, short, spooky, sure, tall, tense, true, untrustworthy, vicious, your

Qualities of both Scully and Mulder: least, our, sure, tense

Social Movements[41]

The second group of messages to be examined in this chapter were sent to newsgroups that participate in what could be called "new social movements." In the 1960s and 1970s a set of "new social movements" gained momentum; e.g., environmentalism, feminism, consumer protection groups, nuclear disarmament, etc. These social movements were not new in the sense of without precedent. Rather they were new in the sense that they focussed on new goals and tactics for organizing especially a rethinking of the habits and practices of everyday life and the production of new kinds of collective identities.[42]

Sociologists of the time found the workings of these new social movements difficult to explain. An Althusserian[43] vocabulary of "interpellation" and "ideological state apparatus" or a Parsonian[44] functionalist description turned out to be inadequate to the task of description. The problem with these older conceptions of social movements was that they did not adequately explain the important role of individual agency and the interactions of individuals in the workings of new social movements. Thus, for instance, where previously it was possible to talk about "the labor movement" in sweeping institutional terms or in the terms of aggregate class politics, such terms miss the point for these new social movements. That point was quite adeptly encapsulated in the aphorism coined by feminists at the time: *The personal is political.*

Since the 1970s transformations in communication technologies, especially the introduction and proliferation of computer networks, has had large implications

---

[41] The work reported in this section is joint work with Joseph Dumit, Professor of Anthropology in the MIT Science, Technology, and Society Program. But any errors in the presentation of this material here are entirely my fault. For an alternative presentation see, for instance, Warren Sack and Joseph Dumit, "Very Large Scale Conversations and Illness-based Social Movements" presented at *Media in Transition*, (Cambridge, MA: MIT, October, 1999).

[42] E.g., Alberto Melluci, "The Process of Collective Identity," in *Social movements and culture*, edited by Hank Johnston and Bert Klandermans (Minneapolis: University of Minnesota Press, 1995).

[43] E.g., Louis Althusser, "Ideology and Ideological State Apparatuses (Notes towards an Investigation)," in *Mapping Ideology*, edited by Slavoj Zizek (New York: Verso, 1994).

[44] E.g., Talcott Parsons, *Social Systems and the Evolution of Action Theory* (New York: Free Press, 1977).

for new as well as "old" social movements. Geographically defined borders are no longer insurmountable for groups of people for whom such borders were impassable even a few years ago. Consequently, completely new groups of people are coming together as social movements where previously -- even if they had been considered a group -- they could never have been considered a movement.

Illness-based social movements are a good example of these sorts of new-new social movements. People suffering from a common, even debilitating condition, can gather and interact through computer networks to voice their collective opinions and organize for social, political, and economic action. I have been attempting to assist the anthropologist Joseph Dumit in his exploration of a set of social movements built around, what he calls, "illnesses you have to fight to get."[45] Attention deficit disorder (ADD), chronic fatigue syndrome (CFS), multiple chemical sensitivity (MCS), and Gulf War syndrome (GWS) are all examples of illnesses where there exist a set of sufferers who claim the illnesses exist, but where the medical and insurance industries do not want to or refuse to recognize the illnesses as legitimate problems.

Online exchanges between members of these groups are not what define the movements: the movements are constructed through multiple media. Nevertheless, online exchanges are important for the movements. For example, people with CFS often have a hard time participating in a face-to-face conversation because they tire easily. However, in an asynchronous medium like email they can take all day, or longer, to compose a message and still keep up their side of the online conversation. Thus, the online discussions can be pivotal for the construction of a collective identity for the movements.

The Conversation Map provides a means to begin to empirically explore the so-called "collective identity" of the group. Literally, we can begin to ask what do "I" and "we" mean for the group as a whole if we aggregate across everyone's language? Also, how are definitions of the illnesses being negotiated and (re)produced in the group discussions?[46]

---

[45] Joseph Dumit, "Symptomatic, Ill and Structurally Damned: Notes on Liminal Creativity and Social Movements," presented at the *American Anthropological Association*, Philadelphia, PA, 1998. See also http://symptom.mit.edu/information/Dumit-Liminality.PDF.

[46] It is important, both for social scientists interested in the structure and dynamics of new social movements, but also for participants, to be able to visualize the movement: *... understanding new social movements in terms of implications for social action requires visualization as well as thought. Such visualization means to see it taking place and to see its impact upon people. This can happen through demonstration projects carried out in practice, and communication about these projects through interactive electronic technology. From* James Herrick, "Empowerment Practice and Social Change: The Place for New Social Movement Theory," presented at *The New Social Movement and Community Organizing Conference*, University of Washington, Seattle, WA. November 1-3, 1995.

One example Conversation Map will be used to illustrate these explorations. Two thousands messages in 866 threads were posted to the Usenet newsgroup alt.med.cfs between 1 January 1996 and 17 February 1996.[47] The Conversation Map computed for this newsgroup for this time period is shown in Figure 11.

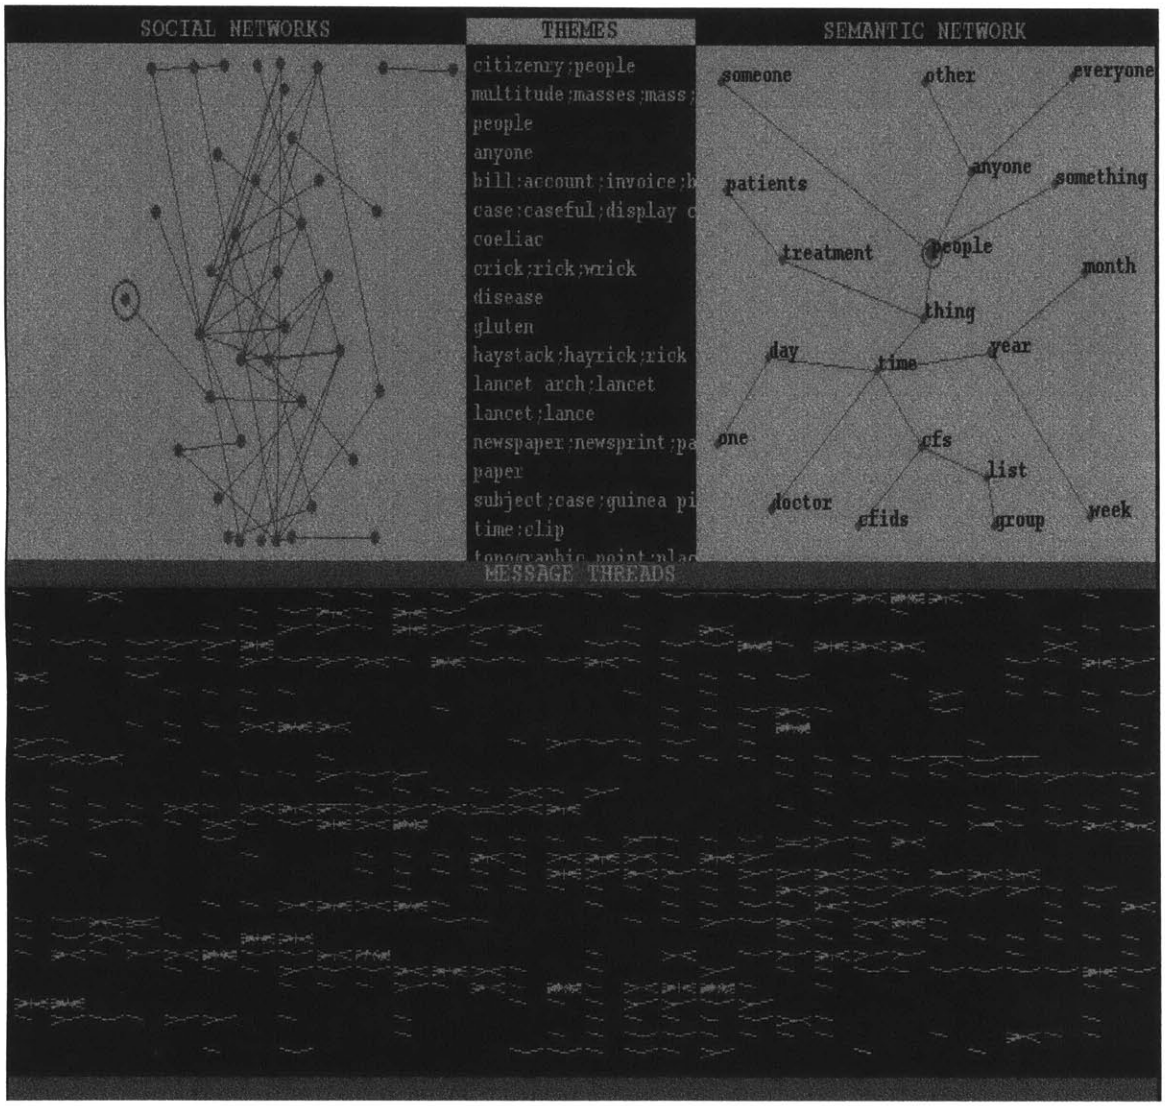

Figure 11: Conversation Map for alt.med.cfs (1 Jan 96 - 17 Feb 96)

See also http://www.interweb-tech.com/nsmnet/herrick.htm

[47] All of the messages analyzed here are publicly archived at http://www.deja.com. The staff at Deja.com was kind enough to provide us with several months of messages from this newsgroup.

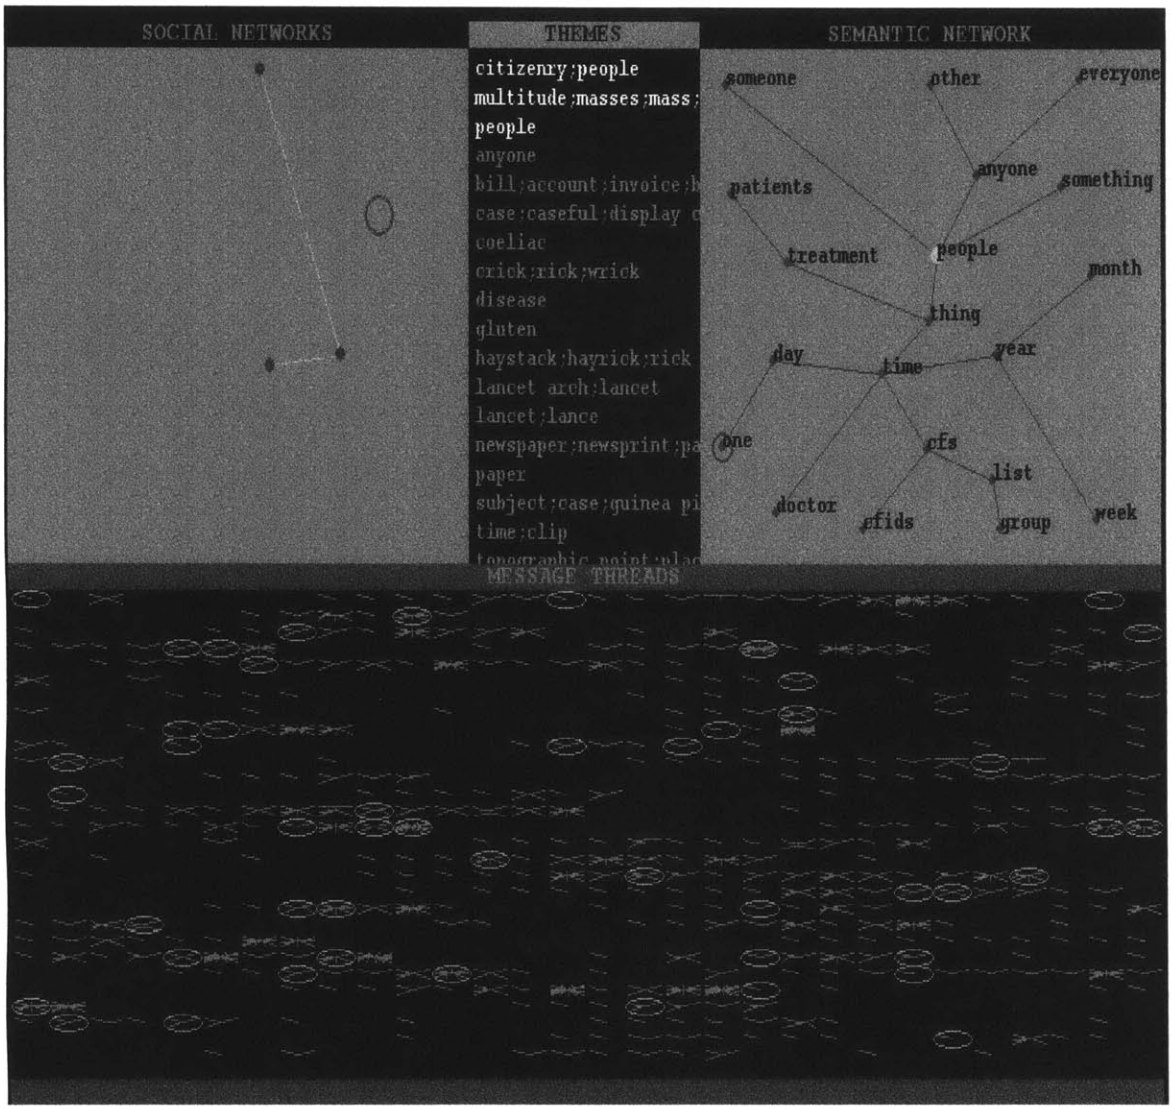

Figure 12: The top three discussion themes for alt.med.cfs

The top discussion themes are all concerned with the issue of people. Selecting the three top discussion themes (by holding down the Control key and clicking on each one in turn; or, simply by clicking on the term "people" in the semantic network) provides a view of the number of threads in which "people" has been a theme of discussion. This view is shown in Figure 12. The view shows us that "people" is used as a theme of discussion in many of the message threads and is one of the central terms of discussion. In other words, "people" is a term important to the group discussion and is not something that only one or two people have mentioned.

The upper left-hand corner of Figure 11 (i.e., the social network pictured there) shows a relatively cohesive social network for the online conversation. However, double-clicking on the term "people" in the semantic network generates a web page like the one shown in Figure 13 and this web page shows that the "profile"

for the term "people" contains many contradictions. For example, all of the following antonyms are listed in the adjectives used to modify the term "people": elderly/young, new/old, healthy/sick, bad/good, and compassionate/nasty. This view from the Conversation Map provides some evidence that "people" is a contended term in the discussion.

Figure 13: Associations for the term "people" in alt.med.cfs

So, if "people" is a term under contention, then in what terms does the so-called "collective identity" reside? The first obvious place to look would be under the term "we." Has the discussion produced a collective voice, a way of speaking for the group using the term "we"? A quick look at the profile for "we" (in Figure 14) shows that "we" has not been used much in the discussion.

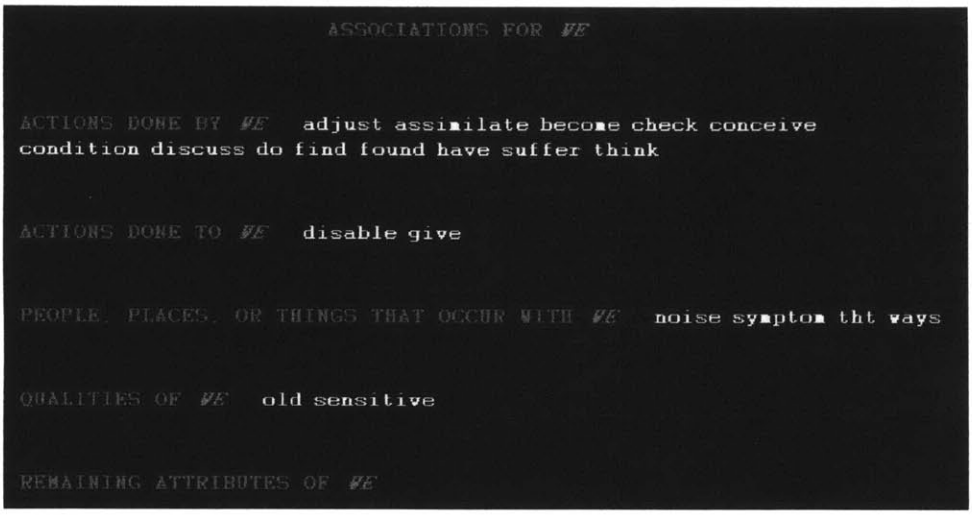

Figure 14: Associations for the pronoun "we" in alt.med.cfs

However, Figure 15 shows that many of the messages are written in the first person and that, therefore, the pronoun "I" has a substantial profile. Moreover, the picture of "I" that comes out of the CFS discussion is a surprisingly coherent. Focussing now on the actions associated with "I," it appears to be the case that the "I" of this newsgroup is quite cerebral and not very physically active. Except for a few verbs (e.g., exercise, kill, march, marry, toss, and urinate) most of the verbs for which "I" functions as a subject, are actions that can be done in an armchair. But, the actions of "I" are also quite eclectic and numerous and so the profile paints a picture of a busy collective identity constructed from the many activities of the individuals in the group. This picture of busyness seems to be at odds with what one might guess a group on chronic fatigue would project. This profile of "I" shows how the Conversation Map system might be used to begin to think about how new social movements sum together the personal contributions of many people into a larger collective statement.

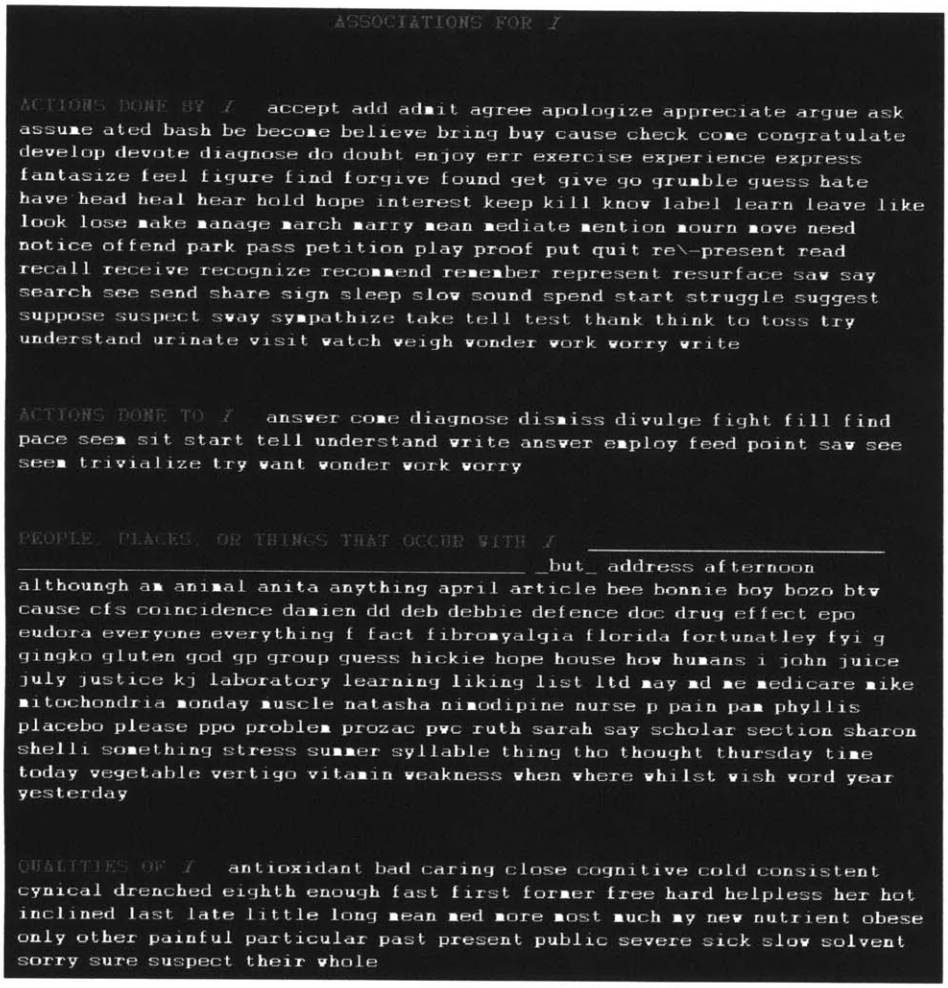

Figure 15: Associations for the pronoun "I" in alt.med.cfs

The other obvious place to begin navigating through the messages of a newsgroup like this one would be to start with the profile for the illness itself, namely CFS (chronic fatigue syndrome). Figure 16 lists the profile for CFS. Although one of the aims of new social movements organized around CFS is to change the medical and insurance industries conception of CFS, nothing in the CFS profile is immediately striking as particularly innovative. However, by looking at the list of actions done to CFS, it appears that a variety of people in the group are talking about CFS as if it were a conventional illness, i.e., as though it had the medically regularized status many CFS sufferers would like it to have. According the profile CFS can be correlated, cured, diagnosed, exacerbated, suffered, and triggered.

Figure 16: Associations for the term "CFS" in alt.med.cfs

Users Groups

The example analyses of audiences and social movements given above showed the use of the Conversation Map system as a kind of "meta-browser." I say "meta" because only the layers built on top of the messages by the system were examined or discussed; e.g., the word profiles, the social networks, the semantic networks, and the themes of discussion. This section will take a more practical turn. The system will be used to find and read a particular message from the archive of messages. In short, we will focus on a task that is still a browsing or navigating activity, but is quite close to a search. This section discusses an activity that is useful in the context of "users groups," i.e., newsgroups devoted to the use of and design with certain kinds of computer software and/or hardware.

The lego-robotics newsgroup is an example of such a users group. The lego-robotics group is devoted to a technical discussion on robotics, specifically the engineering of robots using the components of the various Lego robotics kits, like LegoMindstorms. The archive discussed here contains over 1800 messages

from about 340 participants that were posted over a four month period (4 February 1999 - 5 June 1999).[48]

Imagine entering this discussion as a new participant. You are interested in building a robot with your newly acquired Lego robotics kit, but you don't know what programming language to use in the construction of your robot. You have heard that the programming language NQC (which stands for Not Quite C) would be a good choice. Therefore you would like to query the archive of messages with a question something like this: Were there any discussions in the last few months about NQC in which one of the main contributors to the newsgroup had something to say? You are interested in a discussion, not just a mention of NQC because you want an analysis of its merits. And, you want to hear from one of the main contributors to the newsgroup because then you would probably be getting the opinion of someone who is well respected and knowledgeable. So, the task is almost a search task, but it is not exactly a search task because it does not necessarily have a correct answer.

A navigation task of this sort would be difficult, if not impossible, with a conventional newsgroup browser. With the Conversation Map system it is relatively straightforward to do. However, what follows is a rather long string of images because what is presented here illustrates the results of a sequence of mouse clicks. In Figure 17, is a picture of the lego-robotics Conversation Map as it appears when the browser is first started.

[48] To subscribe to this group go to http://www.lugnet.com/robotics. That site also publicly archives all of the messages of the lego-robotics discussion.

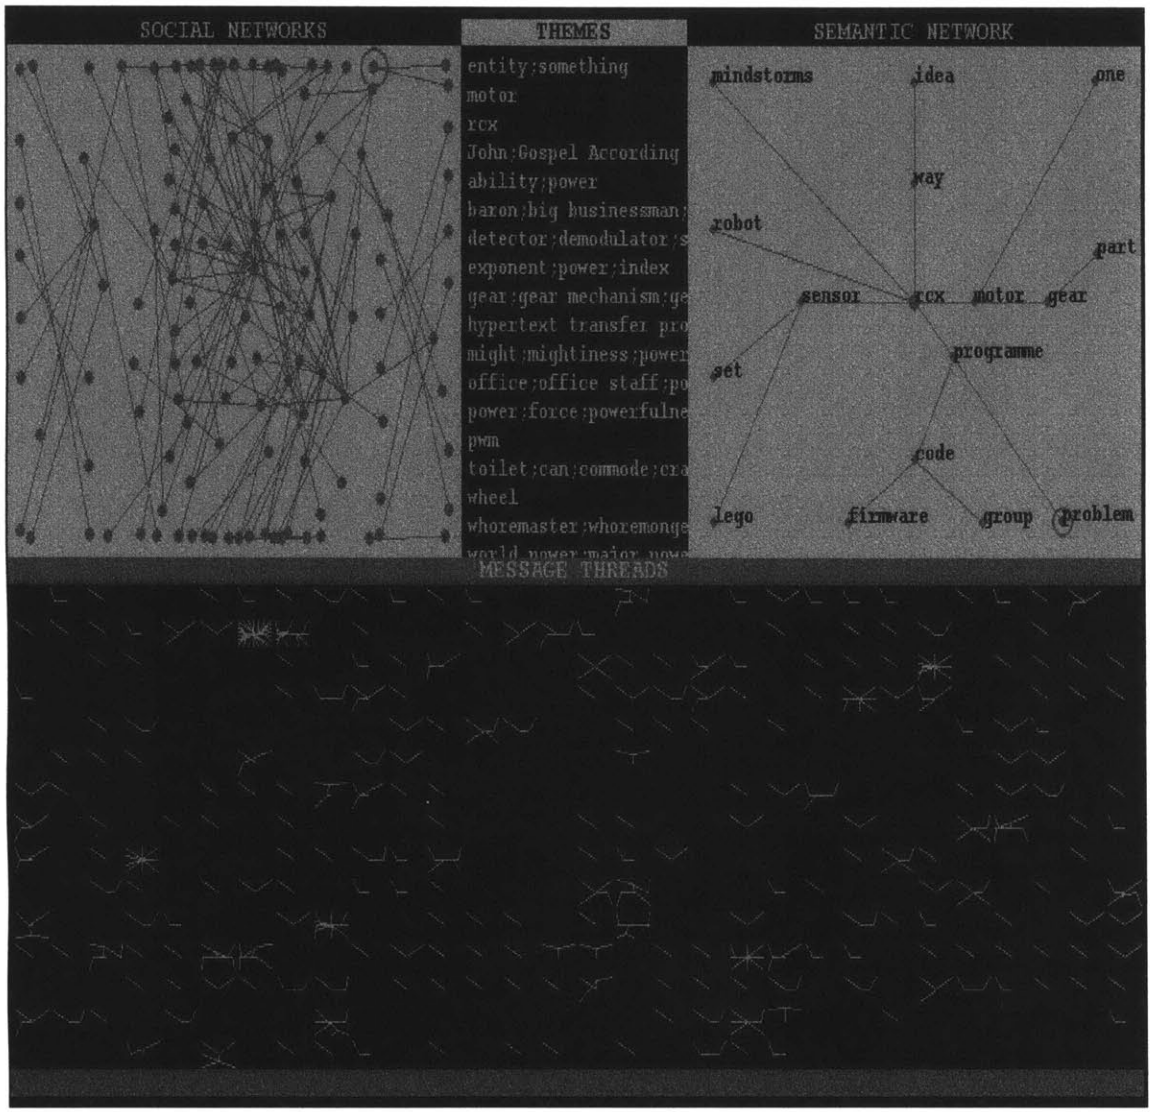

Figure 17: Conversation Map for the lego-robotics newsgroup (4Feb99 - 5Jun99)

In Figure 18 the menu of discussion themes has been scrolled down to the theme of interest (NQC) and that theme has been selected with a mouse click to reveal which threads in the archive include NQC as a theme of discussion. Moreover, note that one pair of participants in the social network is connected by the theme NQC.

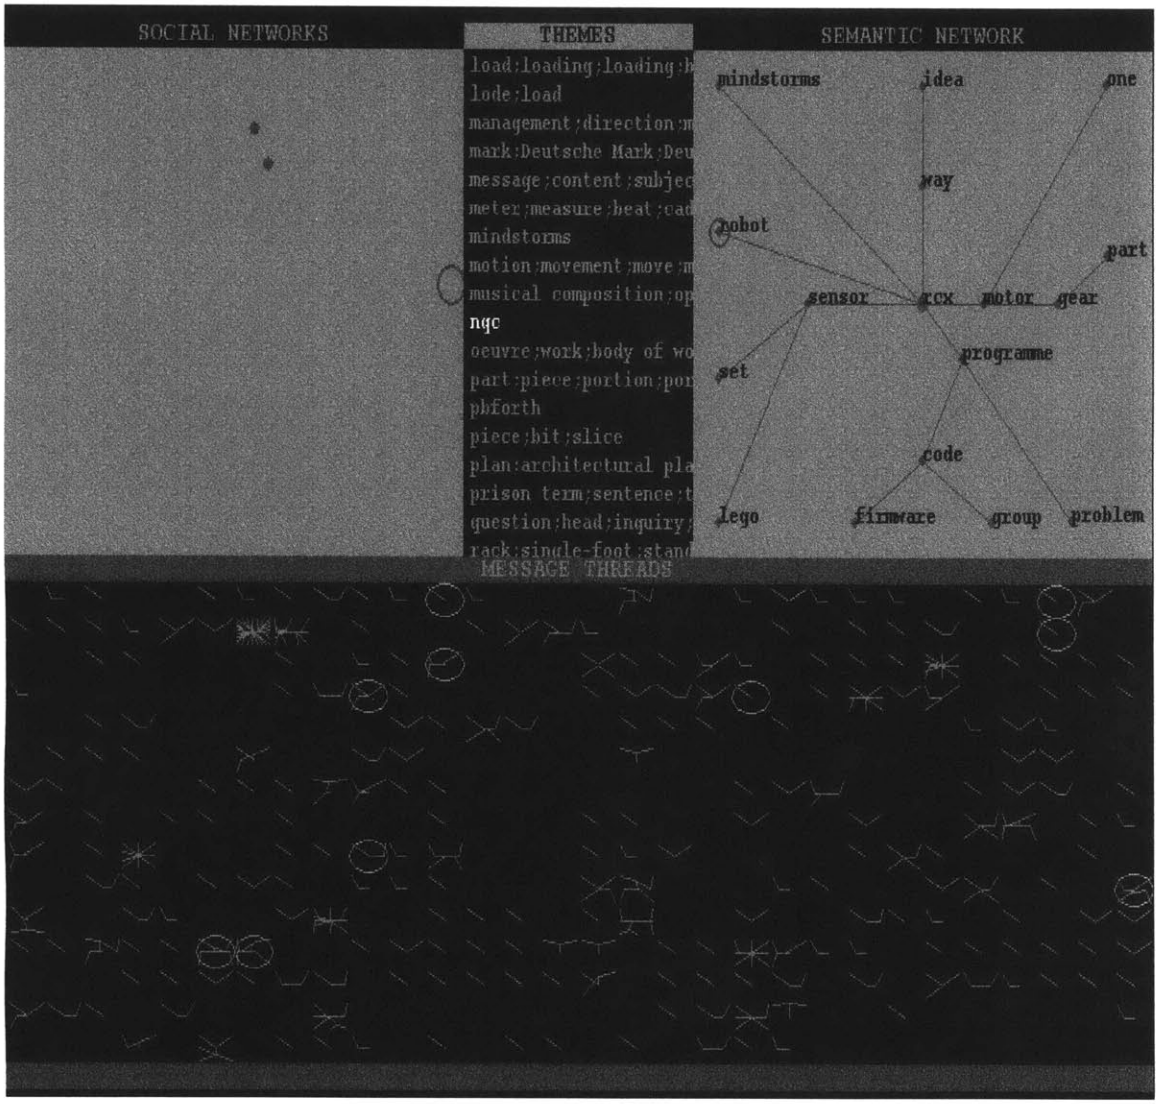

Figure 18: NQC as a theme of discussion in the lego-robotics newsgroup

By clicking on the highlighted participants in the social network it is possible to see if one or both of them is connected in conversation with a large number of other participants. Figure 19 shows that the one selected is in fact connected to many other participants in the social network. This connectivity is a good indication that the selected participant is one of the main contributors to the newsgroup.

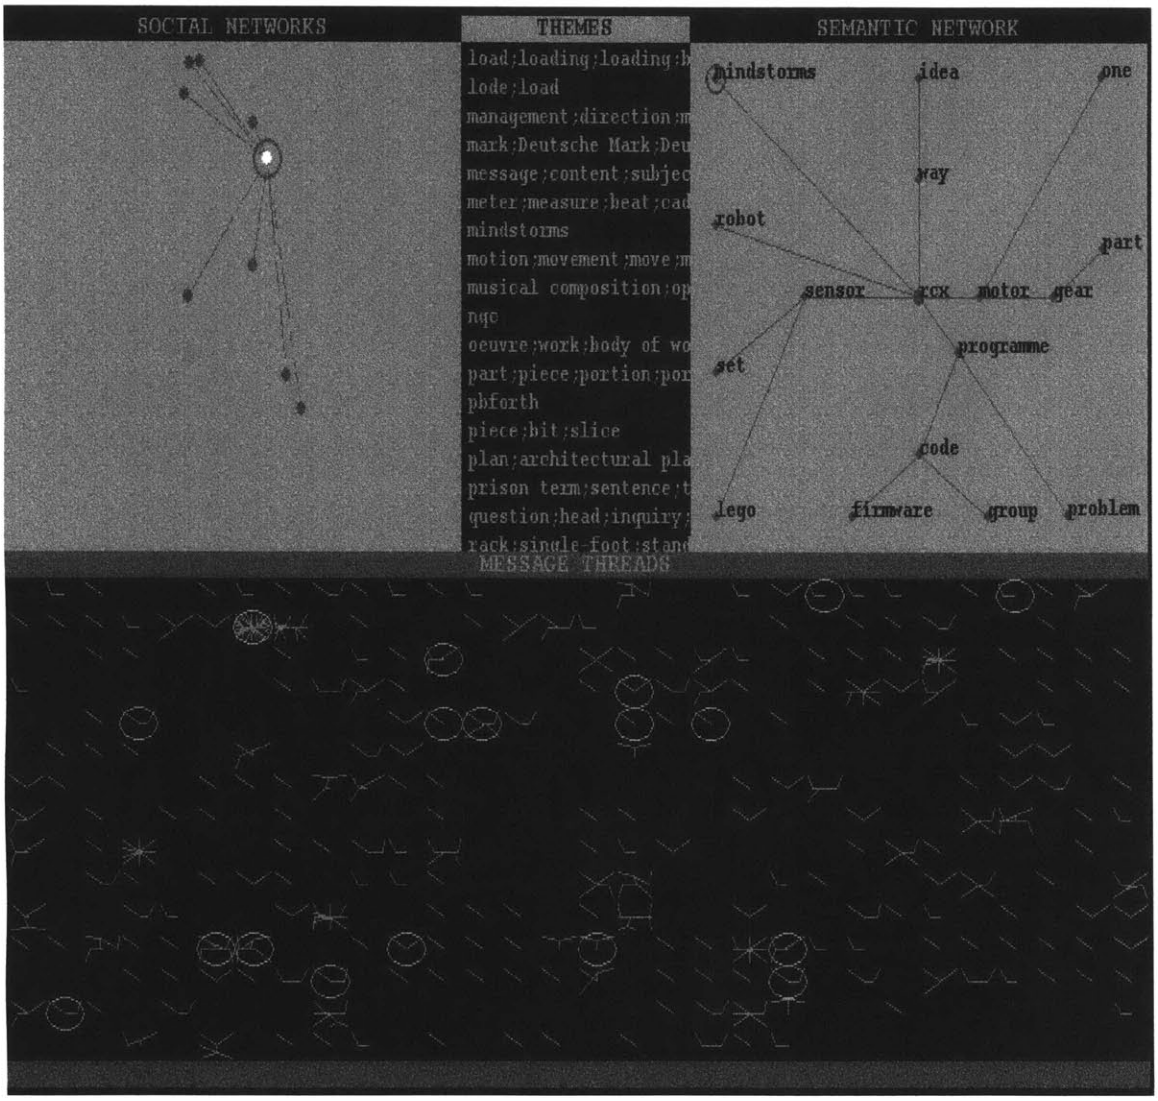

Figure 19: Highlighting one of the main contributors to the newsgroup

Now, if the other participant in the pair is also highlighted (by holding down the Control key and clicking on the second participant's node in the social network) the threads where these two exchanged messages are highlighted. Also, in Figure 20, one can see that these two participants have discussed NCQ and several other subjects listed in the menu of discussion themes.

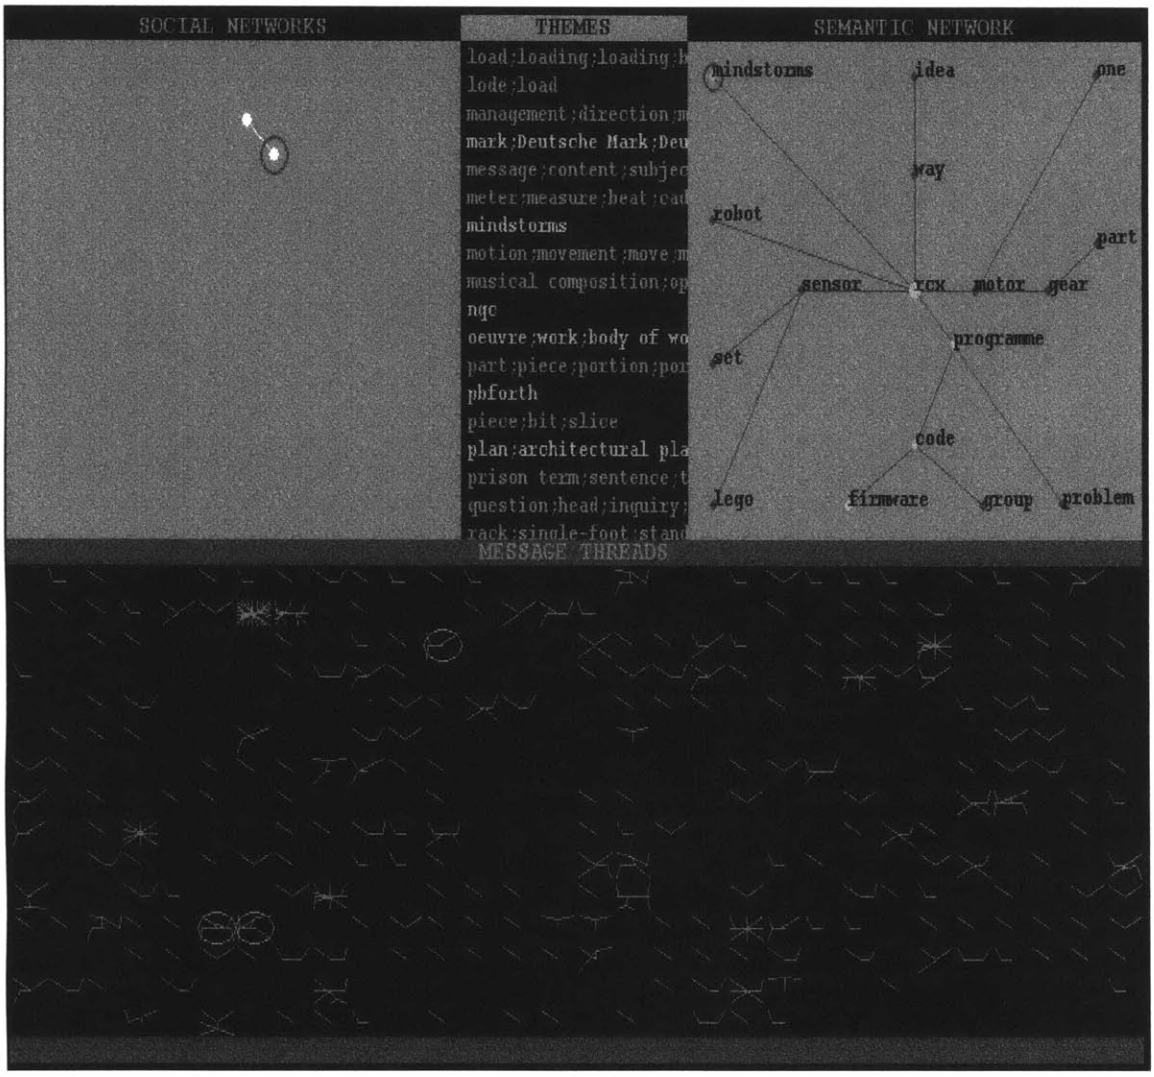

Figure 20: Highlighting threads and themes of two participants

Comparing Figure 20 with Figure 18 it is possible to see that all of the three threads highlighted in Figure 20 are possible places where a suitable discussion about NQC might have taken place. Selecting one opens the thread as shown in Figure 21. The menu of themes contained in the thread is scrolled down until NQC can be seen. NQC is selected with a mouse click and the part of the thread that contains NQC as a theme of discussion is automatically highlighted.

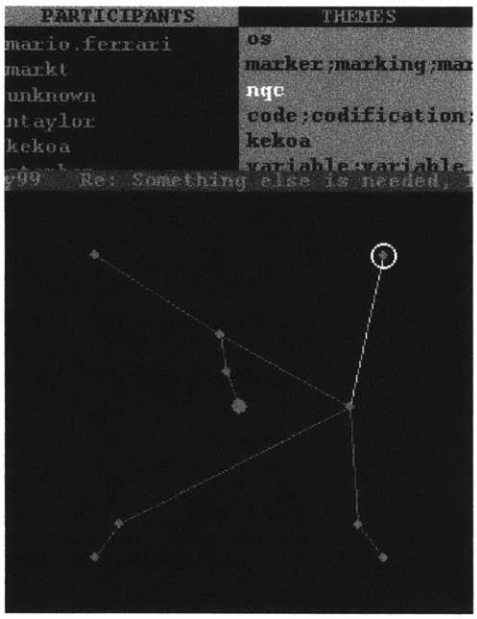

Figure 21: A pair of messages where NQC is a theme is highlighted in the thread

Clicking the mouse on the circled message in the thread opens a new web page containing the text of the message. As can be partially seen in Figure 22, the message contains an interesting discussion of the strengths and weaknesses of the NQC programming language for the rigors of robotics. Note that the subject line of the message says nothing about NQC.

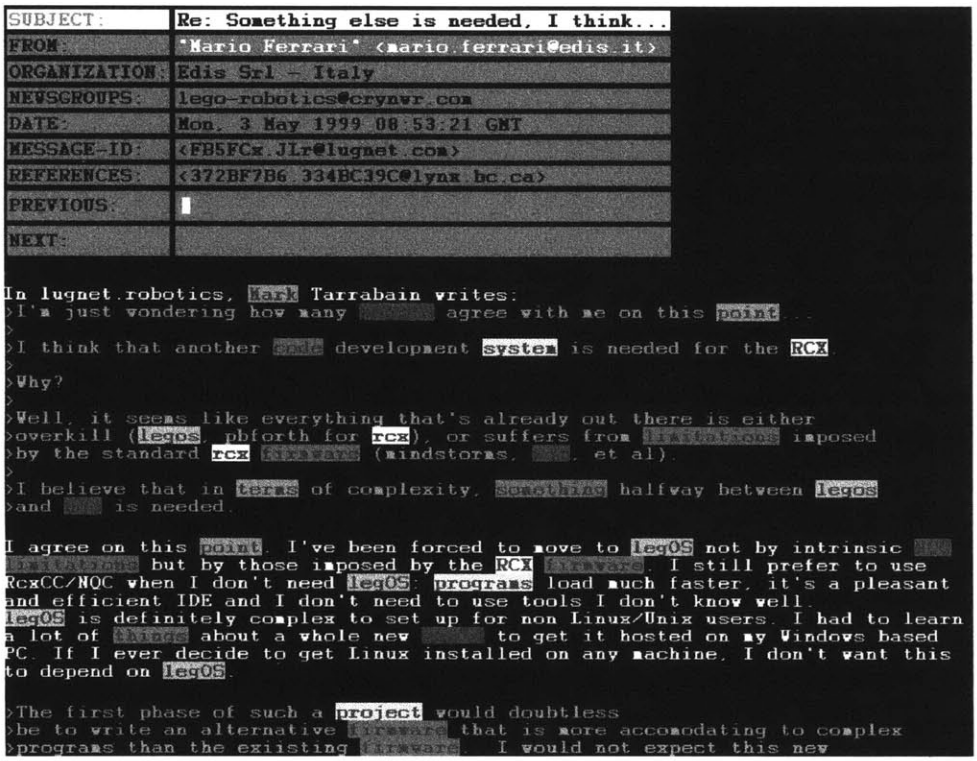

Figure 22: A message in which NQC is discussed

The example given above does not prove that combined social-semantic navigations of this sort are always easy to do with the Conversation Map system. However, it does demonstrate how it is possible to perform a practical but computationally non-obvious task of this sort with the system.

Citizen Diplomats

The final VLSC to be considered is one populated by a group that I will call "citizen diplomats." One unprecedented activity that the Internet has made possible is the debate of international politics by ordinary citizens in different countries on a daily basis in a public "space" where people do not necessarily know one another before the debate begins. Such a debate occurred during the spring of 1999 during the war in Kosovo in the Usenet newsgroup devoted to discussion of Albanian culture (soc.culture.albanian). At the time when the messages analyzed were being exchanged (16 April 1999 - 4 May 1999) the newsgroup was focused on a vicious argument over the war in Kosovo. The archive discussed here contains over 1200 messages from about 260 participants.[49]

[49] Current messages from this newsgroup can be read here: news:soc.culture.albanian. The messages analyzed below can also be found in the public archives of Deja.com: http://www.deja.com.

The Conversation Map for this archive was used as the main example for the Interface chapter and so it will not be displayed again in this chapter. Only the semantic network is shown in Figure 23. The semantic network will be examined to provide a taste for how the navigation of this new sort of political exchange might provide some insights into the micro-politics of international relations.

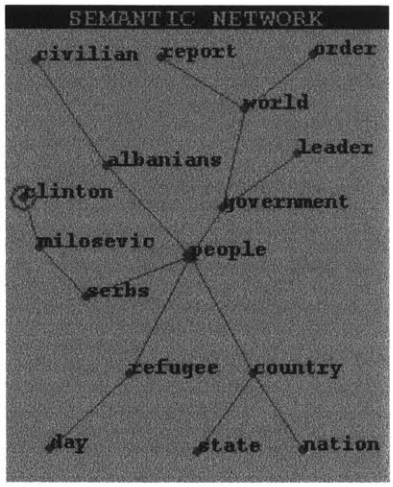

Figure 23: Semantic network for soc.culture.albanian (16 Apr 99 – 4 May 99)

One optimistic way of reading the semantic network computed by the Conversation Map system for the soc.cult.albanian group is this: "people" is a neutral term: Serbs, Albanians, refugees, countries, and governments are all "talked about" like people. This is a sort of thin humanism ("after all we are all people") that begins to explain why any sort of exchange can happen in this VLSC concerning Kosovo, even though, admittedly, the exchange is a very heated and argumentative one.

However, it must be kept in mind that no one in the newsgroup necessarily wrote "we're all people." The comparison, in fact, 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 just 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.

If the 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.cult.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: Serbia." (2) "It is not enough to be alive, *people need* normal life."

Similar verbs 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 overlap designate possible neutral 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.

Conclusions

Messages from four VLSCs involving four different sorts of groups (audiences, social movements, users groups, and citizen diplomats) were explored using the navigation capabilities of the Conversation Map system. The Conversation Map system provides a means for empirically exploring a variety of outstanding social and political issues from the fields of media studies (e.g., identification), sociology and anthropology (e.g., collective identity), and political science (e.g., shared attributes which might indicate neutral terminology for use in discussion). Also, through an extended example, it was shown how the Conversation Map can be used to perform combined social and semantic navigation tasks for more pragmatic purposes (e.g., looking for reasoned opinions about the merits of a piece of software).

ANALYSIS

The analysis procedures of the Conversation Map system use, extend, and combine in novel ways techniques from computational linguistics and quantitative sociology. The goal of the analysis procedures is to generate a graphical summary of archives of online, email-based VLSCs. Three new analysis techniques of potential interest for computational linguistics and sociology are developed: (1) Inter-message lexical cohesion: An algorithm has been created for the computation of lexical ties between email messages that takes into account information about both threading and quotation in the messages (thereby producing an approximation of the themes of conversation). (2) Social cohesion: A definition of social cohesion has been developed. Social cohesion is the overlay of lexical cohesion information on top of a social network (thus creating a rough description of what the conversations between participants are about). And, (3) Social network-based sorting: A means has been designed and implemented to use the computed social networks to weight the importance of word entries in automatically computed rough draft thesauri (which is useful for selecting the terms which are important as discussion themes and thus important to the social structure of the group). A step-by-step description of the analysis procedures is presented. Improvements over related work in CSCW, sociology, and computational linguistics are discussed.

Any text or dialogue establishes a linguistic context within which subsequent utterances must be understood. And beyond the linguistic context is the participatory context. A speaker or a writer directs an utterance or text toward a hearer or reader with a particular purpose – to inform, to amuse, to collaborate in a task, perhaps.[50]

Over the past forty years a large number of computational text summarization technologies have been invented.[51] However, none of the existing technologies take into account what might be called the *participatory context* of the texts summarized. This is a valid oversight when the texts have all been written by one person or even when the summarized texts have been written by multiple authors in a more or less anonymous fashion; e.g., encyclopedia articles and newspaper stories. But, this oversight becomes more problematic when the texts to be summarized are letters, email messages, transcripts of telephone conversations, and other writings where the texts are replies to or commentaries on other texts and are produced to reflect the opinions of the authors. With these sorts of texts – texts of a more conversational character – the participatory context reflected or produced by the writings can be very important. This chapter outlines a set of text analysis procedures that have been developed to automatically generate graphical summaries of very large-scale conversations (VLSCs) like Usenet newsgroup exchanges involving hundreds or thousands of people.

A variety of text analysis and generation methods have been implemented that do take into account the participatory context of conversation-like texts; although, these computational methods have not been applied to the tasks of summarization of large amounts of text. Instead, they have been implemented to, for example, follow the psychological point of view in narrative texts,[52] or participate in one-on-one dialogues with a user.[53] It has yet to be determined whether the techniques and technologies developed for the analysis of participatory context in intimate conversations and small group discussions can be scaled up for the analysis of discussions involving hundreds or thousands of people.

Outside of computational linguistics, within the field of sociology, a number of computational approaches to understanding the participatory or social context of texts have been developed. Most prominently these methods have been applied

---

[50] Janyce Wiebe, Graeme Hirst, & Diane Horton, "Language use in context," *Communications of the ACM (Special issue on natural language processing) 39*, 1 (1996), 102-111.

[51] Inderjeet Mani and Mark T. Maybury (editors) *Advances in Automatic Text Summarization* (Cambridge, MA: MIT Press, 1999).

[52] E.g., Janyce M. Wiebe, "Tracking point of view in narrative," *Computational Linguistics 20*, 2 (1994), 233-287;

[53] J. Moore and C. Paris, "Planning texts for advisory dialogues: Capturing intentional and rhetorical information," *Computational Linguistics 19*, 4 (1993), 651-694.

to the literatures of science. For example, the methods of co-citation analysis[54] are routinely applied to determine the relative importance of a scientific article: its significance is thought to be a function of the number of other articles that cite it. The methods of social network theory[55] and actor-network theory[56] provide technologies akin to co-citation analysis, but have their own particular strengths and weaknesses. The techniques of co-citation analysis have also been elaborated and improved upon by artificial intelligence, natural language processing researchers.[57]

However, the participatory context of a text is not established solely through citation or quotation. A variety of other ties of cohesion[58] and relations of coherence are used to establish the participatory context. Recent work in cohesion analysis has produced new computational methods for summarizing and cross-indexing documents. New work in cohesion analysis includes, for example, lexical cohesion analysis of single documents;[59] discourse segmentation of single documents;[60] and, coreference analysis of multiple documents.[61] Interesting new work shows how cohesion structure can be used

---

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

[55] See, Stanley Wasserman and Joseph Galaskiewicz (editors) *Advances in Social Network Analysis: Research in the Social and Behavioral Sciences* (Thousand Oaks, CA: Sage Publications, 1994).

[56] Michel Callon, John Law, Arie Rip (editors) *Mapping the Dynamics of Science: Sociology in the Real World* (London: Macmillan Press, Ltd., 1986). See also Bruno Latour and Geneviève Teil "The Hume Machine: Can association networks do more than formal rules" *Stanford Humanities Review (special issue on artificial intelligence)* 4.2 (1995): 47-65. The computational techniques of actor-network analysis -- specifically co-word analysis -- is basically the calculation of mutual probabilities between nouns in scientific abstracts and so this technique probably has more affinities with techniques in computational linguistics than with those developed by other sociologists.

[57] Wendy Lehnert, Claire Cardie, and Ellen Riloff. "Analyzing research papers using citation sentences. In *Proceedings of the 12th Annual Conference on Cognitive Science*, 1990.

[58] Michael A.K. Halliday and Ruqaiya Hasan *Cohesion in English* (New York: Longman, 1976).

[59] E.g., 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).
See also Mark A. Stairmand. "Textual context analysis for information retrieval" in the *Proceedings of the 20th annual international ACM SIGIR conference on Research and development in information retrieval*, August 1997.

[60] Daniel Marcu. "Improving summarization through rhetorical parsing tuning" in the *Sixth Workshop on Very Large Corpora* (Montreal, Canada, August 1998).

[61] Amit Bagga and Breck Baldwin "Entity-Based Cross-Document Coreferencing Using the Vector Space Model" In *Proceedings of ACL-COLING'98* (Montreal, Canada, June 1998).

in the inference of discourse coherence relations.[62] Nevertheless, the goal of work in computational cohesion analysis – like most of the work done in automatic summarization -- has largely been to automatically find intra- and inter-*textual*[63] relations between documents. The work described in this chapter is an attempt to analyze inter-*textual* relations between documents in conjunction with *ideational* and inter-*personal* relations. In other words, the goal is to analyze the ideas, the textual, and the participatory structure of a discussion involving a large number of people.

The terms – textual, ideational, and interpersonal – are the terms of systemic functional linguistics. Another way of differentiating the work reported here from other related work in quantitative sociology and computational linguistics is in these terms of the systemic functional linguist Michael Halliday. According to Halliday[64], language has at least three meta-functions: (1) *ideational*: language can represent ideas; (2) *interpersonal*: language functions as a medium of exchange between people; and, (3) *textual*: language functions to organize, structure, and hold itself together; this function allows the various devices of cohesion, including citation, ellipsis, anaphoric reference, etc. to be used. Unlike most other text summarization technologies, the procedures described in this chapter include some representation of the interpersonal and so operate on all three meta-functions of the texts summarized. The procedures reported here also differ from those used in, for instance, co-citation analysis because those of co-citation analysis do not contain information about the ideational or textual aspects of the texts analyzed.

In addition to these differences from other text summarization work, the work described here is aimed at the generation of a graphical – rather than a textual – summary. Why a graphical summary? Because the work described here is intended to function as the backend to a graphical, newsgroup browser: the *Conversation Map* system. The automatically generated graphical summaries function as graphical interfaces for browsing and navigating large archives of newsgroup messages. With a graphical interface generated by the Conversation Map system one can browse a set of newsgroup articles according to who is "talking" to whom, what participants are "talking" about, and the central terms and possible emergent metaphors of the conversation.

The goal is therefore to supply summaries that are "sufficient for the day," for example, where the role of summaries is to facilitate browsing in an interactive

---

[62] Sanda Harabagiu. "WordNet-Based Inference of Textual Cohesion and Coherence" in the *Proceedings of FLAIRS-98* (Sanibel Island, FL, May 1998).
[63] Michael A. K. Halliday. *An Introduction to Functional Grammar, Second Edition* (London: Edward Arnold, 1994).
[64] *Ibid.*, 179.

information retrieval environment where there are other search tools available, as well as access to the source documents;[65]

It is already the case that browsers for email, newsgroups, streaming audio, and the World Wide Web are a multi-billion dollar business and a huge cultural phenomenon. While some of the technologies developed for text summarization are now being used for web page search engines, relatively little work has been done towards the type of automatic summarization that Karen Sparck Jones envisions as being good for browsing and navigating large bodies of information. The contributions described in this chapter can be understood as interpersonal and graphical improvements to the state of the art of text summarization in order to support the design and implementation of a new kind of "content-based," graphical, newsgroup browser.

To allow a combination of social and semantic navigation[66] the Conversation Map system computes a social network[67] corresponding to who is replying to (or citing) whose messages. The Conversation Map system also parses and analyzes the contents of the newsgroup articles to calculate a semantic network[68] that highlights frequently used terms that are similar to one another in the newsgroup discussion. The design and implementation of these two functionalities required the research and development of three new techniques of potential interest for computational linguistics and sociology:

(1) **Inter-message lexical cohesion**: An algorithm has been created for the computation of *lexical ties* between email messages that takes into account information about both threading and quotation in the messages (thereby producing an approximation of the themes of conversation).

(2) **Social cohesion**: A definition of *social cohesion* has been developed.[69] Social cohesion is the overlay of lexical cohesion information on top of a social network (thus creating a rough description of what the conversations between participants are about).

(3) **Social network-based sorting**: A means has been designed and implemented to use the computed social networks to weight the importance

---

[65] Karen Sparck Jones, "Automatic summarizing: factors and directions," in Mani and Maybury, 7.

[66] This distinction is described in Paul Dourish and Matthew Chalmers. "Running Out of Space: Models of Information Navigation." Short paper presented at *HCI'94* (Glasgow, UK, 1994).

[67] Wasserman and Galaskiewicz.

[68] Cf., M.R. Quillian "Semantic Memory" In M. Minsky (editor) *Semantic Information Processing* (Cambridge, MA: MIT Press, 1968).

[69] Warren Sack, "Diagrams of Social Cohesion," In *Descriptions of Demonstrated Systems, ACL'99* (University of Maryland, College Park: Association for Computational Linguistics, June 1999).

of word entries in automatically computed rough draft thesauri (which is useful for highlighting certain regularities regarding how some terms are discussed over the course of an online conversation).

These three contributions will be discussed in the context of a description of the VLSC summarization procedure that generates a graphical summary given an input of a large archive of newsgroup messages.

Overview of the VLSC Summarization Procedure

Input A file containing a number of email messages posted to a newsgroup

Output
(1) a list of *lexical ties* between messages
(2) a set of *social networks*: nodes are participants; arcs indicate mutual response and/or quotation and are labeled with lexical tie information; the lexical tie information is sorted and presented as a menu in the final summary
(3) a set of "social network sorted" *semantic networks*: these are a rough draft thesaurus for the nouns which appear in the corpus of messages

The analysis procedure of the Conversation Map system performs the following steps on an archive of newsgroup messages in order to compute the graphical summary described above. Steps N, O and P correspond to the three contributions listed above.

(A) Messages are threaded.

(B) Quotations are identified and their sources (in other messages) are found.

(C) A table of newsgroup participants (i.e., those people who have posted to the newsgroup) to messages is built.

(D) For every participant, the set of all other participants who replied to the participant is recorded. Similarly, for every participant, the set of all other participants who quoted from the messages of the participant is recorded. Participants who reciprocally reply to and/or quote from one another's messages are linked together into a social network.

(E) The "signatures" of participants are identified and distinguished from the rest of the contents of each message.

(F) The words in the messages are divided into sentences.[70]

(G) Sentences are divided into clauses according to punctuation (e.g., parenthetical phrases are subordinated to main clauses) so that the parser (applied in step I) does not need to handle punctuation.

(H) Discourse markers (e.g., connecting words like "if", "therefore", "consequently", etc.) are tagged in the messages.[71]

(I) Every word of every message is tagged according to its part-of-speech (e.g., "noun", "verb" "adjective", etc.)[72]

(J) Every word is morphologically analyzed and its root is recorded.[73]

(K) Every noun is "term expanded" using the thesaurus WordNet. Recorded along with each noun are its synonyms, antonyms, hypernyms (generalizations), hyponyms (specializations), meronyms (parts; e.g., a table has legs and a top), and holonyms (wholes for which the noun might be a part; e.g., a leg can be a part of a table).

(L) The words of the messages are parsed into sentences using a partial parser.[74]

(M) The lexicosyntactic profile of every noun in the archive is compared to the lexicosyntactic profile of every other noun in the archive.[75] Nouns that are used or discussed in the same manner are calculated to be similar and are placed close to one another in the semantic networks. One can understand this semantic network as a crude approximation to the sorts of metaphors of discourse identified by linguists like George Lakoff and Mark Johnson.[76]

---

[70] The tool described in the following paper is used: Jeffrey C. Reynar and Adwait Ratnaparkhi. "A Maximum Entropy Approach to Identifying Sentence Boundaries." In *Proceedings of the Fifth Conference on Applied Natural Language Processing* (Washington, D.C., March 31-April 3, 1997).

[71] We use a list of discourse markers compiled by Daniel Marcu, *The Rhetorical Parsing, Summarization, and Generation of Natural Language Texts*, Ph.D. Thesis (Toronto: Department of Computer Science, University of Toronto, December 1997).

[72] A simple trigram based tagger is used to accomplish the part-of-speech tagging.

[73] The database containing morphological and syntactic information comes from the University of Pennsylvania: Daniel Karp, Yves Schabes, Martin Zaidel, and Dania Egedi. "A Freely Available Wide Coverage Morphological Analyzer for English" in *Proceedings of COLING-92*, (1992).

[74] The partial parser is a re-implementation and revision of the parser described here: Gregory Grefenstette, *Explorations in Automatic Thesaurus Discovery* (Kluwer Academic Publishers: Boston, 1994).

[75] An algorithm similar to the one described in Grefenstette, *Ibid.* is used.

[76] George Lakoff and Mark Johnson. *Metaphors We Live By* (University of Chicago Press: Chicago, 1980).

Thus, for example, if the noun "economy" and the noun "plant" are often associated with the same verbs and adjectives (e.g., "plants grow", the economy grows", "plants have roots", "the economy has root", "we have a healthy economy", "we have a healthy plant" etc.) the two words will be closely coupled in the word associations network and one can read that network as stating something like "the economy is like a plant."

(N) An analysis of lexical cohesion[77] is performed on every pair of messages where a pair consists of one message of a "thread" and another message that either immediately follows the first message in the thread (i.e., is a reply to the first message) and/or follows the first message in the thread and contains a quotation from the first message. This analysis produces a series of lexical ties between messages that can be understood as a crude approximation to the theme of the conversation in a sequence of messages. The lexical database WordNet is used in the lexical cohesion procedure.[78]

(O) By using the index created in step (C) with the results of step (N) a set of lexical ties are computed for every pair of participants who have replied to and/or quoted from one another over the course of time represented by the newsgroup archive under analysis. These aggregated lexical ties are layered on top of the social network computed in step (D). The result is that most of the links between pairs of participants are labeled with one or more lexical ties (i.e., one or more "discussion themes"). The combination of social networks and lexical cohesion results is called *social cohesion*.

(P) Once all of the noun-noun pairs have been compared in step (M) and a nearest neighbor for each noun computed, a subset of the semantic networks computed are selected for display by ranking the semantic networks. The top-ranked semantic network contains a set of terms (used as "discussion themes") that connect the greatest number of participant pairs linked in step (m). In this manner, information about the social networks of the newsgroup is used as a kind of "lens" to select an important subset of the semantic information. Effectively, this type of interlacing of the social and semantic information supports social and semantic navigation in the interface generated for the newsgroup.

---

[77] See Michael A.K. Halliday and Ruqaiya Hasan *Cohesion in English* (Longman: New York, 1976). The lexical cohesion analysis procedure we have developed is akin to, but different than, the one described here: 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 (MIT Press, Cambridge, MA, 1998).

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

Preprocessing

All of the software (for steps A through P) has been implemented in about 8000 lines of code written in the Perl programming language.[79] Several lexical resources compiled by others have been incorporated into the system; specifically a machine-readable lexical database, a morphology database, and some part-of-speech frequency information compiled from a hand-tagged corpus of news stories. Steps A through M are necessary preprocessing steps. Some of A through M have been implemented in other systems and so are not a focus of this chapter. However, some of steps A through M are novel. Steps N through P inclusive (the steps that compute the three main results of the analysis) will be later explained in detail.

The most important of the novel preprocessing steps is a set of methods to identify quotations in messages. Although, at first, it seems that it would be trivial to recognize quoted text in a message, it turns out not to be so easy. Reasons for why it is difficult are given and an extended, computational definition of quotation is given. Similarly, identifying signatures (i.e., the addresses, aphorisms, ASCII art, etc. that some people use to sign their messages with or after their name) would seem to be a simple problem, but it is not. So, a computational definition of signature is also given. The main results of this chapter (a definition for lexical tie between messages, etc.) are strongly dependent upon the preprocessing steps, especially the steps to identify quotations and signatures.

To be more precise about the work involved in these steps, several of them are defined in pseudo-code. The hope is that some of the important intricacies of the implementation can be more clearly rendered in commented pseudo-code than they could be in just English. However, the descriptions can be roughly understood simply by reading the comments and skipping the code. The pseudo-code is more for the very interested reader who might want to re-implement the system.

Two more comments for those readers interested in the pseudo-code: (1) The pseudo-code is written in loose variant of the lambda calculus.[80] Roughly speaking, it is written in a functional programming style using a restricted form of Lisp. Some of the definitions have been left in English to avoid getting lost in implementation details, but most of the rest of the pseudo-code could be easily translated into a lambda calculus interpreter (e.g., a form of Scheme[81]) and

---

[79] See Larry Wall, Tom Christiansen, and Randal L. Schwartz, *Programming in Perl, Second Edition* (Cambridge, MA: O'Reilly & Associates, Inc., 1996).

[80] Alonzo Church, "The Calculi of Lambda-Conversation," *Annals of Mathematical Studies, Vol. 6*, Princeton University Press, 1941.

[81] W. Clinger, J. Rees, H. Abelson, R.K. Dybvig, C.T. Haynes, G.J. Rozas, N.I. Adams IV, D.P. Friedman, E. Kohlbecker, G.L. Steele Jr., D.H. Bartley, R. Halstead, D. Oxley,

executed as is. Other than some of the English-language-only definitions, I have not abused the lambda calculus notion too much other than to (a) use a couple of loop constructs (e.g., the construct *dolist*); and, (b) overload function definitions so that, in some cases, several functions with the same name have been written, each intended to do the same sort of work for different argument types. Neither of these notational excesses is difficult to clean up either for implementation purposes or for purposes of making the definitions more formal. (2) Although the pseudo-code is more or less a set of well-formed, executable programs, they are quite different programs than the ones that actually constitute the analysis steps of the Conversation Map system. In the actual system a great number of intermediate data structures are maintained to increase the processing speed. In the pseudo-code, most data is defined with one or more function calls. This difference in efficiency could be partially bridged by directly implementing the functions of the pseudo-code as memoized functions (i.e., functions that cache their results so that they do not have to be recomputed when they are called again with the same arguments).[82] However, a great number of details concerning the actual form of data structures have also been elided in the pseudo-code and these need to be included to get reasonable performance from the system. For example, some of the definitions rely on set intersection and set union operations and, in the pseudo-code, it is implied that these operations are operations on lists when, in the actual implementation, sets are implemented as bit vectors and intersection and union handled as bit vector operations. So, the pseudo-code is also supplemented with a variety of comments on some of the more important implementation details.

Steps A through M will be listed again and, beneath the listing of each step, a set of comments will be added.

(A) Messages are threaded.

Message threading is performed by all conventional email readers and newsgroup browsers (e.g., Eudora, Netscape Messenger, RN, etc.), so the fact that the Conversation Map system also does it is of no surprise. Messages are assumed to be in the standardized format defined in RFC1036.[83] In the RFC1036 format every message has a *message-id* slot that is filled with a unique identifier. Also, every message contains a *references* slot. The *references* slot should contain the unique identifier of another message if the message was sent in response to the other message. The conventional method for threading the messages is simply a means to chain together message identifiers by examining the references slots.

G.J. Sussman, G. Brooks, C. Hanson, K.M. Pitman, M. Wand, *Revised(4) Report on the Algorithmic Language Scheme*, http://www-swiss.ai.mit.edu/~jaffer/r4rs_toc.html

[82] See Peter Norvig, *Paradigms of Artificial Intelligence Programming: Case Studies in Common Lisp* (San Mateo, CA: Morgan Kaufmann Publishers, 1992), 269-275.

[83] http://www.rfc-editor.org

However, the Conversation Map system includes an auxiliary method for threading if the *references* slots are not present or have been stripped out by a remailer. Unfortunately, this is common occurrence in non-Usenet newsgroups; e.g., in newsgroups administered from a private listserv. To handle these occurrences, the contents of the *subject* slots of the messages are compared. If the subject lines of two messages match and message1 was sent before message2, then message2 is assumed to be a reply to message1. The *references* slots of the messages are reconstructed in this manner. Matching subject lines requires a bit of flexibility in the matcher so that if message1's subject line reads "important news today" and message2's reads "RE: important news today," the two should, nonetheless, be matchable to one another. Also, determining if message2 was sent after message1 requires a comparison of the dates and times listed in the messages' respective *date* slots. Since messages can be sent from different time zones, the date comparison routines normalize the dates to some standard time (e.g., GMT) before comparison. The addition of this auxiliary threading method makes the Conversation Map system capable of partially threading messages that cannot be threaded by most conventional newsgroup readers.

(B) Quotations are identified and their sources (in other messages) are found.

One might imagine that finding the quotations in a message would be as simply as looking for those lines in the message that are either indented or which have some sequence of punctuation marks prefixed to the quoted line. Typically, quoted texts in messages look something like the following:

```
>>> On 31 February 2001 Warren Sack
>>>                <wsack@media.mit.edu>                wrote:
>>> Hi guys!  The future is really great!
```

Unfortunately, simple methods (like look for the lines prefixed with punctuation marks) do not work very well for two reasons (1) even though such a method might be able to identify some of the quotations in the message, such a method does not resolve the question of where the quotation comes from (i.e., which message was the quoted section taken from?); (2) many conventional newsgroups readers break original lines into multiple lines, some of them do not prefix quoted lines with any punctuation, and some introduce empty lines between existing lines in the quoted material; moreover, message authors, when quoting material, oftentimes rearrange its textual layout so that it does not easily match the original message lines.

Consequently, what is needed is a more sophisticated procedure for identifying quotations. Humanists and copyright legislators (among others) have debated for centuries over the question of what constitutes a quotation, so it should come as no surprise that this matter is not as simple as it might first appear.

The general idea behind the procedure proposed below and implemented in the Conversation Map system is this: to determine if a line has been quoted from somewhere else, look at all of the messages that precede the current one in the current thread and try to find lines in previous messages that can be almost, if not exactly, matched to the current line. Also, the system uses punctuation prefixes as an auxiliary method, so, if line A has been determined to be a quote from some previous message M, and B is of unknown origin, but has the same punctuated prefix as A, then a guess is made that B comes from M too.

It might appear to be the case that this method is too expensive to be practically useful. However, there is an implementation detail not mentioned in the pseudo-code below that makes the method practical. To identify quotes, put the lines of the current message and the lines of all the messages that precede the current one in a file. At the end of each line, mark its source message. Delete any punctuation prefixes on the lines and then sort them alphabetically. If two lines are similar and have different sources, then the oldest message can be determined to be the origin for the line found in a newer message.

One conceptualization which will be repeatedly used in the explanations which follow is this: message threads are conceived of as "trees" where the first posted message is the root, replies to the first are its "children" (i.e., they are the next ply down in the tree), replies to replies are children of the children, etc. Conversely, every message except the first one in a thread has a "parent," i.e., a message it was sent in response to. So, to find all of the messages that preceded the current one in a thread, one simply needs to follow the "parent" pointers to the oldest "ancestor," the "root" or first message of the thread.

Immediately below is a set of pseudo-code definitions to more precisely define what it means to find the source of a quoted line. Anyone who has already read the rough description given above will not get many more immediate insights about the quotation identification procedure by reading the pseudo-code below. However, the pseudo-code below defines a function *source* which, when given a line of text from a message, computes where it comes from. This function is necessary to the main results of the chapter concerning the computation of lexical ties, the definition of social cohesion, and the use of social cohesion links to sort the output of a corpus-based linguistics procedure. So, the extent to which these main contributions can be articulated precisely depends upon a precise rendering of the function *source* to rigorously define what is meant by quotation in email messages.

The function *source* will be defined in pseudo-code, but so, presently, will a set of other functions. The definition of source and all of the subsequent definitions rely on a small set of preliminary definitions which are stated immediately below. After these "utility" functions have been defined, source and its associated subfunctions is defined.

Preliminary Definitions

```
(define (length list)
  "Return the number of elements in the list.")

(define (min i j)
  "Given two numbers return the smallest.")

(define (max i j)
  "Given two numbers return the largest.")

(define (append x y)
  "Given two two lists, x and y,"
  "make one list in which the elements"
  " of x come before the elements of y.")

(define (union x y)
  "Combine the lists x and y into a single set.")

(define (intersection x y)
  "Returns the elements that occur in"
  "both x and y.")

(define (member element list)
  "Return true if the list contains the"
  "given element.")

(define (sort predicate list)
  "Use the predicate to sort the list."
  "Thus, for example, if the predicate is <"
  "and the list is a list of numbers"
  "sort would return a new list in which"
  "the smallest number is at the"
  "beginning and the largest at the end.")

(define (string<= string1 string2)
  "Return true if, in alphabetical order,"
  "string1 comes before string2."
  "Return false otherwise.")

(define (string= string1 string2)
  "Return true if string1 and string2 are equal.")

(define (string-append strings)
  "Accept a list of strings and"
  "append them all together into"
  "one string.")

(define (includes-nonwhitespace? string)
  "Return true if the string contains"
  "alphanumeric characters or any other"
  "characters that are visible; i.e.,"
  "characters that are not spaces, tabs,"
```

```
  "newlines, etc.")

(define (front-trim paragraph)
  "Given a list of strings, delete all of the strings at the"
  "beginning of the list that contain only whitespace."
  (if (not (includes-nonwhitespace? (car paragraph)))
      (front-trim (cdr paragraph))
      paragraph))

(define (back-trim paragraph)
  "Given a list of strings, delete all of the strings at the"
  "end of the list that contain only whitespace."
  (reverse (front-trim (reverse paragraph))))

(define (trim paragraph)
  "Given a list of strings, delete all of the strings at the"
  "beginning and end of the list that contain only whitespace."
  (front-trim (back-trim paragraph)))

(define (filter predicate list)
  "Remove all elements of a list that do not satisfy"
  "the given predicate."
  (if (predicate (car list))
      (cons (first list)
            (filter predicate (cdr list)))
      (filter predicate (cdr list))))

(define (position element list)
  "Return the position in the list, coded as an integer,"
  "where the element first occurs.  The first element of"
  "a list is in position 0."
  (let ((position -1))
    (dolist (e list)
        (set position (+ 1 position))
        (if (equal e element)
            (return position)))
    -1))
```

Definition of the function Source and its necessary subfunctions

```
(define (lines message)
  "Return the lines of text in the body of the message.")

(define (message line)
  "Return the message from which the line of text comes from.")

(define (parent message)
  "Return the message that the given message was sent in reply to.")

(define (children message)
  "Return the messages that were replies to the given message.")
```

```
(define (behead line)
  "Removes the prefix of any"
  "punctuation marks from the string"
  "and return the remainder of the string.")

(define (prefix line)
  "Return the substring of punctuation marks"
  "that prefixes the given line.")

(define (headless-string= string1 string2)
  "Compare string1 and string2 after removing"
  "any punctuation mark prefixes."
  "See if string1 and string2 are equal."
  (string= (behead string1) (behead string2)))

(define (headless-string<= string1 string2)
  "Compare the alphabetic ordering of string1"
  "and string2 after removing any punctuation"
  "mark prefixes.  See if string1 comes before"
  "string2."
  (string<= (behead string1) (behead string2)))

(define (ancestor? message1 message2)
  "See if message2 is a reply (i.e., is a 'child') of"
  "message1.  If not, see if it is a child of one of"
  "message1's children, children's children, etc."
  "If so, return true, otherwise, return false."
  (if (member message2 (children message1))
      true
      (dolist (child (children message1))
          (if (ancestor? child message2)
              (return true)
              false))))

(define (oldest-ancestor message)
  "Find the 'root' (i.e., the initial message posted)"
  "of the thread containing the message."
  (if (parent message)
      (oldest-ancestor (parent message))
      message))

(define (headless-matches line lines)
  "Remove any of the lines that don't match"
  "the given line."
  (filter (lambda (l) (headless-string= line l))
          lines))
```

```
(define (threads-from message)
  "Given a message, return a list of lists of messages."
  "Every list of messages returned is a thread.  Starting with"
  "the initial message, each message in a thread is"
  "a response to the message that precedes it in the list."
  (if (children message)
      (let ((threads ()))
        (dolist (child (children message))
          (let ((subthreads (threads-from child)))
            (dolist (subthread subthreads)
              (set threads (cons (cons message subthread)
                                 threads)))))
        threads)
      (list message)))

(define (containing-thread message)
  "Return the thread that contains the message."
  (let ((threads (threads-from (oldest-ancestor message))))
    (dolist (thread threads)
      (if (member message thread)
          (return thread)))
    nil))

(define (source line)
  "Given a line, that might or might not be a quote,"
  "return the line itself if it is not a quote; otherwise,"
  "return the line from the message from which it"
  "was quoted; i.e., return its original 'source.'"
  (let ((thread (containing-thread (message line)))
        (lines ()))
    (dolist (message thread)
      (append (lines message) lines))
    (car (sort ancestor? (headless-matches line lines)))))
```

Note that, in practice, a function like *source* does not always work because lines in previous messages are not matchable to lines in the current message (because, for instance, the ones in the current message are pieces of lines or lines reorganized from previous messages). In such a case, an auxiliary function, *prefix-source*, is employed that uses the prefix of the current line along with a database of known prefixes and sources of other lines in the message.

```
(define (prefix-source line message)
  "Given a line and a message, if another line from"
  "the message has the same prefix and has a known"
  "source, then guess that the given line has the same"
  "'source.'"
  (let ((prefix (prefix line))
        (lines (lines message)))
    (dolist (l lines)
      (if (and (string= prefix (prefix l))
               (source l))
          (return (source l))))
```

nil))

(C) A table of newsgroup participants (i.e., those people who have posted to the newsgroup) to messages is built.

This step is easily implemented by a function that examines the header of each message and records the *message-id* of each message under the name of the participant listed in the *from* slot of the message.

(D) For every participant, the set of all other participants who replied to the participant is recorded. Similarly, for every participant, the set of all other participants who quoted from the messages of the participant is recorded. Participants who reciprocally reply to and/or quote from one another's messages are linked together into a social network.

After step (B) and step (C) have been completed a database that records threading and citation dependencies between messages has been built. This database is used to create a new database where the dependencies recorded are dependencies between the authors of the messages rather than simply between the messages.

(E) The "signatures" of participants are identified and distinguished from the rest of the contents of each message.

The means used to perform this step is largely analogous to the means used to find quotations. All of the messages authored by a given participant are put into one file, the file is sorted, and repeated sequences of lines are found. If some sequence of lines is repeated in every message authored by a participant, then that sequence is identified as the author's signature.

In practice this method works quite well. The naïve method for finding signatures (i.e., look at the bottom of messages to see if there is something there that looks like a signature) does not work because signatures do not necessary resemble one another. Some people sign their messages with aphorisms, quotations, and/or advertisements. Some people just sign their name. Some people include elaborate ASCII art designs in every message sent. In short, a signature does not necessary look like something in particular. Again, like with quotations, it is no surprise to find that signatures are hard to recognize: this issue has been under examination and debate for centuries in philosophy, linguistics, and legal studies.[84]

The identification method used in the Conversation Map is defined operationally to be that text which a participant includes in every message sent. Usually this operational definition finds what most would call a signature (at least, that is,

[84] E.g., Jacques Derrida, "Signature Event Context," in *Limited Inc* (Evanston, IL: Northwestern University Press, 1988).

what people on the Internet recognize as a signature), however, sometimes it yields non-intuitive results. For instance, sometimes the only regular feature of a participant's messages is the repetition of a salutation like "Hi everybody!" The system picks these regularities as signatures.

It is important for the system to be able to find signatures so that, when it is computing connections between messages and participants (e.g., when it is computing what will be presently defined as social cohesion) it does not count a person's signature as some sort of commentary on a quote that immediately precedes it.

Again, since the signature identification procedure is important for the main results – at least insofar as the main results of this chapter depend upon a means for identifying quotations and signatures – a pseudo-code definition for a function called *signature-line?* is presented below.

```
(define (from message)
  "Given a message return the name of the participant who"
  "posted the message.")

(define (messages-from participant)
  "Given the name of a participant, returns all messages"
  "in the corpus authored by the participant.")

(define (line-number line)
  "Return the position in the message of the given line.")

(define (common-lines messages)
  "Find all of the lines of text that appear in every"
  "one of the messages."
  (let ((common-lines (lines (car messages))))
    (dolist (message messages)
        (set common-lines
            (intersection common-lines
                          (lines message))))
    common-lines))

(define (contiguous? line paragraph)
  "Return true if the line appears above the paragraph"
  "in a message."
  (and (equal (message paragraph)
              (message line))
       (= (- (line-number (car paragraph)) 1)
          (line-number line))))
```

```
(define (select-signature paragraphs)
  "Given a list of lists of lines -- i.e., paragraphs --"
  "trim each list of lines (see above for the definition"
  "of trim) and then pick the longest list of lines"
  "as a potential signature."
  (let ((trimmed-paragraphs ()))
    (dolist (paragraph paragraphs)
        (set trimmed-paragraphs
             (cons (trim paragraph)
                   trimmed-paragraphs)))
    (car (sort (lambda (p1 p2)
                  (>= (length p1)
                      (length p2)))
               (filter (lambda (p) p)
                       trimmed-paragraphs)))))

(define (signatures participant)
  "Look through all the messages posted by the participant"
  "to identify the lines that occur in every message and"
  "which might, therefore, constitute the participant's signature."
  (let* ((messages (messages-from participant))
         (common-lines (common-lines messages))
         (signatures ()))
    (dolist (message messages)
        (let ((next-paragraph ())
              (paragraphs ()))
          (dolist (line (lines message))
              (if (member line common-lines)
                  (cond ((contiguous? line next-paragraph)
                         (set next-paragraph
                              (cons line next-paragraph)))
                        (else
                         (set paragraphs
                              (cons (reverse next-paragraph)
                                    paragraphs))
                         (set next-paragraph
                              (list line))))))
          (if next-paragraph
              (set paragraphs
                   (cons (reverse next-paragraph)
                         paragraphs)))
          (set signatures
               (cons (select-signature paragraphs)
                     signatures))))
    signatures))

(define (signature-line? line)
  "Given a line, determine whether it is a line from a participant's signature."
  (let* ((participant (from (message line)))
         (signatures participant))
      (dolist (signature signatures)
          (if (member line signature)
              (return true)))
```

false))

(F) The words in the messages are divided into sentences.

This is the one step in the Conversation Map system that has been implemented using an existing tool. The tool described in the following paper is used: Jeffrey C. Reynar and Adwait Ratnaparkhi. "A Maximum Entropy Approach to Identifying Sentence Boundaries." In *Proceedings of the Fifth Conference on Applied Natural Language Processing*, March 31-April 3, 1997. Washington, D.C. The tool is implemented in Java 1.1 and is interfaced to the rest of the analysis procedure (which is implemented in Perl) by having the tool write the results of the end-of-sentence analysis out to a file. The file is then read in by the Perl programs. People who are not computational linguists often underestimate the difficulty of this step too. Suffice it to say that Reynar and Ratnaparkhi's tool does far better than any naïve approach (e.g., a procedure which just looks for periods). This problem is especially hard in the often agrammatical or alternative-grammatical world of online prose.

(G) Sentences are divided into clauses according to punctuation (e.g., parenthetical phrases are subordinated to main clauses) so that the parser (applied in step L) does not need to handle punctuation.

This is a stop-gap measure necessary because most contemporary parsers do not include methods for parsing punctuation. Integrating punctuation into parsing rules is currently a cutting edge research area that is far from solved.[85]

(H) Discourse markers (e.g., connecting words like "if", "therefore", "consequently", etc.) are tagged in the messages.

We use a list of discourse markers compiled by Daniel Marcu.[86] The analysis method used to identify the discourse markers is a simple one based on standard methods in the literature.[87] More or less, the method boils down to this: look for a word or phrase from a database of discourse markers at the beginning of sentences and/or close to punctuation marks like commas and semi-colons.

These discourse markers could be used for something interesting. E.g., Marcu uses them as an essential ingredient in his computer program that can parse a text into rhetorical relations between parts of the text. In a future version of the Conversation Map system something will be done with them. Right now, they

---

[85] See, for instance, Christine Doran, *Incorporating Punctuation into the Sentence Grammar: A Lexicalized Tree Adjoining Grammar Perspective*, Ph.D. Dissertation and Technical Report IRCS-98-24 (Philadelphia, PA: University of Pennsylvania, 1998).

[86] Marcu.

[87] E.g., Diane J. Litman. "Cue Phrase Classification Using Machine Learning," *Journal of Artificial Intelligence Research*, 5:53-94, 1996.

are being identified mostly so that the parser (explained in step L) does not spend any time trying to parse them into noun or verb phrases.

(I) Every word of every message is tagged according to its part-of-speech (e.g., "noun", "verb" "adjective", etc.)

A simple trigram-based tagger is used to accomplish the part-of-speech tagging.[88] In earlier versions of the Conversation Map system several existing tools to perform this step were used. However, a variety of rather uninteresting implementation issues made it necessary to have the rest of the analysis procedure interact with the tagger in a manner that was difficult or impossible with the existing tools. Consequently, I built new part-of-speech tagger. It is a simple Perl program that uses a database of trigrams and other statistics that were automatically compiled from a the Linguistic Data Consortium's distributed version of the hand-tagged Wall Street Journal Corpus of several hundred thousand words. The Perl program runs as a server and the rest of the analysis procedure uses a specially-built client to interact with it.

(J) Every word is morphologically analyzed and its root is recorded.

The morphological analyzer is also a newly built piece of software, but it relies on a database containing morphological and syntactic information that was created at the University of Pennsylvania.[89] It is mildly interesting to note why the software that was available for using the database was not integrated into the Conversation Map system. In 1992, when the database software was published by the team at the University of Pennsylvania who created the database, machines were rather slow and disk space rather small, so the software built for the database was a very sophisticated and hard to port application (i.e., hard to move to a different kind of machine). After spending a couple of days trying to port the software I realized that today's machines would allow the database to be simply kept as one big alphabetically-sorted list and that the look-up procedure could be about 10 lines of simple Perl code (rather than thousands of lines of C). The Perl code that implements the morphological analysis procedure is more or less this piece of code:

```
use Search::Dict;
open(MORPHOLOGY, "sorted_english_morphology_dictionary.txt");
print "\n morphology>> ";
while(<>) {
    look(*MORPHOLOGY,$_,0,0);
    $entry = <MORPHOLOGY>;
    print $entry."\n";
    print "\n morphology>> ";
```

[88] For a more detailed explanation of this method see, for instance, Eugene Charniak, *Statistical Language Learning* (Cambridge, MA: MIT Press, 1993), 45-51.
[89] Karp, Schabes, Zaidel and Egedi.

(K) Every noun is "term expanded" using the thesaurus WordNet. Recorded along with each noun are its synonyms, antonyms, hypernyms (generalizations), hyponyms (specializations), meronyms (parts; e.g., a table has legs and a top), and holonyms (wholes for which the noun might be a part; e.g., a leg can be a part of a table).

This step is essentially to the definition of the lexical cohesion analysis procedure defined below. To implement this step a package of Perl procedures has been constructed to access the WordNet database.

(L) The words of the messages are parsed into sentences using a partial parser.

The partial parser is a re-implementation and revision of the parser described by Gregory Grefenstette.[90] The parser was designed with the intention that its output would be used as the input to a statistical analysis procedure for large text corpora. Consequently, the parser is fast and simple.

It needs to be fast to be able to handle large amounts of text in a reasonable amount of time. Its output needs to be simple so that the patterns of text it identifies will recur a significant number of times. Before re-building this parser, I experimented with several more sophisticated parsers. Unfortunately they were both too slow and their output was so heavily annotated with special syntax tags that it was difficult to determine if, for example, one noun phrase was similar to another because the noun phrases had so much internal structure to them. This is not to say that the current parser is ideal. It could be improved upon and future versions of the system will implement a more sophisticated parser.

The current parser can be understood as a four step process: (1) words of a sentence are gathered into noun, prepositional, and verb phrases; (2) verb phrases are assigned a type (active, passive, or attributive); (3) the noun and prepositional phrases are examined to see if they function as subject, direct objects, or indirect object; and, (4) the annotated noun and prepositional phrases are traversed and individual nouns are assigned a set of attributes (e.g., subject of a given verb, object of a certain verb, modified by a particular adjective, etc.)

The first step of the parser can be understood as a simple grammar which is run on every sentence to gather words into phrases. The grammar rules use the part-of-speech tags assigned to the words in step (i).

---

[90] Grefenstette.

```
noun_phrase_or_prepositional_phrase     →    preposition noun_phrase
noun_phrase_or_prepositional_phrase     →    noun_phrase
noun_phrase          →    noun
noun_phrase          →    proper_noun
noun_phrase          →    pronoun
noun_phrase          →    verb_present_particle_gerund
noun_phrase          →    adjective
noun_phrase          →    cardinal_numeral
noun_phrase          →    determiner
noun_phrase          →    interjection
noun_phrase          →    noun_phrase noun_phrase_or_prepositional_phrase
verb_phrase          →    verb
verb_phrase          →    infinitive_marker
verb_phrase          →    verb_present_particle_gerund
verb_phrase          →    modified_verb_phrase
modified_verb_phrase      →    adverb verb_phrase
modified_verb_phrase      →    auxiliary_verb verb_phrase
modified_verb_phrase      →    verb_phrase conjunct verb_phrase
modified_verb_phrase      →    verb_phrase verb_phrase
```

Obviously, this grammar has a lot of shortcomings, but since the goal is to analyze large amounts of text in order to get a rough overview, a summary, of how people are talking about certain subjects, the shortcomings of the grammar – from a linguistic point of view – are not necessarily the same as its shortcomings from the perspective of the overall goal of the system.

After the words of a sentence have been gathered into phrases, each verb phrase is assigned a type. The type assignment routine is also simple. It looks at the last verb in a verb phrase to see if it is a form of the verb "to be." If it is, then the sentence is assigned to be of type *attributive*. Otherwise, it is assumed to be of type *active* unless a form of "to be" is found somewhere in the middle of the verb phrase, in which case it is assigned to be of type *passive*. Here is the pseudo-code for the verb phrase type attribution routine:

```
(define (tag word)
  "Return the part-of-speech tag for the word.")

(define (verb-phrase-type vp)
  (if (equal (tag (last vp)) 'infinitive_marker)
      'attributive
      (let ((type 'active))
        (dolist (word vp)
          (cond ((equal (tag word)
                        'infinitive_marker)
                 (set type 'passive))
                ((equal (tag word)
                        'verb_present_particle_gerund)
                 (set type 'active)))
          type))))
```

After the verb phrases have been assigned a type another set of rules is run to assign the noun and prepositional phrases their roles of subject, direct object or indirect object. These rules are approximated by the following grammar:

```
a_verb_phrase →     active_verb_phrase
a_verb_phrase →     attributive_verb_phrase
sentence            →   subject a_verb_phrase direct_object
sentence            →   subject a_verb_phrase direct_object indirect_object
sentence            →   indirect_object passive_verb_phrase subject
```

Finally, the noun and prepositional phrases are traversed and attributes are assigned to individual nouns. For example, the head noun of a noun phrase (usually the last noun of a noun phrase) is assigned the attribute "subject_for_jump" if the noun_phrase is the subject phrase for a verb phrase in which a form of "to jump" is the main verb of the verb phrase. For example, in the verb phrase "Jack jumped" "Jack" is the head noun and is the subject of the verb "to jump."

The composite of all of these assignments make up what will here be called the *lexicosyntactic profile* for the noun. A noun's lexicosyntactic profile is a long list of word and grammatical relation pairs (e.g., subject_for_jump, object_for_love, etc.). Each pair is labeled with frequency data; i.e., the number of times in which the noun appeared in the context of the lexicosyntactic pair. These profiles are stored as vectors of integers.

When these heuristics of the parser fail, the overall performance of the system becomes similar to one where no parsing is performed. Instead, in many systems for automatic thesaurus compilation (a step of the procedure described next) "text windowing" is employed instead of parsing. "Text windowing" techniques sample, for example, the two words preceding and the two words following every noun in order to build a profile for every noun. This profile is then used like the lexicosyntactic profiles generated by the parser. So, even in the worst case, the parser's output simulates text windowing, a technique found to be useful for the summaries that are top-level goal for the Conversation Map system.

(M)The lexicosyntactic profile of every noun in the archive is compared to the lexicosyntactic context of every other noun in the archive. Nouns that are used or discussed in the same manner are calculated to be similar and are placed close to one another in the semantic networks. One can understand this semantic network as a crude approximation to the sorts of metaphors of discourse identified by linguists like George Lakoff and Mark Johnson.[91]

[91] Lakoff and Johnson.

To compare the profiles, An algorithm similar to the one described in Gregory Grefenstette is used.[92] The method used is a well-known one in the literature of information retrieval although here it is being used on vectors of lexicosyntactic pairs rather than on vectors of simple words. The following pseudo-code defines two functions and a set of other necessary subfunctions and terms. The function *similarity* returns a numerical weight to represent how similar two terms are given a parsed corpus where every term – i.e., every noun – has a lexicosyntactic profile that has been built for it. The function *similarity* implements a weighted Jaccard measure. The function *most-similar*, given a noun, finds the other noun with a lexicosyntactic profile most similar to its own profile.

```
(define terms
  "Terms is a constant that lists all of the unique nouns in a given corpus.")

(define attributes
  "A constant that lists all of the unique attributes (e.g., verbs and adjectives)"
  "in a given corpus.")

(define (attributes term)
  "Return the list of unique attributes for the term.")

(define nao
  "Nao is an integer constant that holds the total number of attribute occurrences.")

(define (freq term attribute)
  "Return the number of times the term and the attribute occurred together.")

(define nterms (length terms)
  "Nterms is a constant that records the"
  "total number of unique nouns in the corpus.")

(define nattributes (length attributes)
  "Nattributes is a constant that records the"
  "total number of unique attributes in the corpus.")

(define (p term attribute)
  (/ (freq term attribute)
     (length (attributes term))))

(define (global-weight attribute)
  (let ((weight 0))
    (dolist (term terms)
       (let ((p (p term attribute)))
          (set weight
               (+ weight
                  (/ (* p (log p))
                     (log nao))))))
    (- 1 weight)))
```

[92] Grefenstette.

```
(define (local-weight term attribute)
  (* (global-weight attribute)
     (log (+ 1 (freq term attribute)))))

(define (similarity term1 term2)
  (let ((ia (intersection (attributes term1)
                          (attributes term2)))
        (min-sum 0)
        (ua (union (attributes term1)
                   (attributes term2)))
        (max-sum 0))
    (dolist (attribute ia)
        (set min-sum
             (+ min-sum
                (min (local-weight term1 attribute)
                     (local-weight term2 attribute)))))
    (dolist (attribute ua)
        (set max-sum
             (+ max-sum
                (max (local-weight term1 attribute)
                     (local-weight term2 attribute)))))
    (/ min-sum max-sum)))

(define (most-similar term)
  (let ((max-sim 0)
        (nearest nil))
    (dolist (next terms)
        (if (and (not (equal next term))
                 (> (similarity next term) max_sim))
            (set nearest next)))
    nearest))
```

The output of this step provides a means for beginning to understand which terms are comparable to which other terms in a given VLSC. The output of the Conversation Map system is a far cry from the reasoned opinions of analysts, like Lakoff and Johnson, who describe how terms in a discourse are comparable. However, it does provide an automated means for one to begin to pose the sorts of hypotheses that analysts like Lakoff and Johnson pose.

Finally, with the preprocessing now described, the three main results of the chapter can be stated: (1) inter-message lexical cohesion; (2) social cohesion; and, (3) social network-based sorting.

Inter-Message Lexical Cohesion

An algorithm has been created for the computation of *lexical ties* between email messages that takes into account information about both threading and quotation in the messages (thereby producing an approximation of the themes of conversation). Most previous computational work on lexical cohesion analysis has been done on documents like newspaper stories or encyclopedia articles

where the main goal was to determine which words, within a document, might be thematically tied to which other words. Consequently, the structure of the document *per se* was not very complicated. Lexical ties were searched for between sentences and the distance between sentences was easy to define. Even in newer work on inter-document coreference[93] the structure of the documents and the relationships between the parts of related documents have not been too complicated. Quotations, signatures, and threading of messages makes the case of lexical tie analysis for email messages more complicated because the structure of documents is more complicated than these other sorts of texts. Admittedly, one could simply ignore the internal structure of the messages, but this would efface the information necessary for the higher goal pursued here: to garner some information about which themes of conversation link which people in an online discussion. For example, one might simply count the quotations within a message as text original to the message, but that simplification would ignore the fact that a certain participant had quoted and commented on the text of another participant. Consequently, quotations, signatures, and threading must be taken into account for the definition of lexical ties in online conversations conducted through email exchanges.

In what follows a more precise definition of lexical ties for threaded email messages is presented in commented pseudo-code form. Given the previously defined functions for identifying signatures and quotations, the definition of lexical tie is relatively straightforward, although a little longwinded because of the multiple text segments within messages (here called "paragraphs") that must be negotiated.

First, we define how a message can be broken up into paragraphs. Text that is a part of a participant's signature is ignored. Contiguous text that is quoted from other messages is grouped together into a paragraph. Contiguous text that is from the current message is also grouped together into a paragraph. Participants often author messages that are broken into multiple paragraphs: first a quotation from another message, then a comment on the quoted text, then another quotation and a comment on it, etc.

---

[93] E.g., Amit Bagga, Breck Baldwin, and Sara Shelton (editors) *Coreference and Its Applications, Proceedings of the Workshop*, New Brunswick, NJ: Association for Computational Linguistics, 22 June 1999.

```
(define (paragraphs message)
  (let ((next-paragraph ())
        (paragraphs ()))
    (dolist (line (lines message))
      (if (not (signature-line? line))
        (if (and next-paragraph
                 (equal (source (car next-paragraph))
                        (car next-paragraph)))
            (cond ((equal (source line) line)
                   (set next-paragraph (cons line next-paragraph)))
                  (else
                   (set paragraphs
                        (cons (reverse next-paragraph)
                              paragraphs))
                   (set next-paragraph
                        (list line))))
            (cond ((equal (source line) line)
                   (set paragraphs (cons (reverse next-paragraph)
                                         paragraphs))
                   (set next-paragraph
                        (list line)))
                  (else
                   (set next-paragraph (cons line next-paragraph)))))))
    (if next-paragraph
      (set paragraphs (cons (reverse next-paragraph) paragraphs)))
    paragraphs))
```

The next definition is simply a means to access the paragraph that contains a given line of text in a message.

```
(define (paragraph line)
  (let ((paragraphs (paragraphs (message line))))
    (dolist (paragraph paragraphs)
      (if (member line paragraph)
        (return paragraph)))))
```

Given a line of text, find the paragraph that follows it (i.e., find the paragraph that follows the paragraph that contains the given line of text).

```
(define (following-paragraph line)
  (let* ((paragraphs (paragraphs (message line)))
         (paragraph (paragraph line))
         (position (position paragraph paragraphs)))
    (if (< position (length paragraphs))
        (nth (+ 1 position) paragraphs)
        nil)))
```

With these subfunctions, a definition of lexical tie can now be stated for message-to-message links. However, since a lexical tie between messages is ultimately defined by a set of lexical ties between terms in one message and

terms in the second message, definitions of lexical tie for smaller text segments than the message must also be given.

First a definition of lexical tie is given for two tokens. This is a very conservative function insofar as terms are only expanded to the set of terms immediately related to them in the machine-readable thesaurus WordNet. Most previous work on lexical cohesion includes a more expansive definition of term expansion. If the expansions of the terms intersect, then a lexical tie is said to exist between them.

```
(define (lexical-ties term1 term2)
  (let ((term1-synonyms (synonyms term1)))
    (cond ((equal term1 term2)
           (list term1))
          ((intersection term1-synonyms (synonyms term2))
           (intersection term1-synonyms (synonyms term2)))
          ((intersection term1-synonyms (antonyms term2))
           (intersection term1-synonyms (antonyms term2)))
          ((intersection term1-synonyms (hypernyms term2))
           (intersection term1-synonyms (hypernyms term2)))
          ((intersection term1-synonyms (hyponyms term2))
           (intersection term1-synonyms (hyponyms term2)))
          ((intersection term1-synonyms (holonyms term2))
           (intersection term1-synonyms (holonyms term2)))
          ((intersection term1-synonyms (meronyms term2))
           (intersection term1-synonyms (meronyms term2)))
          (else
           ()))))
```

A line-to-line definition of lexical tie uses the above definition to compare every noun in line1 with every noun in line2.

```
(define (lexical-ties line1 line2)
  (let ((lexical-ties ()))
    (dolist (term1 (nouns line1))
      (dolist (term2 (nouns line2))
        (set lexical-ties
             (union (lexical-ties term1 term2)
                    lexical-ties))))
    lexical-ties))
```

The line-to-paragraph definition uses, analogously, the definition of line-to-line lexical ties. And then, the paragraph-to-paragraph definition simply generalizes the definition one step further.

```
(define (lexical-ties line paragraph)
  (let ((lexical-ties ()))
    (dolist (l paragraph)
      (set lexical-ties
           (union (lexical-ties line l)
                  lexical-ties)))
    lexical-ties))

(define (lexical-ties paragraph1 paragraph2)
  (let ((lexical-ties ()))
    (dolist (l paragraph1)
      (set lexical-ties
           (union (lexical-ties l paragraph2)
                  lexical-ties)))
    lexical-ties))
```

Finally, we are in a position to define a lexical ties relation between messages. There are two means for messages to be tied lexically. In the first case, there may be lines in message2 that are quoted from message1. In this case, the function searches each paragraph following a quoted paragraph and attempts to find lexical ties between the paragraphs. In the second case, message2 may have been sent in reply to message1. In this case, lexical ties between the last paragraph of message1 and the first paragraph of message2 are searched for.

```
(define (lexical-ties message1 message2)
  (let ((quoted-lines ())
        (last-paragraph-of-1 (last (paragraphs message1)))
        (first-paragraph-of-2 (first (paragraphs message2)))
        (lexical-ties ()))
    (dolist (line (lines message2))
      (let ((source-message (message (source line))))
        (if (equal source-message message1)
            (set quoted-lines (cons line quoted-lines)))))
    (dolist (line quoted-lines)
      (set lexical-ties
           (append (lexical-ties line (following-paragraph line))
                   lexical-ties)))
    (if (and (equal message1 (parent message2))
             (equal (source (car last-paragraph-of-1))
                    (car last-paragraph-of-1))
             (equal (source (car first-paragraph-of-2))
                    (car first-paragraph-of-2)))
        (set lexical-ties
             (union (lexical-ties last-paragraph-of-1
                                  first-paragraph-of-2)
                    lexical-ties)))
    lexical-ties))
```

While the above definition of lexical ties between messages produces some interesting results (see, for instance, the discussion themes menus of those example Conversation Maps discussed in chapter 2), it does not handle all

possible cases. In some cases, the quoted text follows the comments on the quote. Determining when the quote follows the comments or vice-versa would require one more level of sophistication than the definition provided here.

Social Cohesion

A definition of *social cohesion* has been developed.[94] Social cohesion is the overlay of lexical cohesion information on top of a social network (thus creating a rough description of what the conversations between participants are about and who interacts with whom given a certain theme of conversation). As has been repeatedly discussed throughout this chapter, the point of these analyses is to integrate interpersonal, social network, information of the online conversation together with other semantic information. The detailed definition of lexical tie provided above is therefore of interest to this high-level goal insofar as it makes a definition of social cohesion possible.

In the following pseudo-code a novel definition of tie of social cohesion is provided that depends upon the definition of lexical tie between messages given above (and, thus, takes into account information about quotations, signatures, and threading between messages).

First we define the variable *participants* as the list of all those participants who have posted a message to the newsgroup being examined. Step (C) described above provides us with the means for gathering this list of participants.

```
(define participants
  "a list of the name of every participant"
  "who has posted a message to the newsgroup")
```

To find all of the ties which connect participant1 to participant2 we find all of the lexical ties which connect participant1's messages (either through quotation or reply) to participant2's messages. Then we collect all of the ties which connect 2's messages to 1's. The intersection of these connections defines the set of all lexical ties between their messages that have been used in both directions (i.e., the ties are some indication of the discussion themes that have characterized 1's replies to 2 and vice versa).

[94] Sack.

```
(define (social-cohesion-ties participant1 participant2)
 (let ((lexical-ties-from-1-to-2 ())
       (lexical-ties-from-2-to-1 ()))
    (dolist (message1 (messages-from participant1))
       (dolist (message2 (messages-from participant2))
          (set lexical-ties-from-1-to-2
               (union (lexical-ties message1 message2)
                      lexical-ties-from-1-to-2))
          (set lexical-ties-from-2-to-1
               (union (lexical-ties message2 message1)
                      lexical-ties-from-2-to-1))))
    (intersection lexical-ties-from-1-to-2
                  lexical-ties-from-2-to-1)))
```

For completeness, we define the set of all social cohesion ties between all pairs of participants. The resultant list is what is shown, in the Conversation Map interface, in the middle menu, the so-called "discussions themes" menu.

```
(define social-cohesion-ties
 (let ((all-ties ()))
    (dolist (p1 participants)
       (dolist (p2 participants)
          (if (not (equal p1 p2))
             (set all-ties
                  (union all-ties
                         (social-cohesion-ties p1 p2))))))
    all-ties))
```

These ties of social cohesion are far less detailed than the rules of conversation articulated by conversation analysts for small stretches of one-to-one conversation. But, by examining the ties of social cohesion, one can begin to get a feel for the importance of themes of conversation for given, large, online groups.

These social-cohesion-ties are used to label the edges of the social network computed in step (D). Examining the social network labeled with these ties of social cohesion, one can begin to reckon who is talking with whom about what.

Social Network-based Sorting

A means has been designed and implemented to use the computed social cohesion ties to weight the importance of word entries in automatically computed rough draft thesauri (computed in step M). The weight of a social cohesion tie is defined by the number of pairs of participants who it connects together (i.e., for the pairs for whom the tie has been a reciprocated "theme" of conversation). A pseudo-code definition for the weight of a social cohesion tie is as follows:

```
(define (social-weight tie)
  (let ((weight 0))
    (dolist (p1 participants)
      (dolist (p2 participants)
        (if (not (equal p1 p2))
          (let ((ties (social-cohesion-ties p1 p2)))
            (if (member tie ties)
              (set weight (+ 1 weight)))))))
    (/ weight 2)))
```

A more efficient definition would recognize and incorporate the fact that (social-cohesion-ties p1 p2) = (social-cohesion-ties p2 p1), but, as was discussed above, the composition of the pseudo-code is not an attempt to show how these computations are done efficiently.

With this definition of social weight it is now possible to order the social cohesion ties. This ordering is reflected in the order of the ties as they appear in the Conversation Map's "discussion themes" menu. Also, with this definition it is possible to order the term comparisons computed for the rough draft thesaurus of step M.

This ordering by social weight is novel in the literature of computational linguistics. Also, it is a novel contribution in the social network theory where the computational linguistics results have not been related to the computed social networks.

To select a portion of the rough draft thesaurus for display as a semantic network in the upper right hand corner of the Conversation Map interface, every term in the thesaurus is ordered by the *social-weight* function. (To do this first requires the terms be translated into a form comparable to the ties which are, for the most part, defined as WordNet names for sets of synonyms, synset-ids. This is a straightforward translation operation using WordNet.) After the weights have been computed and the terms ordered, the top-ranked term, call it TRT, is taken and a transitive closure operation is performed on it to select the part of the semantic network that surrounds it. The term computed to be *most-similar* to the TRT (see step M for a definition of the function *most-similar*) and all terms for which TRT was the most-similar term constitute the first expansion of TRT. This cluster of terms is recursively expanded by forward chaining with the most-similar and the inverse of the most-similar function until the cluster contains as many

terms as can be connected. The term within this cluster that has the greatest number of most-similar (or inverse most-similar relations; i.e., the term in the network with the largest "degree" to use the terminology of graph theory) is the term that is selected as the "root" of the semantic network. The resultant tree-shaped semantic network is then drawn into the appropriate space of the Conversation Map interface.

Related Work

Several other content-based Usenet newsgroup readers have been built with text analysis procedures simpler than those incorporated into the Conversation Map system discussed in this chapter. For example, Isahara and Ozaku[95] describe an intelligent network news reader that performs a sort of example-based, relevance feedback procedure to select small collections of messages from an archive given an example message. The intelligent network news reader also contains a method for identifying sub-threads within larger threads by analyzing the content of the messages in a thread.[96] However, systems of this sort[97] are mostly concerned with filtering messages rather than with one of the problems addressed by the Conversation Map system: How can all of the messages in an archive be graphically displayed and organized according to content of the messages and the social structure representative of the participants' interactions?

Many of the computational techniques developed for the analysis of Usenet newsgroups do not take the linguistic content of the messages into account at all using, instead, exclusively information that can be garnered from the headers of the messages.[98] Other work does employ some keyword spotting techniques to identify and sort the messages into categories but does not involve the analysis of grammatical or discourse structures.[99]

Work that does use the contents of the messages for analysis often does not take the threading of the messages into account, or, if it does, does not pay

---

[95] Hitoshi Isahara and Hiromi Ozaku. "Intelligent Network News Reader," in *Proceedings of IUI'97* (Orlando, FL: Association for Computing Machinery, 1997).

[96] Kiyotaka Uchimoto, Hiromi Ozaku, and Hitoshi Isahara. "A Method for Identifying Topic-Changing Articles in Discussion-type Newsgroups within the Intelligent Network News Reader HISHO," in *Proceedings of Natural Language Processing Pacific Rim Symposium* (Phuket, Thailand, December 2-4, 1997).

[97] Cf., Beerud Sheth *NEWT: A Learning Approach to Personalized Information Filtering,* Master's Thesis (Cambridge, MA: MIT Media Laboratory, 1993).

[98] See, for example, Marc Smith. "Netscan: Measuring and Mapping the Social Structure of Usenet" Presented at the *17th Annual International Sunbelt Social Network Conference* (San Diego, CA: February 13-16, 1997).

[99] See, for instance, Judith Donath, Karrie Karahalios, and Fernanda Viegas "Visualizing Conversations," in *Proceedings of HICSS-32* (Maui, HI: Association for Computing Machinery, January 5-8, 1999).

attention to the social network produced by newsgroup participants.[100] Or, if the work does take the threading and citation information into account it does not necessarily use any of the linguistic contents of the messages to compute the graphical display.[101]

Research that has combined content analysis with an analysis of co-referencing of messages and discussion participants has often employed non-computational means to categorize the contents of messages.[102] Some of the most interesting work that analyzes message threading, participant interaction, and the form and content of messages is often ethnographically-oriented, sociolinguistic analyses of newsgroup interactions that is done without the assistance of computers and is so, necessarily, based on a reading of only a small handful of messages.[103] Ideally one could program the computer to emulate the latter sort of analysis, but that will require many advances in the field of computational linguistics. What is unique to the text analysis procedures of the Conversation Map system is the automatic construction and combination of social and semantic networks that, together, provide a means for exploring both the social and semantic structure of a VLSC, e.g., a Usenet newsgroup.

The novel text analysis procedures in combination with a graphical interface make the Conversation Map system an example of a new sort of content-based browser. Earlier examples of content-based browsers[104] used simpler text analysis procedures akin those employed in information retrieval systems. New content-based browsers, clients, and readers (like the Conversation Map system) will incorporate more sophisticated text analysis (and probably, eventually, image analysis) techniques.

---

[100] Michael L. Best. "Corporal ecologies and population fitness on the net." *Journal of Artificial Life*, 3(4), 1998.

[101] Steve Cannon and Gong Szeto, Parasite, http://parasite.io360.com/index.html and http://www.cybergeography.org/atlas/topology.html, 1998.

[102] Michael Berthold, Fay Sudweeks, Sid Newton, Richard Coyne. "It makes sense: Using an autoassociative neural network to explore typicality in computer mediated discussions" In *Network and Netplay: Virtual Groups on the Internet,* edited by F. Sudweeks, M. McLaughlin, and S. Rafaeli (Cambridge, MA: AAAI/MIT Press, 1998).

[103] E.g., Susan Herring, Deborah A. Johnson, Tamra DiBenedetto. "'This discussion is going too far!': Male resistance to female participation on the Internet," in *Gender Articulated: Language and the Socially Constructed Self*, edited by K. Hall and M. Bucholtz (New York: Routledge, 1995).

[104] Earl Rennison, "Galaxies of News: An Approach to Visualizing and Understanding Expansive News Landscapes," in *Proceedings of UIST'94* (1994).

Conclusions

A procedure for analyzing and summarizing the archives of an email-based VLSC has been presented. The results of the analysis procedure include (1) a set of lexically labeled social networks; (2) a set of semantic networks; and, (3) a list of lexical ties – or "themes" – of discussion ranked according to their frequency of use in reciprocated conversational exchanges. These results are rendered as graphs and displayed in the Conversation Map interface as graphical summaries of the newsgroup analyzed and can be used to browse and navigate the archive of analyzed messages.

PART TWO: THEORY

ETHICS

The good of new computer network-based technologies for conversation is that they can potentially provide us with new means for free speech and therefore new powers of self-governance. If these new forms of conversation work effectively as new technologies of self they should make it possible for us to better understand how and where we are located in a variety of social and linguistic networks. Furthermore, by making these networks visible they should give us the tools to plan how we want our interconnections to be between us, other people, and non-human entities, like the machines we depend upon everyday. I argue for a means to visualize the emergent social and semantic networks that are now hard to perceive. By making them visible for everyone to see, it is hoped that it will be possible for them to be used as resources for public discussion. Who is excluded from the linked social networks of public view? Who is central? Why? What are the predominant associations connected with the signs of public life and public discourse? Who and what are accorded positive even glorious associations? Who are given more modest or even negative associations? Why? Why not? A great amount of effort is currently devoted to gathering online information about people for the work of marketing, advertising, etc. At the same time, many people are building home pages on the World Wide Web and creating other electronic projections of themselves. The ethics of software design must examine how these two activities coincide and sometimes clash.

In addition to ecology, the question of ethics of media and the future direction of new communication technologies of artificial intelligence and command-and-control constitutes one of the two axes in which to rethink the idea of progress for today's planet.[105]

Ideally any new technology for conversation should support free speech. It is therefore the ethics of free speech technologies and techniques that are the foci of this chapter.

Amendment I: Congress shall make no law respecting an establishment of religion, or prohibiting the free exercise thereof; or abridging the freedom of speech, or of the press; or the right of the people peaceably to assemble, and to petition the Government for a redress of grievances.

The Constitution of the United States is a document that attempts to lay the ground rules for a free and equal society. It is a document of political and social organization. The functions of free speech and a free press are means to allow the people to organize and sort themselves into groups, parties, and coalitions for choice and action. In other words, the First Amendment of the Constitution does not so much order society as much as it establishes conditions necessary for society to order itself.

Computer technologies – especially network technologies – have similar powers to order people or to provide the means to allow people to order themselves. Phil Agre -- using some terminology from Larry Lessig – describes the emerging world of computational politics like this:

As increasingly complex social relationships are mediated by networked information technology, we are becoming accustomed to the idea that the protocols of these mediated interactions -- the "code" in Larry Lessig's terms -- constitutes a set of working rules... Computers, like institutions generally, both enable and constrain, and both computers and institutions are, in one important aspect anyway, discourses made material -- made, that is, into machinery that governs to some degree the lives of the people who use it.[106]

Some theorists imagine that this "machinery that governs" can be nothing other than antidemocratic. Oscar Gandy is a strong champion of this pessimistic view of technological determinism. He examines the information gathering activities of

---

[105] *La question de l'éthique des médias et de l'orientation prospective de nouvelles technologies de communication, d'intelligence artificielle et de commande constitue, avec la problématique écologique, un des deux axes de recomposition d'une pensée de progrès pour la planète d'aujourd'hui.* Félix Guattari, "Pour une éthique des médias," *Le Monde*, Mecredi 6 novembre 1991, p. 2.

[106] Phil Agre. "Growing a Democratic Culture: John Commons on the Wiring of Civil Society" presented at *Media in Transition* (Cambridge, MA: MIT, October 1999). See http://dlis.gseis.ucla.edu/people/pagre/commons.html

corporations and governments aimed at profiling their customers and citizens respectively and sees in these profiling activities a new, totalitarian form of surveillance. Borrowing a phrase from Michel Foucault[107] and Jeremy Bentham[108], he calls these new information technologies for profiling and sorting people into categories the *Panoptic Sort*:

Can the technologies developed during the control revolution in late capitalism be transformed to serve a democratic purpose, or is such a system of control inherently antidemocratic? Does workplace democracy need a system of disciplinary surveillance? Does a democratic public sphere need political strategists armed with sharply focused citizen profiles? Does an efficient market need consumer research? Advertising and promotion? Segmentation and marketing? Or are these activities incompatible, mutually inconsistent, contradictory, and antagonistic to the notion of free acting, fully informed rational producers and consumers?

It has been and remains my view that the panoptic sort is an antidemocratic system of control that cannot be transformed because it can serve no purpose other than that for which it was designed – the rationalization and control of human existence.[109]

The silver lining in this gloomy pronouncement is Gandy's subscription to the powers of design. He states that the "system of control cannot be transformed because it can serve no purpose other than that for which it was designed." Such a statement implies that a differently designed computer technology – a democratic, anti-panoptic sort technology -- would necessarily "serve no purpose other than that for which it was designed," namely the purposes of democracy. Such a technology might be designed – not to automatically sort people, but rather – to provide people with the tools they need to sort themselves; e.g., the tools of free speech and the right to assembly mentioned in the First Amendment of the Constitution.

Gandy's blind faith in the power of design and sense of impotence in the face of implemented technology is an unlikely mix of an overestimation of the human agency of designers and an underestimation of the agency of non-designers. In contrast, Larry Lessig's analogy between the design of computer code for networked systems and the design of a constitution is simultaneously more hopeful and more realistic than Gandy's cocktail of over- and underestimation of human agency in the face of technology:

---

[107] Michel Foucault, *Discipline and Punish: The Birth of the Prison*, translated by Alan Sheridan (New York: Vintage Books, 1979).

[108] Jeremy Bentham, *The Works of Jeremy Bentham* (London: W. Tait, Simpkin, Marshall, 1843).

[109] Oscar H. Gandy, Jr., *The Panoptic Sort: A Political Economy of Personal Information* (San Francisco, CA: Westview Press, 1993), 227.

The code of cyberspace – whether the Internet, or net within the Internet – the code of cyberspace defines that space. It constitutes that space. And as with any constitution, it builds within itself a set of values, and possibilities, that governs life there ... I've been selling the idea that we should assure that our values get architected into this code. That if this code reflects values, then we should identify the values that come from our tradition – privacy, free speech, anonymity, access – and insist that this code embrace them if it is to embrace values at all. Or more specifically still: I've been arguing that we should look to the structure of our constitutional tradition, and extract from it the values that are constituted by it, and carry these values into the world of the Internet's governance – whether the governance is through code, or the governance is through people.[110]

To follow Lessig's prescription – to design code that embodies the values of free speech – two things are necessary: (1) We need to have a working definition of free speech; and, (2) We need to be able to compare software design practices and software design vocabularies against this working definition of free speech. After all, if software could be implemented with inspiring essays, Lessig would have finished the necessary implementation. Unfortunately, software – as Gandy points out – has a vocabulary and legacy deeply indebted to systems of domination. To translate the language of democracy into the language of software is a difficult goal. But, I will argue, it is not an impossible goal.

Free Speech

The Constitution of the United States presumes, but does not provide, a working definition of free speech. Consequently, the First Amendment illustrates a principle of legislation design that is akin to what, in software design, might more-or-less correspond to object-oriented software design: definition by precedent. In order for a statement like "Congress shall make no law ... abridging the freedom of speech" to make sense, a well-known and respected precedent of freedom of speech must exist. To find this precedent, working definition we must do some history.

According to Arnaldo Momigliano's dictionary entry "Freedom of Speech in Antiquity," freedom of speech had two definitions in ancient Athens, *parrhesia* and *isegoria*:

Isegoria *implied equality of freedom of speech, but did not necessarily imply the right to say everything. On the other hand,* parrhesia *looks like a word invented by a vigorous man for whom democratic life meant freedom from traditional inhibitions of speech. ... The two words* parrhesia *and* isegoria *point to the*

---

[110] Larry Lessig, "Open Code and Open Societies: Values of Internet Governance (Draft 3)," presented as the *1999 Sibley Lecture* (Athens, GA: University of Georgia, February 1999). See http://cyber.law.harvard.edu/works/lessig/kent.pdf

conflict between democracy as liberty and democracy as equality that was to concern later political thinkers.[111]

Contemporary, American practices of voting and polling illustrate how isegoria – equality of speech – can be supported without parrhesia – liberty of speech. On election day, the voters are presented with what is – for all practical purposes – a multiple choice questionnaire: choose Mr. A, Ms. B, or Ms. C for office 1; choose Ms. D, Mr. E, or Mr. F for office 2; etc. While every voter has equal right to speak, voters have an extremely limited liberty to say anything more than "yes" or "no" to each of the preprinted choices on the ballot.

Of course, the minutes in the ballot box do not constitute the entire, American democratic experience. The inclusion of constitutional rights to, for instance, "assemble" and "petition the Government" insure that a more expansive definition of liberty of speech – of parrhesia – can be exercised. Thus, before the ballots are printed, it is possible for the people to organize themselves to put desired candidates, petitions, etc. on the ballot.

Parrhesia is that function of free speech that is demonstrated when the electorate produces a collectively authored statement perhaps in dialog with (or against), but not because of, the government in power. In other words, this function of free speech is what makes a citizenry's self-governance possible. The self-governance facilitated by free speech can be permitted by a constitution and/or by code, but it cannot be guaranteed. The only guarantee is the one demonstrated by an active citizenry willing and able to think and work together to produce collectively authored statements and goals. Phil Agre describes this activity of thinking together as "collective cognition."[112]

As a type of free speaking, parrhesia *was something more than permitted speech; its "freedom" was more than a right to speak. For a state to permit or allow it is not sufficient for it to exist, just as for a state to seek to suppress it is never enough to eliminate its possibility. Thus, we should not restrict our understanding of it to our own modern constitutional or legal view of the right to free speech. Rather we may take it as "ethical" in the sense [the philosopher*

---

[111] Arnaldo Momigliano. "Freedom of Speech in Antiquity" In *Dictionary of the History of Ideas, Volume 2*, edited by Philip P. Wiener (New York: Scribner's Sons, 1973), 260.

[112] *The Internet makes visible a layer of social process that is more fundamental than organizations, and just as fundamental as institutions, namely the customs by which people who have something in common think together. Before collective bargaining comes collective cognition, and collective cognition in its various modes is greatly facilitated by the various community-building mechanisms of the Internet. Ideologies can form in the networked community of computer programmers; news can spread in the networked community of nurses; experiences can be shared in the networked community of cancer patients; patterns can be noticed by the networked community of pilots; agendas can be compared by the networked community of environmental activists; ideas can be exchanged in the networked community of entrepreneurs; stories can be told within the networked community of parents; and so on.* Agre, *Op.Cit.*

Michel] Foucault gave to the term: as a matter of how people recognize themselves as "subjects" of free critical discourse, or of what relations to themselves and to one another they must have in order to speak truly about themselves. In this way we may see in ancient parrhesia *a start of the activity of critical* Wahrsagen *of which Foucault dreamt: the attempt to open an experience of the freedom that conditions our participation in the history that "constitutes" us as individuals, as members of communities, and of states.*[113]

John Rajchman's exposition of parrhesia concisely interrelates two aspects of parrhesia which are sometimes characterized as two different forms of free speech:

(1) parrhesia as political virtue (e.g., you told the prince the truth even if it cost you your head); and,

(2) parrhesia as moral virtue (e.g., you admitted the truth even if it cost you your self image).[114]

This bifurcation of parrhesia is understandable if we remember that the purpose of free speech is to facilitate self-governance. Usually, the term "self-governance" is used to describe a group's control over its own affairs. Thus, "parrhesia as political virtue" is the good of free speech when the "self" (in "self-governance") is understood to be a group of people.

When the "self" is understood to be an individual, then "self-governance" means more-or-less the same thing as "self-control" or "self-discipline." Free speech, in the case of the individual (as compared to the case of a group member) is "parrhesia as moral virtue," i.e., the uncensored expression of the "free association" narrative of a psychotherapy patient or a penitent in the Christian sacrament of confession.

Michel Foucault's history of the practices and transformations of parrhesia details its beginnings in ancient Greece – with Socrates and others – and its subsequent adaptations by ancient Roman society, the early Christians.[115] This history has

---

[113] John Rajchman *Truth and Eros: Foucault, Lacan, and the Question of Ethics* (New York: Routledge, 1991), 119

[114] This distinction and these examples are from Thomas Flynn, "Foucault as Parrhesiast: His Last Course at the Collège de France (1984)" in The Final Foucault, edited by James Bernauer and David Rasmussen (Cambridge, MA: MIT Press, 1987).

[115] See especially, Michel Foucault, *Discourse and truth: the problematization of parrhesia*, Joseph Pearson (editor), Six lectures given at the University of California at Berkeley, Oct-Nov. 1983. The text was compiled from tape-recordings made of the lectures delivered, in English, by Michel Foucault at the University of California at Berkeley in the Fall Term of 1983, transcribed and edited in 1985. The lectures were given as part of Foucault's seminar, entitled "Discourse and Truth". Since Foucault did not write, correct, or edit any part of the text transcripts, they lack his imprimatur and do

been extended to later philosophers, medical practitioners (e.g., Freud[116] and other psychoanalysts), and political thinkers. Foucault's writings show how the practices of free speech have been incorporated into or linked to a variety of technologies and techniques, especially those of contemplation, confession, writing, and conversation. This larger set of techniques and technologies are referred to by Foucault as *technologies of the self*:

My objective for more than twenty-five years has been to sketch out a history of the different ways in our culture that humans develop knowledge about themselves: economics, biology, psychiatry, medicine, and penology. The main point is not to accept this knowledge at face value but to analyze these so-called techniques that human beings use to understand themselves.

As a context, we must understand that there are four major types of these "technologies," each a matrix of practical reason: (1) technologies of production, which permit us to produce, transform, or manipulate things; (2) technologies of sign systems, which permit us to use signs, meanings, symbols, or signification; (3) technologies of power, which determine the conduct of individuals and submit them to certain ends or domination, an objectivizing of the subject; (4) technologies of the self, which permit individuals to effect by their own means, or with the help of others, a certain number of operations on their own bodies and souls, thoughts, conduct, and way of being, so as to transform themselves in order to attain a certain state of happiness, purity, wisdom, perfection, or immorality.[117]

---

not present his own lecture notes. The transcripts are the notes of one of his auditors and are online here: http://perso.club-internet.fr/kmille/discourse.htm

[116] Patrick H. Hutton, "Foucault, Freud, and the Technologies of the Self," in *Technologies of the Self: A Seminar with Michel Foucault*, edited by Luther H. Martin, Huck Gutman and Patrick H. Hutton (Amherst, MA: University of Massachusetts Press, 1988).

[117] Michel Foucault, "Technologies of the Self," in *Ethics: Subjectivity and Truth; Essential Works of Foucault 1954-1984, Volume I,* edited by Paul Rabinow (New York: The New Press, 1997), 224-225.

If we put this list of technologies together with the technologies and techniques discussed previously, we can outline a hierarchical list of technologies and techniques that contextualizes and elaborates the idea of free speech:

technologies of human understanding
1. technologies of production
2. technologies of sign systems
3. technologies of power
4. technologies of the self
    A. free speech
        i. isegoria
        ii. parrhesia
            a. parrhesia as political virtue
            b. parrhesia as moral virtue

At the root of the outline shown above are the technologies, techniques, and arts of human understanding. These are divided into four technologies (of production, of sign systems, of power, and of the self). The technologies of the self are refined into a number of activities. Although here we only mention the technology and techniques of free speech, Foucault also describes other technologies of the self, such as abstinences, memorizations, examinations of conscious, meditations, silence, listening to others, and diary writing.[118] Free speech is, in turn, divided into isegoria and parrhesia. Parrhesia is specialized into two types: parrhesia as political virtue and parrhesia as moral virtue.

What distinguishes one technology from another is a combination of design and use. This assertion is in sharp contrast with Gandy's statement discussed above that a system "can serve no purpose other than that for which it was designed." Gandy's statement is especially concerned with the design of technologies of power to be used with databases of material collected through surveillance and/or observation of particular individuals. Gandy's statement is incorrect because a technology is a combination of design *and* use. For instance, a practice which makes good sense to do for oneself -- e.g., seeing a doctor on a regular basis and keeping detailed records of one's health – can shift from being a technology of the self to becoming a technology of power if a third party – like an insurance company – is allowed to collect and analyze health records. However, a technology of the self can be designed in such a way that makes it more resistant against such a transformation. Thus, it is a good idea to encrypt medical records stored in computer databases and design the database system

[118] Michel Foucault, "On the Genealogy of Ethics," in *Ethics: Subjectivity and Truth; Essential Works of Foucault 1954-1984, Volume I,* edited by Paul Rabinow (New York: The New Press, 1997), 273.

so that any third parties must request the permission of the patient to get the "key" to records.[119]

Functionalities and Resistances

Since, for instance, a technology of the self could possibly be transformed into a technology of power, of production, or of sign systems, it is necessarily the case that two different kinds of design criteria need to be followed to create a robust technology of the self. One criterion concerns the *functionalities* of the system; e.g., a medical records system must be able to store and retrieve records. The second concerns the *resistances* of the system; e.g., a medical records system designed for patients must be hard to use as a technology of power for an insurance company attempting to quantify the risks associated with insuring patients.

A technology's resistances are not exactly the inverse of what many information systems designers now refer to as a technology's "affordances."[120] The so-called affordances are those functionalities of a technology that are easy and obvious to exercise (from the perspective of some sort of idealized, "normal" user).

In contrast, a technology's resistances have nothing to do with "ease of use." Resistances are designed into a technology so that it is difficult or impossible to use the technology for certain activities. Thus, for instance, many designs are "child proofed" (e.g., the back windows of cars only roll down so far, prescription drugs come in containers that are not easy to open, etc.) so that they are "user unfriendly" for children. If a technology is designed with only the criteria of "affordances" and "ease of use" and not with the criteria of functionality and resistances, it is quite likely to be appropriated for purposes that the designer would be unwilling to explicitly design for. The designer who does not care what their work is used for can afford to ignore the practice of designing resistances; the rest of us cannot.

Between each of the four named technologies are six sorts of resistances that keep one sort of technology from becoming another sort of technology. Foucault calls these resistances "encounters":

The encounter between the technologies of domination of others and those of the self I call "governmentality."[121]

---

[119] Cf., D.M. Rind, I.S. Kohane, P. Szolovits, C. Safran, H.C. Chueh and G.O. Barnett, "Maintaining the Confidentiality of Medical Records Shared over the Internet and World Wide Web," *Annals in Internal Medicine* 1997;127(2):138-141.
[120] James Gibson, *The Ecological Approach to Visual Perception* (New York: Houghton Mifflin, 1979).
[121] Foucault, 273.

While Foucault only names one of these resistances, we will examine it (i.e., "governmentality"[122]) and the other two resistances which separate technologies of the self from its siblings (namely, the resistances of automation and simulation).[123] After examining these resistances, some of the functionalities of a technology to support free speech will be discussed. The point of discussing each of these resistances and functionalities is to create a vocabulary that bridges the description of technologies of free speech with computer technologies.

One might wonder how such a discussion will address issues of privacy, censorship, and property rights that are so firmly a part of contemporary, American political discussions about the ethics and morals of the Internet. Self-governance is at the center of the ethics of technology that I am arguing for. What is advocated here are tools to help people (and the people) govern themselves. Privacy, non-censorship, and property rights can be secured through the use of such tools for self-governance, but these are consequences of -- not the causal agency behind – the technologies.

Analogously, in the Constitution of the United States, there is no declared right to privacy. However, such a right is more or less an emergent consequence of the various amendments included in the Bill of Rights which were designed to insure self-governance and preserve the integrity of the self (i.e., both the self as individual and the self as collective). The principle of putting the care of the self at the center of an ethics of design, or an ethics in general, has been the key ingredient for centuries in the West. Normally this ingredient is voiced in the imperative; e.g., "Know yourself" and/or "Take care of yourself." [124]

---

[122] Michel Foucault, "Governmentality," in *The Foucault Effect: Studies in Governmentality*, edited by Graham Burchell, Colin Gordon and Peter Miller (Chicago: University of Chicago Press, 1991), 87-104.

[123] Resistances between other technologies are also interesting, but our focus here is on technologies to support freedom of speech and technologies of the self in general. So these three other pairs will not be discussed. For instance, the ways in which technologies of the sign resist the powers of mastery exercised by an author or speaker is the resistance philosopher Jacques Derrida labels *différance*. See Jacques Derrida, *Positions* (Chicago: University of Chicago Press, 1981), 38-39. Labor politics separate the technologies of power from the technologies of production. Aspects of performance, execution, and materiality separate sign systems from productions. (Idioms expressing this insight include "Put your money where your mouth is.")

[124] *The concern with freedom was an essential and permanent problem for eight full centuries of ancient culture. What we have here is an entire ethics revolving around the care of the self; this is what gives ancient ethics its particular form. I am not saying that ethics is synonymous with the care of the self, but that, in antiquity, ethics as the conscious practice of freedom has revolved around this fundamental imperative: "Take care of yourself."* Michel Foucault, "The Ethics of the Concern for Self as a Practice of Freedom," in *Ethics: Subjectivity and Truth; Essential Works of Foucault 1954-1984, Volume I,* edited by Paul Rabinow (New York: The New Press, 1997), 285.

Automation and Free Speech: Resisting Technologies of Production

How can a vocabulary concerning the self, self-governance, and free speech be translated into a software design vocabulary? In fact, it already has been; it just needs some extensions and a little tweaking.

The Greek root *autos* (used in automation, automatic, etc.) refers to the self. Obviously, automation has a very well-defined meaning in the present, computer age. Usually, when we speak of automation we are referring to a process of industrial production (i.e., to a technology of production, in Foucault's terms).

However, the technical, engineering meaning of the term "automation" has most of the ethical connotations beat out of it. From a technical perspective, automation – if it is successful – implies the implementation of a faster or more efficient machine to replace human labor. The ethical decision still embedded in the engineering notion is the choice to replace or change the job of some person. However, what has often been overlooked – at least since the invention of cybernetics[125] -- is the fact that automation both replaces a human-powered process with a machine while it also allows one group of people to become autonomous from another. For instance, industrial automation is often imposed by management in order to reduce its dependence on labor. But, independence from one set of people (e.g., labor) (re)produces reliance on another group (e.g., other corporations who build and maintain machinery).

Unfortunately, both the engineering and management understandings of "automation" are exclusively concerned with it as a technology of production and say nothing about it as a technology of the self. What is the self, the "auto," in this "automation"? Clearly, it is a job specification that, if it can be automated, can be described in mechanical and/or computational terms. This "auto" of industrial automation is just one manifestation of a much larger movement in the western world to describe the self in computational terms.[126]

Sherry Turkle has examined this larger cultural movement to describe the self in computation terms. Turkle calls the computer the "second self" in her book *The Second Self: Computers and the Human Spirit*.

---

[125] "Cybernetics" is a term which, itself, comes from the ancient Greek word *kubernhthj* meaning steersman or governor. The inventor of the term, Norbert Wiener, was acutely aware of the ethics of automation and worked closely with labor unions to examine the politics of industrial automation. See Norbert Wiener, *Cybernetics, science, and society: Ethics, aesthetics, literary criticism; Collected Works of Norbert Wiener, Volume 4*, edited by P. Masani (Cambridge, MA: MIT Press, 1985).

[126] In academia this movement is called cognitive science. See Howard Gardner, *The Mind's New Science: A History of the Cognitive Revolution* (New York: Basic Books, 1985).

...as computers become commonplace objects in daily life – in leisure and learning as well as in work – everyone will have the opportunity to interact with them in ways where the machine can act as a projection of part of the self, a mirror of the mind.[127]

In Turkle's older work on computers, and newer work on computer networks, the "self" refers to an individual person and not a collective, a group, a corporation, a nation, etc. In Turkle's narratives the computer plays the role of the contemplative object or environment that helps us reflect on our selves, our identities. In her terms, the computer is analogous to a Rorschach inkblot test insofar as they are both powerful, projective media.[128] This conceptualization is helpful for understanding the computer as a technology of the self, but it is limited. Its limitations are revealed by Turkle's interviewing methods. She interviews individuals to determine how the computer or a computer network has changed the individual's "sense of self." She does not interview groups, she interviews individual members of groups. She does not posit a group "self." She does not examine how the links or boundaries between communities of people are changed or erased by the introduction of computer technology. Such a conceptualization of computer as technology of the individual self is not entirely adequate to the task of articulating an ethics of design because ethics is centrally concerned with how the self responds to or changes in the face of the other.[129]

Why is this important? What difference does it make for ethics, free speech and self-governance if we talk about selves, people, populations, groups, etc. using a limiting or a flexible vocabulary? Because, we need a vocabulary for describing -- not only what is, but also -- what could be, what alternatives to the present might be, so that we can decide how to make things better rather than simply accepting them as they are. We need a vocabulary that allows us to be prescriptive, to make proposals, as well as one to be descriptive. Such a vocabulary is a vocabulary of goals, plans, and self-governance. But, more specifically, we need a vocabulary adequate to the task of prescribing and proscribing a range of selves and their interrelationships; i.e., the range from the self as individual self to the collective self of self-governance and self-rule.

---

[127] Sherry Turkle, *The Second Self: Computers and the Human Spirit* (New York: Simon and Schuster, Inc., 1984), 15-16.

[128] *Ibid.*, 14.

[129] *For Levinas ... the ethical relation – and ethics is simply and entirely the event of this relation – is one in which I am related to the face of the Other (le visage d'autrui), where the French word autrui refers to the other human being, whom I cannot evade, comprehend, or kill and before whom I am called to justice, to justify myself.* Simon Critchley, *The Ethics of Deconstruction: Derrida and Levinas* (Cambridge, MA: Blackwell, 1992), 5.

Artificial Intelligence and Autonomous Systems

Vocabularies for the design of self-governing, autonomous systems are, obviously, a mainstay of computer science in general, but of artificial intelligence (AI) in particular. AI inherited this concern from its father field, cybernetics. But, there are also inadequacies in the AI vocabulary of planning and self-governance that need to be addressed if good vocabulary for thinking about the design of technologies of the self, especially technologies for free speech, is to be formulated.

AI's problems are congenital, not inherited. Cybernetics was a field formed in opposition to various types of behaviorism (which emphasized how external stimuli rather than internal goals governed the behavior of individuals and systems). AI was also formed in opposition to behaviorism and distinguished itself as a field from cybernetics by concentrating on a certain approach to autonomous systems that was, at least early in its formation, characterized by a concentration on what are known as *symbol systems*.

During and immediately after World War II the science of cybernetics gained increased prominence. Cybernetics differs from most work done within the confines of a strict behaviorism in at least two ways: (1) Whereas behaviorists postulated linear relationships between an external stimulus and an organism's response, cybernetics introduced the ideas of recursive (i.e., circular) relations between perception or sensation and action known as positive and negative feedback circuits. (2) While behaviorists avoided labeling any behavior "goal-directed" (because it would imply the postulation of internal representations), cyberneticians (re)introduced teleology into scientific descriptions of behavior.[130]

Subsequently, the earliest work in AI elaborated on the cyberneticians' usage of "goal-directed behavior" and de-emphasized external contexts and empirically observable stimuli, the pre-occupation of the behaviorists. Consequently, AI immediately began to diverge from cybernetics due to AI's neglect of an analysis of feedback from the environment. Some contemporary work addresses this early neglect, but early work in AI – e.g., the work of Newell, Simon, and Shaw on the General Problem Solver (GPS)[131] – only explored feedback insofar as the "external world" could be internalized in the computer. To work, GPS required that a full and accurate model of the "state of the world" (i.e., insofar as one can even talk of a "world" of logic or cryptoarthimetic, two of the domains in which GPS solved problems) be encoded and then updated after any action was taken (e.g., after a step was added to the proof of a theorem). This assumption – that perception was always accurate and that all of the significant details of the world could be modeled and followed – was incorporated into most AI programs for

[130] Steve Heims, *The Cybernetics Group* (Cambridge, MA: MIT Press, 1991), 15.

[131] Alan Newell, J.C. Shaw, and Herbert A. Simon, "GPS, A Program That Simulates Human Thought" in *Computers and Thought*, ed. by Edward A. Feigenbaum and Julian Feldman (New York: McGraw-Hill, 1963), 279-293.

decades and resulted in what became known to the AI community as the "frame problem;" i.e., the problem of deciding what parts of the internal model to update when a change is made to the model or the external world.[132] Not surprisingly, AI robots built around strict internal/external divisions sometimes exhibited extremely erratic behavior when the robots' sensors were even slightly inaccurate in the measurement of the external world.

This failure of AI design to produce autonomous robots has in the past decade or so spawned a return to a cybernetics-like interest in the context or environment of the system. Robotic systems are now being designed with sensors and feedback looks wired into their planning machinery. This new research is often called *situated action*[133] or *behavior-based AI*[134] It couples the "neural nets" of connectionism to robotic (hardware and software) bodies with sensors. The research agenda of these groups is, in many ways, a direct descendent of cybernetics insofar as it insists on the employment of feedback circuits and the disruption of internal representation versus external world dichotomies created in and by early AI work.

This new "context sensitive" AI avoids dichotomizing internal representations and the external world by writing systems that compute with a fundamentally different kind of signs than the signs employed in the computer programs of early-AI practitioners. Early AI used symbol systems, contemporary, situated robotics uses, what has been called indexical systems.[135] Symbols are supposed to be signs that are independent of what they represent. For example, the U.S. flag is a symbol for the country, but it is not caused by the country, it is not in some ways intrinsically linked to the country such that it will fall apart if, say, the economy of the country falls apart. On the other hand, indexes are signs that point at something and they change their meaning if the thing they point at changes. Thus, if I point my finger at a woman and say, "That's my friend Jennifer," but then Jennifer moves while I look away to speak, my interlocutor might understand the intent of my spoken phrase, but will be confused by the meaning of my pointing finger if to now points at no one he can see. In robotics, the indexical signs used to construct machines have sensors on them to act as indexes into the current state of the environment and the machine computes on these constantly changing signs.

---

[132] J. Martins, "Belief Revision" in *Encyclopedia of Artificial Intelligence, Second Edition,* editor-in-chief Stuart C. Shapiro (New York: Wiley, 1992), 111.

[133] Philip E. Agre and David Chapman, "Pengi: An Implementation of a Theory of Activity," in *Proceedings of the Fifth National Conference on Artificial Intelligence* (Seattle, WA: Morgan Kaufmann, 1987), 268-272.

[134] Rodney Brooks, "Intelligence Without Representation," *Artificial Intelligence* 47 (1991): 139-160.

[135] Philip E. Agre. "Representation and Indexicality" in *Computation and Human Experience* (New York, NY: Cambridge University Press, 1997); William J. Clancey. *Situated cognition: On human knowledge and computer representations* (New York, NY: Cambridge University Press, 1997).

Ironically, this insight concerning symbol versus indexical systems has been most influential in computational thinking about robotics and other physical (or simulated physical, e.g., virtual reality) systems. This is ironic because the distinction between signs which are symbols and signs which are indexes was first made by the semiotician Charles Sanders Peirce in the early-twentieth century and later elaborated by linguists.[136] According to Pierce, an index is a sign that is directly -- for instance, causally -- linked to its object such as smoke is linked to fire, or as a thermometer is linked to the surrounding temperature. A symbol, by comparison, is a sign linked to its object more indirectly, for instance, through an association of ideas or by habitual connection. The meaning of an index is thus more dependent upon its immediate context than the meaning of a symbol. In other words, AI robotics researchers have begun to build systems with "context sensitive" parts.

However, even though this remarkable development within AI robotics has had great success for producing autonomous robots, it has had little impact on the production of systems to give people more independence, more autonomy, more opportunities for free speech. This, can be explained in terms of what has been the conventional funding structure of AI. The robots are a very powerful technology of production and thus of great use to the funders of AI (the U.S. Department of Defense, large corporations, etc.) most interested in robotics as a tool for military and industrial automation. Consequently, the innovation of indexical systems as an insight into potential technologies of the self as, for instance, technologies for free speech, has not been vigorously pursued. It is the lack of resistance, the lack of will and the lack of knowledge about how to differentiate, technologies of production from technologies of the self that has made this original insight from linguistics and analysts of language and speech into a very useful tool for industry, but has left the discovery underexplored in the area of contemporary, computational technologies of the self.

Indexical Systems and Governmentality

The technology of indexical systems is of potential ethical interest because they provide a vocabulary for describing how semi-autonomous systems (like us and some of our machines) are closely connected to one another and to our environment. Connections and dependencies between ourselves and our linguistic environments also exist. In short, we connect to one another through speech and text. To have the opportunity to do this, to connect and self-organize through the various media of "speech," is the purpose of free speech in a democratic politics.

---

[136] Charles Sanders Peirce, *Collected papers of Charles Sanders Peirce*, edited by Charles Hartshorne and Paul Weiss (Cambridge: Harvard University Press, 1931-1958).

Certain social scientists, especially ethnomethodologists[137] and some linguistic anthropologists[138] have fully explored the idea of indexical systems as a means of description but not as a vocabulary of computational design. In fact, the idea of indexical representations was introduced to AI practitioners by the anthropologist Lucy Suchman in a book she wrote (entitled *Plans and Situated Actions*) critiquing symbolic AI technologies of planning and more-or-less launching the new area that has come to be known as situated-action research in AI.[139] In her book, Suchman points out how the idea of situated actions and indexical signs can be used to rethink the workings of symbolic AI-based technologies.

Following Suchman's lead, it is possible to re-examine old AI experiments in symbol systems using the descriptive powers of indexical systems. There are a number of old technologies that, when examined in this fashion, can be seen to be interesting precursors to a future set of technologies of the self even when they were, in the past, dismissed because, effectively, they were not understood as technologies of the self, but rather expected to be technologies of power or of production.

From the perspective of indexical systems and technologies of the self, one especially interesting early experiment in AI was a project to simulate the interviewing techniques of a Rogerian[140] psychotherapist. As we know from the history of free speech, the sets of techniques and technologies of psychotherapy derive their power from a long history of similar techniques of parrhesia. This history includes Socratic dialogues, Christian rituals of confession, developments in psychoanalysis like the theory of transference, and more contemporary insights of sociology, anthropology, and ethnomethodology (e.g., the documentary method) useful to social scientists for interviewing people and eliciting information. However, ELIZA has rarely been examined in this light even though it is one of the most heavily discussed projects of early AI.

The simulation, a computer program called ELIZA, was designed and implemented by Joseph Weizenbaum in 1966 at MIT.[141] The anthropologist Suchman has shown how the workings of ELIZA are similar to the ethnomethodologist's, Harold Garfinkel's, 1960s experiments with the

---

[137] See Harold Garfinkel, *Studies in Ethnomethodology* (Cambridge, MA: Blackwell Publishers Inc., 1967).

[138] Michael Silverstein, "Shifters, Linguistic Categories, and Cultural Description," in *Meaning in Anthropology*, edited by K. H. Basso and H. A. Selby (Albuquerque: University of New Mexico Press, 1976), 11-56.

[139] Lucy Suchman, *Plans and Situated Actions: The Problem of Human-Machine Communication* (New York: Cambridge University Press, 1987).

[140] Carl Rogers, *On Becoming a Person* (Boston: Houghton Mifflin, 1961).

[141] Joseph Weizenbaum, "ELIZA – A Computer Program for the Study of Natural Language Communication between Man and Machine," *Communications of the Association for Computing Machinery* (9), 1966: 36-45.

documentary method and indexical functions of language.[142] Suchman's use of Garfinkel's vocabulary to explain how ELIZA works can give one a new appreciation for the ELIZA program as procedure of elicitation rather than as a failed symbolic "model of the mind" as most artificial intelligence histories of ELIZA would have us look at it. Here is a snippet of a dialogue with ELIZA. The transcript is taken from Weizenbaum's original 1966 paper. The text in capital letters was generated by the machine, ELIZA.

Men are all alike.
IN WHAT WAY
They're always bugging us about something or other.
CAN YOU THINK OF A SPECIFIC EXAMPLE
Well, my boyfriend made me come here.
YOUR BOYFRIEND MADE YOU COME HERE
He says I'm depressed much of the time.
I AM SORRY TO HEAR YOU ARE DEPRESSED[143]

Clearly, ELIZA is quite effective as a procedure for the elicitation of free speech even though, at the time, it was reviled as an empty, sophistic procedure that did not "really know" how to symbolically compute the answer to a question.

While this shows ELIZA's powers in the elicitation of an individual's voice in a private space, it is instructive to see ELIZA function in a less private, more public space. These public spaces of VLSC where thousands or more can all speak is ultimately the focus of this thesis, but the following anecdote, taken from the same time period, illustrates ELIZA's failure (as a technology of the self) in a smaller, semi-public space (e.g., a corporate network).

Weizenbaum's former colleague at MIT, Daniel Bobrow, told the following story at a conference in 1969. At the time of the story Bobrow was across town from MIT at a company called BBN.

One Saturday morning about 9 a.m., a BBN Vice President in charge of selling our Telcomp commercial service arrived at our PDP-1 Computer Room hoping to be able to use the computer for a demonstration. He found a note on the PDP-1 teletype that read: "TYPE ALL MESSAGES INTO THIS TELETYPE, AND END THEM WITH A PERIOD." This note was left there by one of the programming bums who had just succeeded in making run a LISP version of the DOCTOR program (based on ELIZA as developed by Weizenbaum). However, the VP thought that this note indicated that this teletype was linked with the one in my home, since I often used the computer all night. Therefore, he started typing to "me", but in fact was typing to this version of the DOCTOR program. The following conversation ensued:

---

[142] Suchman, 64.
[143] Weizenbaum, 36.

VP: Do you think that I can use the computer this morning?
D: WHY DO YOU ASK?
VP: I might be able to make some additional Telcomp sales.
D: WHY AREN'T YOU SURE?
VP: My prospective customers have never seen a demonstration of the system.
D: DOES THIS MEAN VERY MUCH TO YOU?
VP: Of course it does.
D: YOU SEEM VERY POSITIVE ABOUT THAT.
VP: This is ridiculous.
D: WHAT DO YOU MEAN THIS IS RIDICULOUS?
VP: Is there any reason why I can't use the computer now?
D: TELL ME MORE ABOUT YOUR PROSPECTIVE CUSTOMERS.
VP: Please dial me up on 491-1850

Note that after the last remark the VP did not obey instructions, and left out the period. Therefore, of course, the computer didn't answer him. This so infuriated the VP, who thought I [Bobrow] was playing games with him, that he called me up, woke me from a deep sleep, and said:

VP: Why are you being so snotty to me?
B: What do you mean why am I being so snotty to you?

The VP angrily read me the dialogue that "we" had been having, and couldn't get any response but laughter from me. It took a while to convince him it really was the computer.[144]

I think Bobrow's story is a brilliant allegory for the issues that must be faced in designing a computer program to be a technology of the self in a non-private space. While Bobrow's story can still be understood as humorous, the 30 years between the time of its first telling and the present time make it a very different story today. The VP in Bobrow's story is obviously upset by what is communicated via computer to him, but he is presumably upset mostly because of the impertinence of the questions posed to him. He is not concerned with what Internet users today must be concerned with when faced with a computer display that requests a lot of information from them (e.g., the questionnaires one must fill out to get information from a variety of corporate websites). Today, we are all concerned with what will happen to the information elicited from us via computer networks. We are, in Oscar Gandy's terminology, worried about the panoptic sort, the possibility that the information that we answer with will later be used against us as a technology of power. This is not at all cause of the discomfort displayed by Bobrow's VP in 1969.

---

[144] Reproduced in Margaret A. Boden, *Artificial intelligence and natural man* (New York: Basic Books, 1977).

In short, one reason ELIZA would fail as a technology of the self today is because it has not been designed to have any resistances to being completely assimilated as a technology of power. In fact, a quick look on the Internet shows how ELIZA-like bots are rather commonly used on corporate sites for providing commercial information to customers and eliciting potentially lucrative, personal information from customers.[145]

Secondly, ELIZA leaves hidden all of the social and communication networks it is participating in. Bobrow's VP has no idea, for instance, that he is not connected to Bobrow's house via teletype. ELIZA fails as a technology of the self because it hides rather than highlights and facilitates the organizational and governing networks that it implements and participates in.

Thirdly, in 1969 Bobrow thinks it is funny that a computer program was interpreted as an impersonation of him. Today there is an increasingly large business creating personalized agents to conduct business transactions for individual and corporate users on the Internet. To imagine that impersonation via computer networks is only a light joke now is to miss the point that impersonation by electronic means can now be a very serious violation of the integrity and security of the self.

Previous critics have found ELIZA lacking because it was not a good symbolic model of the human mind. Others, such as Weizenbaum himself, found ELIZA's functionality morally shocking because it could be understood to be a means of replacing people with machines in what should be intimate and caring exchanges, like those that should take place between doctors and their patients. But neither of these kinds of critique take into account the fact that in certain -- e.g., private -- circumstances ELIZA can be a very effective technology of the self and in other -- e.g., more public -- circumstances it fails to be a good technology of the self because it has not been designed with any thought as to how it might be transformed into a technology of power.[146]

By reexamining the events described in Bobrow's story it is possible rethink what ELIZA did or did not do well in terms of its functionalities as an indexical system. Look again at the sequence of actions that the VP took in Bobrow's story. First, he partakes in a dialogue mediated by the computer, then he demands and initiates a phone conversation; during the phone call he provides evidence to Bobrow by quoting from the printed transcript of his dialogue with the computer; finally, we can imagine that Bobow and the VP had a face-to-face conversation on Monday, at the office, to clear up any remaining misunderstandings. This sequence of media – computer to phone to printed page to face-to-face conversation – recapitulates, in reverse chronological order, their technological

---

[145] See for descriptions of various applications Andrew Leonard, *Bots: the origin of new species* (San Francisco: Hardwired, 1997).

[146] I have made a much more detailed critique of the essentialisms of AI critiques in this essay: Warren Sack, "Artificial Human Nature," *Design Issues* (Summer 1997).

development. Each previous medium is thought to guarantee or provide the basis for the following one. So, the VP repeats this sequence of media in the hopes that the older media will provide surer evidence of his position when the younger media fail to do so. It is a flight from the "artificial" to what is considered to be the "real."

Thirty years ago, when computers were relatively new, not much information was kept on-line. The chance that a given computer might contain an identifier for a given person – say, a driving record or a social security number – was small. Consequently, there existed few if any technological or cultural *interpretants*[147] holding the signs or identifiers within the computer together with those outside of the computer. In other words, the computers' identifiers were in an almost arbitrary relation to the rest of the world and *symbol systems* was, therefore, a good name for them. In fact, this was the preferred appellation by some of the most prominent AI researchers of the time.[148]

However, today computer systems are more indexical than symbolic because of cultural and technological developments of the past few years (e.g., the widespread use of email addresses; and, networking and cryptography). The computer representations of today are understood to be signs that are securely attached to the "real." Thus, it is unusual for one to feel a need to guarantee a computer connection with a connection provided by an older medium. For example, if I receive an email message from you, it is unlikely that I will phone you to make sure that it was really you that sent me the email message; I normally believe that the return address on the email (the 'From:' field) is a 'reliable' sign. That computer 'symbols' could later be accepted as 'indices' is hardly remarkable. Every index, even the fingerprint, has a history of development previous to its acceptance as a sign that directly points back to a physical entity or body. Now is a crucial point in the history of computers when computational representations are increasingly accepted as indexes rather than simply symbols.

Conclusions: A Network Design Ethics for Free Speech

The good of new computer network-based technologies for conversation is that they can potentially provide us with new means for free speech and therefore new powers of self-governance. If these new forms of conversation work effectively as new technologies of self they should make it possible for us to better understand how and where we are located in a variety of social and linguistic networks. Furthermore, by making these networks visible they should give us the tools to plan how we want our interconnections to be between us, other people, and non-human entities, like the machines we depend upon everyday. I am arguing for a means to visualize the social and semantic

---

[147] Peirce.

[148] Herbert A. Simon, *The Sciences of the Artificial, Third Edition* (Cambridge, MA: MIT Press, 1996).

networks that are now hard to perceive. By making them visible for everyone to see, it is hoped that it will be possible for them to be used as resources for public discussion. Who is excluded from the linked social networks of public view? Who is central? Why? What are the predominant associations connected with the signs of public life and public discourse? Who and what are accorded positive even glorious associations? Who are given more modest or even negative associations? Why? Why not?

Simple technologies, like the original ELIZA system, can show us the possible powerful and good effects on free speech provided by a mirroring or reflection of people's words. By reflecting their words back to them, people are given a means to reflect themselves and thereby find a voice for themselves. But a simple mirror like ELIZA can easily fall apart as a technology of the self in a networked, highly populated environment. In designing these new technologies of the self it is necessary to remember that every technology of the self can function only if it can resist becoming a technology of power, a technology of production, or a technology of the sign (a possibility that will be more closely examined in the following chapter on aesthetics). If it is not designed with a set of resistances as well as with a set of functionalities, then it can easily cease to work as a technology of the self and, instead, becomes effectively one of these other technologies. As new database and network technologies are put into place to describe, classify, and sort people, a new set of resistances must be designed into more sophisticated, computational technologies of the self necessary for free speech and self-governance.

Newer work in AI has found indexical representations to be a powerful tool for the design and implementation of physically-based, autonomous systems. Using the insights gathered by anthropologists and linguists concerning the function and use of indexes in language-based environments it should be possible to create an analogous design vocabulary for the production and critique of autonomous systems that can support free speech and, thereby, self-governance by participants in very large-scale conversations. By recognizing how quotation, citation, word repetition, and other means of threading words and messages together can connect us into groups with "selves" constituted by thousands or millions of people, a new technology of self might be implemented that resists the new technologies of power and production of market surveys, online demographics, and electronic tagging.

AESTHETICS

Previous software design approaches (especially those of artificial intelligence) are closely tied to a commonsense aesthetics, i.e., an aesthetics that presumes a commonsense, a predictably emergent commonsense, or the uncanny, interference of the commonsense world. An alternative to these approaches must be found if we are to design for VLSCs where a potential, or virtual, commonsense is contingent upon the possible (but not necessarily probably) emergence of a community of people who create their own stable semantic and social structure through continued interaction on the Internet. This new aesthetics, therefore, must be useful for the practices of design for emergent communities.[149]

[149] Parts of this chapter appear in Warren Sack, "Artificial Intelligence and Aesthetics," in *The Encyclopedia of Aesthetics, Volume 1*, editor-in-chief Michael Kelly, editor-in-chief (New York: Oxford University Press, 1998).

Introduction: User-Friendly, Commonsensical Interface Design

In order for an interface to work, the person has to have some idea about what the computer expects and can handle, and the computer has to incorporate some information about what the person's goals and behaviors are likely to be. These two phenomena – a person's "mental model" of the computer and the computers "understanding" of the person – are just as much a part of the interface as its physical and sensory manifestations. ... Faced with this nightmare, our seminar at Atari abandoned the topic and we turned our attention to more manageable concepts, such as the value of multisensory representations in the interface.[150]

Brenda Laurel unearths a potentially mountainous obstacle for interface designers. Most interface designers want to create something that is "user friendly," i.e., easy to use. Some of these designers have taken the approach of graphically-sophisticated, direct manipulation interfaces that are intuitive to use.[151] In contrast, artificial intelligence (AI) researchers often insist that the interface *per se* is not that important for the goal of "user friendly" software. If the computer's "understanding" of the person is a deep and profound understanding, then the computer can anticipate or sensitively perceive what a given person wants and fulfill those wants with minimal interaction with the user. This has been called the "intelligent agents" approach to interface design.[152]

Note, however, that both the agents and the graphical interface approaches require some notion of what might be called *commonsense*, or commonsense knowledge. The AI researchers assume that the commonsense can be coded into a computer program so that the computer can "know" what the person knows. The graphical interface designer assumes that an intuitive interface is one that does not require a user to read a thick manual before the user can use it. In other words, the interface should be designed so that the user – does not have to rely on some specialized knowledge but, rather -- can rely on their own commonsense to use the interface.

Many AI researchers have believed that this commonsense can be coded as a computer program. Graphical interface designers do not necessary think that the commonsense can be coded, but they must at least rely on their own intuitions about what is commonsensical in order to determine if an interface design is in practice easy to use without specialized, non-commonsense, knowledge. But, what is commonsense? Marvin Minsky, one of the founders of AI said the following in a recent interview:

---

[150] Brenda Laurel, *Computers as Theater* (Reading, MA: Addison-Wesley Publishing Company, 1991), 12-14.
[151] Ben Schneiderman, *Designing the User Interface: Strategies for Effective Human-Computer Interaction* (Reading, MA: Addison-Wesley, 1987).
[152] Nicholas Negroponte, *Being Digital* (New York: Knopf, 1995).

Q. How do you define common sense?
A. Common sense is knowing maybe 30 or 50 million things about the world and having them represented so that when something happens, you can make analogies with others. If you have common sense, you don't classify the things literally; you store them by what they are useful for or what they remind us of. For instance, I can see that suitcase (over there in a corner) as something to stand on to change a light bulb as opposed to something to carry things in.[153]

Minsky's definition of commonsense can be discussed using a linguistic terminology. Given a term like "suitcase" it should be possible to associate it with actions like "carry" and "stand." I.e., those who possess commonsense should be able to employ "suitcase" as the indirect object of the verb "stand" and "carry." However, expressed in this terminology, it becomes clear that there are set of cultural dependencies implicit in Minsky's definition. What parts of commonsense are missing in the knowledge of a non-English speaker who doesn't know the word "suitcase"? Probably nothing is missing for speakers of a language that have some equivalent to "suitcase" (e.g., "une valise" in French). But, more importantly, what is different, missing, or added for those whose language or culture contains nothing like a suitcase?

Some have suggested that it might be possible to divide up commonsense into two kinds: a culturally dependent commonsense knowledge and a culturally-independent sort of knowledge:

...I have suggested that people analyze the world, sort it into categories, impose structure on it, in order to avoid being overwhelmed by its richness. I have implied that this procedure is not deliberate: the widely held notion of "common sense" suggests that people believe that their theory of the way the world works is a natural reflection of the way the world does work. If we look at the sources of categories, we find that some are natural in origin, but the majority are social. Research suggests that a number of basic "cognitive categories" do arise in individuals naturally, being a product of the way we are constructed biologically. These include basic colour categories, such as black and white, red and green; certain geometrical figures, such as circle, triangle, and rectangle; notions of movement, such as up, down, forward, backward; logical relationships, such as oppositeness, identity, and causation. But the majority of our ideas are not natural. ... What counts as an instance of a category is subject to negotiation and revision. Can a lion count as a pet? Yes, the magistrates say, provided it is locked up securely. Clearly the idea of "pet" cannot be derived from any list of actual animals; it is not a natural feature of certain animals but a property of the culture's system of attitudes towards animals.[154]

---

[153] Claudia Dreifus, "A Conversation with Dr. Marvin Minsky: Why Isn't Artificial Intelligence More Like the Real Thing?," *New York Times*, July 28, 1998; or see http://www.nytimes.com/library/national/science/072898sci-minsky.html

[154] Roger Fowler, *Linguistic Criticism, Second Edition* (New York: Oxford University Press, 1996), 27.

Such a culturally-dependent/culturally-independent division of commonsense – like the one offered by Roger Fowler in the quote above – might be a workable means for interface and/or AI designers to approach their work with. However, such an approach would still require different designs for different cultures if the software was suppose to operate in a domain that did not occupy a "basic" category of knowledge. Conversation, for instance, is a culturally dependent domain if only because topics of conversation are rarely if ever entirely culturally independent. Very large-scale conversation is an even more eclectic domain because, as it is presently practiced on the Internet, participants can come from a wide diversity of cultural backgrounds and so what is or is not commonsensical cannot be enumerated beforehand.

Instead, what is necessary is a design perspective that allows one to see how, for instance over the course of a long-term conversation, commonsense is produced, reproduced, extended, and changed by a group of – potentially culturally diverse – participants. The political philosopher Antonio Gramsci gives us just such a picture of commonsense:

Every social stratum has its own "common sense" and its own "good sense," which are basically the most widespread conception of life and of men. Every philosophical current leaves behind a sedimentation of "common sense": this is the document of its historical effectiveness. Common sense is not something rigid and immobile, but is continually transforming itself, enriching itself with scientific ideas and with philosophical opinions which have entered ordinary life... Common sense creates the folklore of the future, that is as a relatively rigid phase of popular knowledge at a given place and time.[155]

From this perspective, commonsense is accumulated and transformed through the process and productions of science, philosophy and other powerful conversations, discourses, and practices. This is a perspective that has been useful for understanding the workings of older media (e.g., newspapers, television, film, etc.) and could, potentially, be of use to understand and design new forms of media like VLSCs.[156]

---

[155] Antonio Gramsci. *Selections from the Prison Notebooks* (London: Lawrence and Wishart, 1971), 326; as cited in Stuart Hall. "The rediscovery of 'ideology': return of the repressed in media studies," in *Culture, Society, and the Media*, edited by Michael Gurevitch, Tony Bennett, James Curran, and Janet Woollacott (New York: Routledge, 1982), 73.

[156] According to Stuart Hall, Anglo-American media studies of the early-twentieth century saw the media (newspaper, television, etc.) as producers of content that "reflected" the "common sense" of the larger public. The media was said to objectively write down and distribute the consensus, or *sensus communus*, that was produced by the public independent of the media. Hall argues that, later, media studies came to recognize the media's role in producing, rather than simply reflecting community values and common sense. By being the only "voice" which could reach across the nation and even across the world, the electronic and print media framed public discourse, and thus public

However, this is probably easier said than done. Not just interface designers, but many other kinds of artists and designers have consciously or unconsciously relied on some notion of "culturally-independent" commonsense to make aesthetic decisions. To ferret out this dependency in software design and find a workable alternative for thinking about the aesthetics of VLSC interface design, this chapter will first explore how commonsense has been discussed and used in software, specifically artificial intelligence design. It is shown, historically, the connections between aesthetic decisions and terms central to AI work – especially goals and commonsense – are longstanding concerns. It is thus necessary to get some historical and philosophical perspective on discussions of commonsense and aesthetics in order to propose true alternatives. The main goal for this chapter is the formulation of an approach to the design of an interface for VLSC that can show the production of commonsense (especially the commonsense of conversation) and who is responsible for its production. Rather than depending upon an *a priori* defined notion of commonsense, a workable approach to the aesthetics for VLSC interface design must take into account the fact that commonsense is being produced and changed through the conversation itself. After looking at the history of AI, commonsense, and aesthetics, an alternative approach is outlined.

Artificial Intelligence and Aesthetics

Artificial intelligence (AI) is an area of research and design of "intelligent" computer hardware and software. The term "artificial intelligence" was coined for a conference at Dartmouth College held in the summer of 1956.[157] The Dartmouth conference brought together the majority of researchers who are today considered the founders of the field including John McCarthy, Marvin Minsky, Herbert Simon, Allen Newell, and others. While AI has primarily been a concern of computer scientists, its multidisciplinary membership (including also mathematicians, philosophers, engineers, and social scientists) was evident even at the time of the Dartmouth conference. AI did not have a name before the Dartmouth conference yet it, nevertheless, participates in older intellectual and design traditions which have investigated mechanical and symbolic systems and human cognition and perception for centuries. Consequently, as an area of design concerned with cognition and perception, AI can be understood as the latest manifestation of certain views of aesthetics which have their roots in older philosophical, scientific, and artistic projects.

The purpose of the following chapter sections is to give a short history of AI that highlights its relations with a Kantian (Immanuel Kant) view of aesthetics. Its

---

"common sense," simply through the editorial choice of which stories should be broadcast and which should be left untold. *Ibid.*

[157] Howard Gardner, *The Mind's New Science: A History of the Cognitive Revolution* (New York: Basic Books, 1985).

purpose is *not* to give a comprehensive overview of AI (see, instead, Shapiro[158] for one such overview). Instead, this chapter's focus is the intersection of AI and aesthetics and so it supplements, but does not largely overlap, two different histories that have been repeatedly told about (1) AI and science; and, (2) AI and art. A history of AI concerned with its scientific roots would emphasis its relations to the development of calculating machines, logic, and mathematics.[159] An art history of AI would, by contrast, detail its similarities and differences with ancient and modern myths, literatures, and depictions of robots, cyborgs, and artificially (re)created humans like Frankenstein's monster.[160] For expository purposes, these other histories (of AI, art, and science) are mostly left to the side so that a single, streamlined story, focusing on AI and aesthetics, can be told. At the end of these sections, the "streamlining" is questioned by examining some of AI's relationships to other (i.e., non-Kantian) aesthetics. This "unstreamlining" makes it possible to propose a set of alternatives to a commonsense-based aesthetics to interface design.

Early AI

Throughout its -- now over forty year -- history AI has never been a discipline without internal differences. Nevertheless, until about the mid-nineteen eighties it was possible to say that a large majority of AI researchers were concerned with the elaboration of a *rationalistic* understanding of cognition and perception.[161] Within the rationalistic tradition, human identity and the thinking, calculating mind tend to become conflated. AI's rationalistic bent can be understood by examining it as a reaction against behaviorism,[162] the approach that dominated the social sciences for most the first half of the twentieth century in the United States, and an outgrowth of cybernetics,[163] an interdisciplinary effort born during World War II to study social, biological, and electro-mechanical systems as systems of control and information.

Behaviorism and AI

Behaviorists' preference for studying external, empirically observable behaviors rather than, for example, a method of introspection or the analysis of verbal reports of others' thinking, effectively divided psychology (and other social sciences) from closely-related disciplines, like psychoanalysis, which were

---

[158] Stuart C. Shapiro (editor-in-chief) *Encyclopedia of Artificial Intelligence, Second Edition* (New York: Wiley, 1992).
[159] E.g., Gardner.
[160] Cf., Pamela McCorduck, *Machines who think: a personal inquiry into the history and prospects of artificial intelligence* (San Francisco: W.H. Freeman, 1979), 3-29.
[161] Terry Winograd and Fernando Flores, *Understanding Computers and Cognition: A New Foundation for Design* (Norwood, NJ: Ablex Pub. Corp., 1986).
[162] B.F. Skinner, *Science and human behavior* (New York: Macmillan, 1953).
[163]Norbert Wiener, *Cybernetics; or, Control and Communication in the Animal and the Machine* (New York: Wiley, 1955).

founded on the postulation of internal, mental structures and events. As computers became more and more common, the behaviorists' hegemonic position in American social science began to wane. Behaviorists were unwilling to postulate the existence of intentions, purposes and complicated internal, mental mechanisms. Yet, during and after World War II, as computers were built to do more and more complicated tasks, not only computer engineers, but also the popular press began to call computers "electronic brains" and their internal parts and functions were given anthropomorphic names (e.g., computer "memory" as opposed to, for instance, the synonymous term the "store" of the computer). Concomitantly, some social scientists began to take seriously the analogy between the workings of a computer and the workings of the human mind.[164] This set of social scientists went on to found AI and cognitive science as a whole, the area of science that includes AI and a variety of other "computationally-inspired" approaches to cognition in linguistics, anthropology, psychology and neurophysiology.[165]

Cybernetics and AI

At the same time – i.e., during and immediately after World War II – the science of cybernetics gained increased prominence. Cybernetics differs from most work done within the confines of a strict behaviorism in at least two ways: (1) Whereas behaviorists postulated linear relationships between an external stimulus and an organism's response, cybernetics introduced the ideas of recursive (i.e., circular) relations between perception or sensation and action known as positive and negative feedback circuits. (2) While behaviorists avoided labeling any behavior "goal-directed" (because it would imply the postulation of internal representations), cyberneticians (re)introduced teleology into scientific descriptions of behavior.[166]

Subsequently, the earliest work in AI elaborated on the cyberneticians' usage of "goal-directed behavior" and de-emphasized external contexts and empirically observable stimuli, the pre-occupation of the behaviorists. Consequently, AI immediately began to diverge from cybernetics due to AI's neglect of an analysis of feedback from the environment. Some contemporary work addresses this early neglect, but early work in AI – e.g., the work of Newell, Simon, and Shaw[167] on the General Problem Solver (GPS) – only explored feedback insofar as the "external world" could be internalized in the computer. To work, GPS required that a full and accurate model of the "state of the world" (i.e., insofar as one can even talk of a "world" of logic or cryptoarthimetic, two of the domains in which

---

[164] Sherry Turkle, "Artificial Intelligence and Psychoanalysis: A New Alliance." *Daedalus*, 17(1) Winter 1988.
[165] Gardner.
[166] Steve Heims, *The Cybernetics Group* (Cambridge, MA,: MIT Press, 1991), 15.
[167] Alan Newell, J.C. Shaw, and Herbert A. Simon, "GPS, A Program That Simulates Human Thought" in *Computers and Thought*, ed. by Edward A. Feigenbaum and Julian Feldman (New York: McGraw-Hill, 1963), 279-293.

GPS solved problems) be encoded and then updated after any action was taken (e.g., after a step was added to the proof of a theorem). This assumption – that perception was always accurate and that all of the significant details of the world could be modeled and followed – was incorporated into most AI programs for decades and resulted in what became known to the AI community as the "frame problem;" i.e., the problem of deciding what parts of the internal model to update when a change is made to the model or the external world.[168] Not surprisingly, AI robots built around strict internal/external divisions sometimes exhibited extremely erratic behavior when the robots' sensors were even slightly inaccurate in the measurement of the external world.

AI as a Kantian Endeavor

Early AI's anti-behaviorist, inward turn to focus on internal representations (like "goals") led to, what can be understood as, the re-invention of philosophical rationalism's problems and "solutions." Or, since AI has routinely been concerned with the difficulties and sometimes the limits of rationality (expressed, for example, in Herbert Simon's notion of "bounded rationality"[169]), its "re-inventions" more specifically resemble, not rationalism *per se*, but philosophical responses to rationalism like, for example, Immanuel Kant's *Critiques*. Indeed, the rationalistic approach to perception and cognition pursued by a large majority of AI researchers until the mid-nineteen eighties can be explained in Kantian terms.

The following explanation of AI using Kantian terms relies on a well-known reading of Kant formulated by the philosopher Gilles Deleuze.[170] Deleuze's interpretation is akin to several other readings of Kant (notably, the work of Jean-François Lyotard[171]) that, collectively, might be viewed as the "post-structuralist" response to Kant.

Kant's comparison of aesthetical and teleological judgement[172] provides a framework for narrating how AI's original pre-occupations with teleology and neglect of aesthetics caused a series of crises for the field in the mid-nineteen eighties that initiated an "aesthetic turn"[173] in research motivating AI scientists

---

[168] J. Martins, "Belief Revision" in *Encyclopedia of Artificial Intelligence, Second Edition*, editor-in-chief Stuart C. Shapiro (New York: Wiley, 1992), 111.
[169] Herbert A. Simon, *The Sciences of the Artificial, Third Edition* (Cambridge, MA: MIT Press, 1996).
[170] Gilles Deleuze, *Kant's Critical Philosophy: The Doctrine of the Faculties*, translated by Hugh Tomlinson and Barbara Habberjam (Minneapolis, MN: University of Minnesota Press, 1984).
[171] Jean-François Lyotard, *Lessons on the analytic of the sublime: Kant's Critique of Judgment*, translated by Elizabeth Rottenberg (Stanford, CA: Stanford University Press, 1994).
[172] Immanuel Kant, *Kritik der Urtheilskraft, in Immanuel Kant's gesammelte Schriften, Volume V* (Berlin: Bey Lagarde und Friederich ,1790), 165-485.
[173] Warren Sack, "Artificial Human Nature" *Design Issues*, 13 (Summer 1997), 55-64.

and designers to pay increasing attention to issues of the body, the senses, and physical and social environments. While Kant's vocabulary provides a convenient means of describing the problems and achievements of AI, within the literature of AI Kant is rarely mentioned, or, if mentioned, then only represented as formative of AI's parent discipline, cognitive science.[174] Here it is argued that Kant's vocabulary of aesthetics (as "spoken" by post-structuralism) is equal to the task of describing many important research issues in AI. No argument is offered to support the opinion that some sort of "equivalence" exists between AI and Kant's critical project.

In conflict with rationalists (like René Descartes) Kant argued for a limited role for teleological principles to supplement mechanical explanations.[175] Likewise, cyberneticians – in conflict with most behaviorists – argued for a limited use of teleology in coordination with the vocabulary of physics to describe the behavior of complex, non-linear systems. In the nineteen-fifties, when AI took up the vocabulary of teleology (i.e., "final cause") from cyberneticians, what was repressed – or at least foreclosed – for AI was the problematic status of the posited goals and purposes used to understand the behavior of complex systems. For Kant, teleological principles were considered to have no explanatory significance.[176] Instead, teleological judgement – in Kantian terms – was seen as a response to an apparent purposelessness of aesthetic judgement: "purpose" is a projection of the cognitive subject on nature, not an intrinsic property of nature itself. In contrast to Kant's careful re-introduction of teleology within a nuanced discussion of aesthetics – where teleology was seen as a *product of cognition*, but not necessarily an artifact useful for communication and explanation – AI researchers took teleology as a *basis for their scientific explanations of cognition*, problem solving, and learning. This difference of opinion concerning the explanatory significance of goals and purposes is so large that one might assume that any continued narration of AI as a type of "Kantianism" (as AI's history is here described) would be fruitless. However, the way in which AI has struggled with the questions of teleology (e.g., For whom do posited goals and purposes signify something meaningful?) is strikingly Kantian in its (AI's) recurrent appeal to "common sense", a faculty of great importance to Kant's critical philosophy.

Kant and Common Sense

"Faculty" (e.g., a "faculty of common sense") is a crucial yet ambiguous term in Kant's writings. For present purposes it suffices to say that a Kantian "faculty" is a potential or power to realize some end.[177] Computational approaches to philosophy, like those championed by AI, often equate "powers" to

---

[174] Gardner; see also, Andrew Brook, *Kant and the Mind* (Cambridge, UK: Cambridge University Press, 1994).
[175] Howard Caygill, *A Kant dictionary* (Cambridge, MA, Blackwell Reference, 1995), 388.
[176] Kant, para. 61.
[177] Caygill, 190.

"computational processes" or (to use a less theoretical term) "computer programs." Thus, to draw an analogy between the writings of Kant and the writings of AI researchers, it is necessary to imagine that Kant's "faculties" could be re-expressed in a variant, textual form: as computer programs with specific data structures, data flow, and control flow. Although this metaphorical comparison of the human cognitive faculties to computer programs may seem outlandish to many, it is a hallmark, not only of AI, but of all contemporary cognitive science.[178] To compare Kant's work to the research goals of AI it is not necessary to believe that this metaphor (mind as machine) is "true." Rather, it is only necessary for the reader to be able to imagine that AI researchers consider this metaphor to be true; or, if not true, then at least extremely useful.

Kant discusses three cognitive faculties: understanding, reason, and imagination. In the terminology of AI one might explain these faculties as classification (understanding), inference (reason), and schema or pattern matching (imagination). In addition to the three cognitive faculties, Kant describes three sorts of common sense: two legislated and one engendered.[179] The ways in which the three faculties described by Kant (of understanding, reason, and imagination) interrelate with one another are referred to as (1) *logical common sense*; and, (2) *moral common sense* when, respectively, (1) understanding, and (2) reason legislate over the two other complementary faculties. In contrast with these two legislated sorts of common sense, (3) *aesthetic common sense* is engendered when none of the faculties are regent, but when they all, nevertheless, work together even as they function autonomously and spontaneously. In the vocabulary of contemporary computer science, one might say that the differences between these three kinds of common sense are differences in "control structure"; i.e., differences concerning which (or whether) one computer program, program statement, or "faculty" is directing the others.

Kant's theories of *reflective judgement* (which includes both *aesthetic reflective* and *teleological reflective judgement*) function with the support of an engendered, aesthetic common sense. This common sense is engendered when, for example, the faculty of reason compels the faculty of imagination to confront its limits by attempting to schematize a perception of the formless or the deformed in nature (a state referred to as the *sublime*). According to Deleuze, the aesthetic common sense should not be understood as a supplement to logical and moral common sense but that which gives them a basis or makes them possible since the faculties of understanding and reason could not take on a legislative role if it were not first the case (as in the accord of an aesthetic common sense) that they are each capable of operating in free subjective harmony.[180] The implications for AI of this conceptualization of common sense – like Kant's aesthetic common sense – will play an important role in the following discussion.

---

[178] Gardner, 6.
[179] See Deleuze, 49-50.
[180] Deleuze, 50.

AI and Common Sense

The neo-Kantian, Jean-François Lyotard draws an analogy between Kant's description of *reflective judgement* (a mode of thought which works from the particular towards the universal as opposed to *determinant judgement* which proceeds in the inverse direction) and AI researchers' (especially Marvin Minsky's) descriptions of "frame-based" thought and perception.[181] Minsky's "frames" proposal[182] was an attempt to describe common sense thought in humans and its possibilities in computers. Minsky, McCarthy,[183] their students, and colleagues in AI were concerned with the following question about common sense: What is the structure and content of common sense such that it allows one to quickly draw useful conclusions from a vast array of existing knowledge and perceptional data? One of the immediate outgrowths of this research was a series of "frame-based" computer programming languages with control structure statements very unlike previous programming languages.[184] From this, AI, point of view, common sense is a legislative faculty, i.e., a control (or controlling) structure that allows a system to turn its attention away from nonsense so that it can concentrate on the sensical or what is implied by the commonsensical. In other words, in Kantian terms, AI's analysis of "common sense" was largely (and still is in some circles) limited to "legislated common sense" -- "logical common sense" and "moral (i.e., 'reason-legislated') common sense" -- and had (until recently) completely neglected "aesthetic common sense," an unlegislated state of relations between understanding, reason, and imagination.

AI and Non-Military Concerns

Such was the case until the mid-nineteen eighties when two "events" motivated a reappraisal within AI of the issues of aesthetics, including the role of the body and the senses in cognition. The first "event" was the commercial success of a genre of AI computer programs called "expert systems."[185] For the first thirty years of its existence (in the United States) AI was mostly academic research funded by the military. Then, in the mid-eighties, business concerns began funding the development of expert systems to automate a variety of white-collar

---

[181] Jean-François Lyotard, *The Inhuman: Reflections on Time*, translated by Geoffrey Bennington and Rachel Bowlby (Stanford, CA: Stanford University Press, 1991), 15.

[182] Marvin Minsky, "A Framework for Representing Knowledge" in *Mind Design: Philosophy, Psychology, Artificial Intelligence*, ed. by John Haugeland (Cambridge, MA: MIT Press, 1981).

[183] John McCarthy, *Formalizing Common Sense: Papers by John McCarthy*, edited by Vladimir Lifschitz and John McCarthy (Oxford, UK: Intellect, 1998).

[184] E.g., R. Bruce Roberts and Ira P. Goldstein, "The FRL manual," (Cambridge, MA: Massachusetts Institute of Technology, 1977).

[185] Edward A. Feigenbaum and Pamela McCorduck, *The fifth generation: artificial intelligence and Japan's computer challenge to the world* (Reading, MA: Addison-Wesley, 1983).

work. While the U.S. Department of Defense had been content to finance long-term research in which it was presumed that theoretical work might, one day, be of practical interest, the new benefactors of AI demanded empirically evaluated, immediate results. What soon became clear was that many expert systems were "brittle," i.e., they performed competently within a narrow domain of problems, but if the problems were posed in a slightly different manner, or if slightly different types of problems were posed to the systems, the systems responded in erratic and erroneous ways. Moreover, it was noted by users of the systems that the systems were difficult to communicate with: one needed to pose problems in a specially constructed, artificial language and, often, after receiving a solution from a system it was impossible to get the system to explain the rationale for its solution. Expert system adherents claimed the problem was simply that more rules needed to be added to the "brittle" expert systems to make them "flexible." Expert system opponents, often using a philosophical vocabulary of (Martin Heidegger's) phenomenology, claimed that rules were inadequate to the task of articulating the means of human expertise and, thus, no number of rules could allow a machine to match the skills of a human expert.

The second "event" was the U.S. military's loss of funding in the late-eighties due to the end of the Cold War with the Soviet Union. Both "events" pushed AI researchers to look for new funding sources and "applications" in finance, advertising, and entertainment.

Two Strands of Aesthetic AI Research

This exodus from the isolated, military-industrial funded laboratories fostered two strands of research. One strand is attributable to a diverse collection of researchers who, for the purposes of this chapter, will be called the "neo-Encyclopediaists." The second strand of researchers will here be called the "computational phenomenologists." Both of these strands have longer histories of existence, even within the lifetime of AI itself, but they were given more funding and attention after the two, above-mentioned "events." One of the major distinctions between these two strands of researchers is this: while the neo-Encyclopediaists (or at least their predecessors in symbolic AI; e.g., Minsky) feel that "common sense" can be cataloged as a system of rules with intentional content,[186] the computational phenomenologists do not believe that "common sense" can be articulated in the structure of rules.

The "rules" under scrutiny by the computational phenomenologists can be understood as a certain specialized form of computer program articulated as a series of IF-THEN statements (e.g., "IF the water is steaming and bubbling, THEN its temperature is probably 100 degrees Celsius"). But, the term "rules" can also be understood as synecdochically referring to a larger class of computer

[186] Cf., Hubert L. Dreyfus with Harrison Hall (eds.) *Husserl, intentionality, and cognitive science* (Cambridge, MA: MIT Press, 1982), 23.

programs (including, for example, Minsky's "frames", and what others have called "schemata," or "scripts"[187]).

The neo-Encyclopediaists

Motivated by the observation that most AI programs do not contain enough schemata or rules to deal with unforeseen circumstances (c.f., the "brittleness" of expert systems mentioned above), the neo-Encyclopediaists[188] are attempting to produce huge catalogs of "common sense" (i.e., computer programs and databases). Some of these efforts are largely accomplished "by hand" whereby dozens of people are employed for years to encode a myriad of mundane details and rules (e.g., "water is a liquid", "what goes up must come down"). Other efforts are aided by statistical and machine learning techniques to augment or build such catalogs. The most well known of these efforts has been a ten-year project called CYC (originally short for enCYClopedia) financed largely by corporate sponsors.[189] CYC and a handful of other efforts are the contemporary offspring of Minsky's[190] and McCarthy's[191] proposals for representing common sense, oftentimes referred to as "symbolic AI."

However, while work in symbolic AI has always stressed the importance of teleological and "intentional" representation, newer work in "computational linguistics" (a field that intersects with the AI sub-field of "natural language processing") contributes to the neo-Encyclopediaists' efforts without necessarily ascribing the same importance to teleology. Computational linguistic, neo-Encyclopediaist work is often described as the latest extension to the long-standing field of lexicography, the discipline that has historically been responsible for the construction of encyclopedias, dictionaries, and thesauri.[192] This turn away from teleology in recent neo-Encyclopediaist work might be seen as a renewed interest in the freedom of the (Kantian) imagination and its power to schematize without any concept,[193] i.e., an interest in the basis for an aesthetic common sense (taste). One difference, however, is, for instance, the dependence of much recent computational linguistic work on the form of very simple "schemata" or "rules" (e.g., the form and limitations of Markov models) versus the postulation of no schematic influence whatsoever by Kant.

---

[187] Roger C. Schank and Robert P. Abelson, *Scripts, plans, goals and understanding: An inquiry into human knowledge structures* (New York: John Wiley and Sons, 1977).
[188] The first Encyclopediaists were Denis Diderot, Jean Le Rond d'Alembert and their colleagues who wrote and published the *Encyclopédie* from 1751 until 1772.
[189] Douglas Lenat and R. Guha, *Building large knowledge-based systems: representation and inference in the Cyc project* (Reading, MA: Addison-Wesley, 1990).
[190] Minsky.
[191] McCarthy.
[192] Yorick A. Wilks, Brian M. Slator and Louise M Guthrie, *Electric words: dictionaries, computers, and meanings* (Cambridge, MA: MIT Press, 1996).
[193] Kant, para. 35.

The Computational Phenomenologists

While computational phenomenology can be understood to be in opposition to the project of the neo-Encyclopediaists, the neo-Encyclopediaists' turn away from teleology (in favor of lexicography) makes it clear that this opposition is more of a tension than an unbridgeable gap. In fact, the two strands can both be understood as pursuing different forms of phenomenology, one more Edmund Husserl-inspired (i.e., transcendental) and the other more Martin Heidegger-inspired (i.e., existential).

Disbelief in structured rules with intentional content has spawned several different research paradigms some of which will here be subsumed under the label of "computational phenomenology." One paradigm, known as "connectionism"[194] is an attempt to replace rules with digitally simulated "neural nets." Another paradigm, "situated action"[195] or "behavior-based AI"[196] couples the "neural nets" of connectionism to robotic (hardware and software) bodies with sensors. The research agenda of the latter group is, in many ways, a direct descendent of cybernetics insofar as it insists on the employment of feedback circuits and the disruption of internal representation versus external world dichotomies created in and by early AI work. Finally, what is here labeled computational phenomenology, is also meant to encompass recent work in "distributed AI"[197] and "multi-agent systems"; such work takes its metaphors of interaction from social systems (e.g., the systems of various scientific communities for the publication and archiving of journal articles) instead of the metaphors of the isolated thinker preferred by early-AI researchers.

The Aesthetic Turn

The work of the computational phenomenologists constitutes an "aesthetic turn"[198] in AI research since they focus attention on the aesthetic dimensions of cognition including the senses, the body, and the social and physical environment of perception. While the neo-Encylopediaists might be seen as an outgrowth of an older, "symbolic AI," computational phenomenology has been formulated in opposition to symbolic AI. Pivotal to the computational phenomenologists' position has been their understanding of common sense as a

---

[194] James L. McClelland and David E. Rumelhart (eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition; Volumes 1 and 2* (Cambridge, MA: MIT Press, 1986).

[195] E.g., Philip E. Agre and David Chapman, "Pengi: An Implementation of a Theory of Activity," in *Proceedings of the Fifth National Conference on Artificial Intelligence* (Seattle, WA: Morgan Kaufmann, 1987), 268-272.

[196] Rodney Brooks, "Intelligence Without Representation," *Artificial Intelligence* 47 (1991): 139-160.

[197] Les Gasser, "Social Conceptions of Knowledge and Action: Distributed Artificial Intelligence and Open Systems Semantics," *Artificial Intelligence* 47 (1991): 107-138.

[198] Sack.

negotiated process as opposed to a huge database of facts, rules, or schemata. This position is often repeated by the computational phenomenologists: "It should come as no surprise that the area in which [symbolic] artificial intelligence has had the greatest difficulty is in the programming of common sense. It has long been recognized that it is much easier to write a program to carry out abstruse formal operations than to capture the common sense of a dog. This is an obvious consequence of Heidegger's realization that it is precisely in our 'ordinary everydayness' that we are immersed in readiness-to-hand."[199] In other words, common sense is a faculty engendered by our encounters with "nature" and others, i.e., that said by Kant (according to Deleuze) to engender an "aesthetic common sense."

Husserl, Heidegger, and AI

The references to Martin Heidegger used by the computational phenomenologists can be seen as a contemporary manifestation of a debate between AI software designers that began as a philosophical debate initiated by Hubert Dreyfus.[200] Dreyfus and several of his colleagues (especially John Searle) have been critiquing AI (particularly symbolic AI) for over thirty years. Dreyfus has pointed out the close philosophical affinities between the projects of symbolic AI and Edmund Husserl's transcendental phenomenology and its differences from a Heideggerian existential phenomenology.[201] In particular, Dreyfus details the relationship between Husserl's philosophical project and Marvin Minsky's "frames" proposal for encoding common sense.[202] (Dreyfus and Hall, 1982, pp. 19-22).

Given Husserl's deep intellectual debts to Kant, it is understandable that Lyotard would compare Minsky's proposal to Kant's idea of reflective judgement.[203] Thus, these philosophical critiques of AI (e.g., of Dreyfus and Lyotard) give one a means of seeing how symbolic AI's proposals to encode common sense (e.g., Minsky's proposal) inherit the limitations of Kant and Husserl; and, also, the critiques illustrate how Heidegger's critique of Husserl is reflected in the computational phenomenologists' critique of symbolic AI. However, despite the frequent citation of Heidegger's work within the literature of computational phenomenology, it is not clear whether computational phenomenology is a Heideggerian project. In many ways, computational phenomenology is a self-contradictory[204] effort to "enframe"[205] Heidegger's critique of Husserl in a set of technologies.[206]

---

[199] Winograd and Flores, 98.
[200] Hubert L. Dreyfus, *What computers can't do; a critique of artificial reason. 1st edition* (New York: Harper and Row, 1972).
[201] Dreyfus and Hall, 2-27.
[202] Dreyfus and Hall, 19-22.
[203] Lyotard, 1991, 15.
[204] Cf., Richard Coyne, *Designing Information Technology in the Postmodern Age* (Cambridge, MA: MIT Press, 1995), 177.

AI and Cultural Difference

When AI has been dependent upon a Kantian-influenced vocabulary (e.g., the terms "schema," "common sense," and "teleology") its inability to articulate cultural difference is reminiscent of Kant's own limitations or oversights (e.g., with respect to gender differences). For example, in AI discussions of common sense, few researchers have asked *whose* common sense is under consideration preferring, instead, to assume that common sense is common to all humans and not culturally specific.

Even with the "aesthetic turn" in AI, practically no work has been done in AI on culture (e.g., the (re)production of differences of gender, sexuality, class, race, nationality, etc.). A belief in aesthetics as culturally invariant is obviously a useful one for a liberal, Enlightenment politics that Kant's theories of universal subjectivity contribute to. AI and cognitive science, in general, are very much in the vein of Kant's cosmopolitan universalism in their hypothesis of universal cognitive mechanisms "executable" on all sorts of (silicon and carbon-based) "hardware." What this hypothesis of a universal subjectivity leaves unthought is that significant differences between people do exist and, furthermore, the introduction of powerful technologies, like AI, can change people even more by changing their day-to-day lives. As a result, AI and its critics have largely been blind to the ethical implications of AI[207] and its implications for post-Kantian aesthetics.

Nevertheless, some AI work has been done addressing what could be *interpreted* as cultural difference. For instance, ideological difference has been modeled[208] as a difference of teleology (i.e., a difference of goals and the interrelationships between goals); expert/novice differences in education and learning have been modeled as differences of number, detail, type, and interrelationships of rules and schemas;[209] differences between the mentally healthy and the mentally ill (e.g., Kenneth Colby's simulation of a paranoid mind[210]) have been computationally modeled as differences of beliefs and

---

[205] Martin Heidegger, *The question concerning technology, and other essays*, translated by William Lovitt (New York: Garland Pub., 1977).
[206] Cf., Hubert L. Dreyfus, *What Computers Still Can't Do* (Cambridge, MA: MIT Press, 1992).
[207] Sack, 1997.
[208] E.g., Robert P. Abelson and J.D. Carroll, "Computer Simulation of Individual Belief Systems," *American Behavior Scientist*, 8 (1965): 24-30; Jaime Carbonell, *Subjective Understanding: Computer Models of Belief Systems*, Ph.D. Thesis (New Haven, CT: Yale University, 1979).
[209] E.g., Etienne Wenger, *Artificial Intelligence and Tutoring Systems: Computational and Cognitive Approaches to the Communication of Knowledge* (Los Altos, CA: Morgan Kaufmann Publishers, 1987).
[210] Kenneth Mark Colby, "Modeling a Paranoid Mind," *Behavioral and Brain Sciences*, 4 (1981): 515-534.

intentions. Although such work does engage the problematics of such important cultural phenomena as ideology, education, and mental illness, it implicitly assumes that differences of culture are personal differences by attempting to represent them exclusively with "mental," "internal" constructs like goals, plans, beliefs and intentions. Such work reduces the public to the private by ignoring the ways in which social interaction can be (re)productive of cultural difference.

This weakness of AI is not surprising given that the central metaphor of the discipline has been – not *minds* – but *mind*-as-machine. Marvin Minsky's more recent work[211] stretches this metaphor by hypothesizing that a mind is composed of a society of "agents." This work is a shift away from a Kantian vocabulary to a vocabulary of psychoanalysis.[212] Other, newer work in distributed artificial intelligence,[213] multi-agent systems,[214] artificial life,[215] computational models of discourse,[216] and computer-supported cooperative work[217] stretches the central metaphor of AI further by making groups and communities the object of study (rather than the mind of a single individual). Increasingly these new offshoots of AI are not simply stretching the boundaries of AI but, rather, creating independent disciplines.

However, even within these newer disciplines, little attention has been paid to the issue of cultural difference. Instead, what is predominantly stressed is consensus and questions like the following: Within a community of agents how can significant difference and miscommunication be overcome to allow for coordination, agreement, and "common knowledge"?[218]

Turing's Imitation Game

Ironically, cultural difference (specifically gender) is central to what is considered by most AI researchers to be the founding essay of AI. In his essay, "Computing Machinery and Intelligence,"[219] Alan Turing proposes a Wittgensteinian[220]

---

[211] Marvin Minsky, *The Society of Mind* (New York: Simon and Schuster, 1986).
[212] For a description of "agents" and Freud's "realist" model of the ego see Elizabeth A. Grosz, *Jacques Lacan: a feminist introduction* (New York: Routledge, 1990).
[213] E.g., A.H. Bond and Les Gasser (eds.), *Readings in Distributed Artificial Intelligence* (Los Altos, CA: Morgan Kaufmann Publishers, 1988).
[214] E.g., Ronald Fagin, J.Y. Halpern, Y. Moses and M.Y. Vardi, *Reasoning About Knowledge* (Cambridge, MA: MIT Press, 1995).
[215] E.g., Rodney Brooks and Pattie Maes (eds.) *Artificial life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems* (Cambridge, MA: MIT Press, 1994).
[216] E.g., Barbara Grosz and Candace Sidner, "Attention, intentions, and the structure of discourse," *Journal of Computational Linguistics* 12 (3) (1986): 175-204.
[217] Winograd and Flores.
[218] Cf., Fagin et al.
[219] Alan Turing, "Computing Machinery and Intelligence," *Mind*, Volume LIX, No. 236 (1950): 433-460.

(Ludwig Wittgenstein) language game, the "imitation game", to replace the (what he sees as meaningless) question of "Can machines think?" Turing's "imitation game" includes a proposal to program a computer to play the role of a man attempting to imitate a woman; an intriguing proposal concerning the reproduction and representation of gender difference in computational, networked media. Turing's "imitation game" is usually re-named in the literature of AI as the "Turing Test" and re-narrated to exclude any mention of gender difference.[221]

Turing describes the imitation game like this: "It is played with three people, a man, a woman, and an interrogator who may be of either sex. The interrogator stays in a room apart from the other two. The object of the game for the interrogator is to determine which of the other two is the man and which is the woman. ... It is [the man's] object in the game to try and cause [the interrogator] to make the wrong identification. ... The object of the game for [the woman] is to help the interrogator. ... We now ask the question, 'What will happen when a machine takes the part of [the man] in this game?' Will the interrogator decide wrongly as often when the game is played like this as he does when the game is played between a man and a woman? These questions replace our original [question], 'Can machines think?'" [222]

Within the literature of AI, discussions of Turing's imitation game have focused on the role of the machine and the role of the interrogator. The role of the

---

[220] Turing's essay has been intensely discussed for half a century and yet few of these discussions link Turing's method of rephrasing the question "Can machines think" as a (language) game to the methodology that Wittgenstein used more generally to reanalyze the so-called problems of philosophy as language games. Two sorts of evidence make this link between Turing's thinking and Wittgenstein's plausible. First of all, several scholars tell of Turing/Wittgenstein interactions. See, for instance, Otto Neumaier, "A Wittgensteinian View of Artificial Intelligence" in *Artificial Intelligence: The Case Against*, ed. by R. Born (London: Croom-Helm, 1987); see, also, Andrew Hodges, *Alan Turing: The Enigma* (New York: Simon and Schuster, 1983). But, secondly, and most interestingly I believe, there are some of Wittgenstein's own writings that seem to foreshadow by almost twenty years the approach Turing takes: *...This objection is expressed in the question: 'Could a machine think?' I shall talk about this at a later point, and now only refer you to an analogous question: 'Can a machine have a toothache?' You will certainly be inclined to say 'A machine can't have toothache.' All I will do now is to draw your attention to the use you have made of the word 'can' and ask you: 'Did you mean to say that all our past experience has shown that a machine never had a toothache?' The impossibility of which you speak is a logical one. The question is: What is the relation between thinking (or toothache) and the subject which thinks, has toothache, etc.?...* Ludwig Wittgenstein, *The Blue and Brown Books: Preliminary Studies for the "Philosophical Investigations"* (New York: Harper Torchbooks, 1958), 16. Thanks to Joseph Dumit for pointing out this passage of Wittgenstein's writings.

[221] Judith Genova, "Turing's Sexual Guessing Game" *Social Epistemology*, 8(4) (1994): 313-326; Warren Sack, "Painting Theory Machines," *Art and Design*, 48 (May 1996): 80-92.

[222] Turing, 433-434.

woman has been almost entirely ignored. Yet, if one looks more closely at the woman's role in Turing's game, it is clear that variants of this role have been reiterated in popular art and performance for thousands of years. The woman's role, in Turing's game, is to compete with the machine for an identity which is properly hers to begin with (i.e., the role of "woman"). The frequently reiterated, popular fears surrounding AI and its cultural and specifically artistic precedents are the fears of this sort of role; i.e., the fears of loss of identity, fears of replacement by machine, fears of disfiguration, dismemberment, and death.

AI and Aesthetics of the Uncanny

In short, these fears of an AI machine are, specifically, the fears of the "double" as it has been explored in psychoanalytic theory[223] and in the arts, for instance, in literature[224] and film.[225] More generally, these fears can be described as those associated with the *uncanny aesthetic* discussed by Sigmund Freud[226] and others.[227] Fears of the uncanny are often associated with machines, automata, and artificially produced "doubles."

Julia Kristeva has written that "uncanniness occurs when the boundaries between imagination and reality are erased."[228] Some AI researchers have tried to disavow association between their "real" work and the imaginative, artistic tradition that explores the fears of uncanny aesthetics.[229] Yet, any review of AI and aesthetics would certainly be incomplete without mentioning AI's relationship to the aesthetics of the uncanny because popular perception (e.g., as reflected in film, television, literature, and journalists' stories about AI) is often dominated by questions of "doubling:" Will machines replace people?

Limits of the Uncanny

A poststructuralist view of Kant provides a means of understanding some of the relationships between aesthetics and issues central to AI research (e.g., common

---

[223] E.g., Otto Rank, *The Double: A Psychoanalytic Study*, translated by Harry Tucker, Jr. (Chapel Hill, NC: University of North Carolina Press, 1971).

[224] E.g., Ernst Theodor Amadeus Hoffmann, "The Sandman," in *Hoffmann's strange stories*, translated by L. Burnham (Boston: Burnham Brothers, 1855).

[225] E.g., *The Student of Prague* written by Hanns Heinz Ewers and directed by Paul Wegener, 1912.

[226] Sigmund Freud, "The 'Uncanny'," in *The Standard Edition of the Complete Psychological Works of Sigmund Freud, Volume XVII*, translated by James Strachey. (London: Hogarth Press, 1919).

[227] E.g., Julia Kristeva, *Strangers to Ourselves*, translated by Leon S. Roudiez (New York: Columbia University Press, 1991); Anthony Vidler, *The architectural uncanny: essays in the modern unhomely* (Cambridge, MA: MIT Press, 1992).

[228] Kristeva, 188.

[229] E.g., Patrick Hayes and Kenneth Ford, "Turing Test Considered Harmful," in *Proceedings of the International Joint Conference on Artificial Intelligence* (Los Altos, CA: Morgan Kaufmann Publishers, 1995), 972-977 (976).

sense and teleology). Newer offshoots of AI research tend to engage a larger variety of post-Kantian, philosophical, and critical vocabularies (e.g., those of Heideggarian phenomenology and psychoanalysis). Nevertheless, while newer work might move AI outside the limits of a Kantian-inspired aesthetics, the newer work is *not* independent of a larger discourse of aesthetics that includes issues beyond the beautiful and the sublime (e.g., the uncanny).

In fact, the above exploration of the connections between the commonsensical and the aesthetic shows a number of dependencies between seemingly different approaches to designing (AI) software and hardware. The first set of technologies discussed are based upon the assumption that commonsense can be coded, articulated, stated, or at least reasoned about prior to the use of the technology. "User friendly" software presumes a set of associations that are either coded into the software so that it behaves intelligently, or a set of associations that the user will make in trying to figure out and operate the software's interface. A second set of technologies – e.g., those of the so-called computational phenomenologists – are built with the aim that commonsensical behavior is negotiated between an agent and its environment. Thus, in this case, the associations of commonsense – e.g., Minsky's of example of how a suitcase, in certain circumstances, might be categorized as a something to stand on rather as a piece of luggage – emerge from a combination of what is perceived and what is known.

But, both the symbolic AI and computational phenomenologist perspectives assume that there does exist a set of associations (either *a priori* or negotiated) that can be labeled commonsensical. Likewise, so does the uncanny perspective on technology assume a background of the commonsensical. Art and design to produce an uncanny, alienating affect is an attempt to unravel the associations of the familiar and commonsensical. The Russian Formalist Viktor Shklovsky states that the purpose of art is to work against habitualization, familiarization, and automatization:

Habitualization devours works, clothes, furniture, one's wife, and the fear of war... art exists that one may recover the sensation of life; it exists to make one feel things, ... The technique of art is to make objects 'unfamiliar', to make forms difficult, to increase the difficulty and length of perception because the process of perception is an aesthetic end in itself and must be prolonged.[230]

This approach to art, and the common popular presentation of AI technologies (e.g., in Hollywood films) as uncanny, attempts the exact opposite of the aims of user-friendly, commonsensical, "homey" (i.e., canny) design because it specifically makes the designed artifact unfriendly, unsettling, even scary.

---

[230] Viktor Shklovsky, "Art as Technique," in *Russian Formalist Criticism*, ed. by L.T. Lemon and M.J. Reis (Lincoln, NE: University of Nebraska Press, 1965), 11-12.

But it is this very confrontation with social and political practice that the aesthetic theory of estrangement finds an apparently intractable and unyielding test. The formal and critical expression *of alienation, as the first avant-gardes found to their chagrin, does not always neatly correspond to the work of transforming or even ameliorating such conditions in practice. Formal explorations of defamiliarization based on carnivalesque reversals of aesthetic norms, substitutions of the grotesque for the sublime, the uncanny for the domestic, can all too easily be construed as decoration or caricature.*[231]

As Anthony Vidler's comment makes clear, even though this different kind of design is intended to produce scary, unsettling artifacts, sometimes it just produces silly kitsch. Ironically, these aesthetic failures might have the same flaw as the unworkable technologies that were suppose to be immediately recognizable as familiar and user-friendly: it is impossible to determine what, for all people in all circumstances, will be interpreted to be familiar and commonsensical.

Network Aesthetics: After the Uncanny and the Commonsensical

VLSCs are electronic exchanges open to cross-cultural and multi-cultural exchanges. Consequently, a design practice that assumes that the commonsense of VLSC participants can be predicted and/or pre-enumerated is an untenable approach. The bulk of this chapter has been devoted to an illustration of how previous software design approaches (especially those of AI) are closely tied to a *commonsense aesthetics*, i.e., an aesthetics that presumes a commonsense, a predictably emergent commonsense, or the uncanny, interference of the commonsense world. An alternative to these approaches must be found if we are to design for VLSCs where a potential, or virtual, commonsense is contingent upon the possible (but not necessarily probably) emergence of a community of people who create their own stable linguistic and social structure through continued interaction on the Internet. This new aesthetics, therefore, must be useful for the practices of designing for emergent communities.[232]

The term "dialectic" originates from the Greek expression for the art of conversation.[233]

It is worth remembering that a variety of aesthetic practices from art and design have been dependent upon an understanding of conversation or, more

---

[231] Vidler, 12-13.

[232] These emergent communities might, in the lexicon of Gilles Deleuze and Félix Guartarri, be called "minorities." See Gilles Deleuze and Félix Guartarri, *A thousand plateaus: capitalism and schizophrenia*, translated by Brian Massumi (London: Athlone Press, 1988), 469-471.

[233] Paul Edwards (editor-in-chief) "Dialectics," in *Encyclopedia of Philosophy, Volume 2* (New York: The Macmillan Company and The Free Press, 1967), 385.

particularly, an understanding of dialectics. The conviction that the commonsense and stereotypes of mainstream media can be challenged through design of new media has been the foundation for many philosophical and artistic-design projects produced to find new material forms for some recent theory of dialectics. At least since Socrates' time, artists, designers, and philosophers have been inventing new dialectical processes to unravel the forms of each new medium and each new commonsense. New theories of dialectics were developed by Plato, Aristotle, Immanuel Kant, Georg Wilhelm Freidrich Hegel, Karl Marx, Theodor Adorno, and others. Artists and designers have elaborated these dialectical processes for new and existing media. For example, a variety of artistic processes were developed in the early-twentieth century that can be seen as media-specific instantiations of Marx's theory of dialectics.[234] Among these processes might be mentioned Sergei Eisenstein's then-new techniques of editing and film *montage*.[235] Eisenstein's frequent collaborator[236], the Russian Formalist Viktor Shklovsky, described a set of devices used in poetry for making the unconscious conscious by making the familiar strange.[237] Shklovsky's notion of "making strange" (*ostranenie*) with poetry is comparable to Bertolt Brecht's theory of the "estrangement-effect"[238] (*Verfremdung*) in epic theater.[239] Analogous phenomena and devices – called *faktura*[240] -- were researched by the Russian Constructivists for media as diverse as architecture, painting, sculpture, and collage.

But, each of the artistic design practices mentioned differs according to the medium in which it was practiced and according to the theory of dialectics – or conversation – that it incorporated or made into material form. Obviously, as soon as conversation becomes something completely different, the aesthetics of a "dialectic" practice must also renew itself. The question is, therefore, what is the new aesthetics for VLSC, a form of conversation that involves many more people in far more complicated social and semantic dynamics than earlier forms of conversation and dialectic envisioned?

---

[234] See, for instance, Karl Marx, "Critique of Hegel's Dialectic and General Philosophy," in *Karl Marx: Early Writings* (New York: McGraw-Hill, 1983), 379-400.

[235] Sergei Eisenstein, *Film form; essays in film theory*, ed. and tr. by Jay Leyda (New York, Harcourt, Brace, 1949).

[236] See Fredric Jameson, *The Prison-House of Language: a critical account of structuralism and Russian formalism* (Princeton, NJ: Princeton University Press, 1972), 61.

[237] Victor Erlich, *Russian formalism: history, doctrine, 3rd edition* (New Haven: Yale University Press, 1965).

[238] Jameson, 58.

[239] Bertolt Brecht, *Brecht on theatre: The development of an aesthetic*, edited and translated by John Willett (New York: Hill and Wang, 1964), 70-71.

[240] "The Russian term 'faktura' literally means texture but this is inadequate to convey the ideological and artistic overtones which it carries in Russian. Faktura suggests the working of the surface of materials." From Christina Lodder, *Russian Constructivism* (New Haven: Yale University Press, 1983), 280, ff. 64. See also, Benjamin Buchloh, "From Faktura to Factography," *October*, 30 (Fall 1984).

Dialectics has always been a form of conversational interaction, but also a procedure for division or repeated logical analysis of genera into species. For all intents and purposes this activity, the division of objects into kinds and their ordering is one of the main activities of computer sciences, especially that area of computer science that concerns the construction of databases. Thus, the recent approach to, what she calls, *database aesthetics* outlined by the artist and designer Victoria Vesna is one vision of a new form for dialectical aesthetics:

Artists working with the net are essentially concerned with the creation of a new type of aesthetic that involves not only a visual representation, but invisible aspects of organization, retrieval, and navigation as well. Data is the raw form that is shaped and used to build architectures of knowledge exchange and as an active commentary on the environment it depends on – the vast, intricate network with its many faces.[241]

I agree with Vesna, but think that two aspects of her proposal need to be amplified and extended. Firstly, it is important to remember that data, especially, the data of VLSC is never raw. It is always the end result of the writing or speaking activities of some participant in conjunction with a one or more "invisible" computation procedures of organization, retrieval, or navigation. Ever since the invention of things like social security numbers, the organization and indexing of databases has been about the organization and indexing of people as well as data. Like all dialectic, or conversational, processes, VLSC is a process in which people are organized. If a democratic design is pursued, then the VLSC can be about the self-organization, the self-governance of people. In Foucault's terms it can be a *technology of the self* rather than a *technology of power.*

Secondly, the "invisible aspects of organization" are only invisible if design aesthetics adopts a point of view that puts those aspects over the "horizon line" of vision or inspection. In other words, I think a new aesthetics of conversation must have an affinity with many of the older dialectically-motivated aesthetics insofar as they were attempts to make the invisible visible.

To mark these amendments to Vesna's proposal I will refer to this new aesthetics as a *network aesthetics* to emphasize the focus of such an approach. The focus should be to show the production and reproduction of connections and associations between people and data and their interminglings (in, for example, the form of a *social cohesion*, or the emergence of a commonsense or shared set of metaphors or definitions).

There is a well-worked out philosophy of what I am calling network aesthetics. It is the descriptions of rhizomes, networks, desiring machines, and intensities

---

[241] Victoria Vesna, "Database Aesthetics: Of Containers, Chronofiles, Time Capsules, Xanadu, Alexandria and the World Brain" *Journal of Artificial Intelligence and Society* (Fall 1999).

articulated by Gilles Deleuze and Félix Guattari.[242] However, it is difficult to articulate that philosophical approach to aesthetics when the domain of analysis is the new media technologies of information processing machines and networks (rather than, for example, older media like painting, architecture, or film). Deleuze and Guattari's lexicon was largely influenced by the vocabularies of non-linear system dynamics, cybernetics, and other technical artifacts that now need to be redesigned with this new aesthetics. In short, it is hard to distinguish the figure from the ground and many people, even the experts in this philosophical lexicon, fail to distinguish the vocabulary from the domain of analysis.

Various writers have noted a convergence between Deleuze's work and the scientific theories of complexity and chaos (developed in fields like physics and computer science). Brian Massumi, the English-language translator of Deleuze and Guattari's book *Thousand Plateaus*, said the following in a recent article:

[Gilles Deleuze's] work ... could profitably be read together with recent theories of complexity and chaos. It is a question of emergence, which is precisely the focus of the various science-derived theories which converge around the notion of self-organization (the spontaneous production of a level of reality having its own rules of formation and order of connection).[243]

Manuel De Landa, in his book *War in the Age of Intelligent Machines* meticulously expounds on how Deleuze's work intersects with theories of complexity, chaos and self-organization.[244] Indeed, Deleuze emphasizes his own mathematical and scientific "borrowings" in such work as chapter 15 of his book *Logic of Sense*.[245]

---

[242] E.g., Gilles Deleuze and Félix Guartarri, *A thousand plateaus: capitalism and schizophrenia*, translated by Brian Massumi (London: Athlone Press, 1988). Gilles Deleuze and Félix Guattari, *Anti-Oedipus: capitalism and schizophrenia*, translated by Robert Hurley, Mark Seem, and Helen R. Lane (Minneapolis: University of Minnesota Press, 1983). Gilles Deleuze, *Cinema 1: The movement-image*, translated by Hugh Tomlinson and Barbara Habberjam (Minneapolis: University of Minnesota Press, 1986). Gilles Deleuze, *Cinema 2: The time-image*, translated by Hugh Tomlinson and Barbara Habberjam (Minneapolis: University of Minnesota Press, 1986). Félix Guattari, *Chaosmosis: An Ethico-Aesthetic Paradigm*, translated by Paul Bains and Julian Pefanis (Bloomington, IN: Indiana University Press, 1995).

[243] Brian Massumi "The Autonomy of Affect" *Cultural Critique* (The Politics of Systems and Environments, Part II), 31 (Fall 1995), 93.

[244] Manuel De Landa *War in the Age of Intelligent Machines* (New York: Zone Books, 1991), 234-237

[245] Gilles Deleuze *Logic of Sense,* translated by Mark Lester and edited by Constantin V. Boundas (New York: Columbia University Press, 1990), chapter 15.

Conclusions

Rather than presume a sort of commonsense aesthetics to design for VLSC, a new approach must be developed. This is true because VLSCs are often multi-cultural and cross-cultural in composition, so there is no one commonsense that can be identified *a priori*. Instead, the new approach must be based on the understanding that communities can emerge through VLSC communities and so new social and semantic relationships can be produced, reproduced, and transformed. It is possible to visualize these emerging communities and linguistic possibilities as the growth and change of networks or rhizomes.

But, to visualize them as such is problematic or at least overdetermined today. When one thinks today of, for instance, "network" is it not that the Internet itself forces all other ideas of networks out of one's head? Obviously, this is not the case for everyone, but for the many for whom this is the case, the quick association between any mention of network and the Internet creates a problem for a *network aesthetics*. How can network aesthetics be a new vocabulary for the design of network technologies if the vocabulary is immediately conflated with the object of study and design?

Perhaps, this concern for a more abstract theoretical language to describe a more material design practice is simply another misplaced nostalgia for an older kind of aesthetics. John Rajchman puts this counter into play:

What is then abstract? Today the question arises in relation to what is known as the "information" age. Perhaps some new pragmatist will apply the critique of abstractions found in Bergson and James to the very idea of information and the computational paradigm to which it belongs. Two related postulates might be distinguished. The first says that information is independent of the material medium through which it is transmitted; the second says that simulation and reality come to the same thing. Thus one "abstracts" from material support and, by replicating processes, abstracts them from the particularities of their real existence; even "life" becomes only abstract information, which can be replicated and so made artificially. The two postulates of immateriality and irreality then combine in the great conceit of the info era: that electronic devices will abolish real or material space and time and transport us all into another abstract, bodiless "space" or "reality," consummating the triumph of silicon over carbon.
By contrast in Deleuze one finds an abstraction concerned not with extracting information from things (as though the material world were so much clumsy hardware) but rather with finding within things the delicate, complicated abstract virtualities of other things. Such abstractions doesn't entail independence or transferability from material support and doesn't operate according to a logic of simulation. Rather inherent in materials it supposes the subsistence of

connections that exceed the messages of a medium and ourselves as senders and receivers of them.[246]

So, to summarize, an aesthetics for VLSC needs to focus on producing the means for visualizing and understanding how social and semantic relations intertwine and communities and commonsense emerge. It may be the case that the theory and the phenomena of these kinds of self-organizational productions and reproductions are both network-based and, thus, conceptually and materially, mutually recursive in definition; an unsettling but perhaps workable perspective.

[246] John Rajchman, "What is Abstraction?," in *Constructions* (Cambridge, MA: MIT Press, 1998), 73.

EPISTEMOLOGY

Detailed, micro-analyses of face-to-face conversation usually involve a very different kind of work and produce a very different type of research result – i.e., a very different type of knowledge -- than do macro-scale analyses of discourses involving thousands or millions of people. This micro/macro divide is a recurrent one in many of the social sciences and has been widely discussed in, for example, economics and sociology. Bridging this divide for the analysis of VLSCs is necessary because, on the one hand, the phenomenon under examination is macro-scale by definition; but, on the other hand, one of the most important motivations for analyzing VLSCs is to give participants a means to find their way and locate their position in a VLSC. Consequently, standard social scientific methods of dealing with macro-scale phenomenon by working with norms and averages are unworkable because they risk effacing the contributions of particular individuals. It is argued that a bridge can be found between micro and macro scale analyses of online conversations. This bridge is the lexicon, or what might be called the "thesaurus," of a group conversation. On the micro-scale, contributions to a conversation are judged to be coherent and cohesive partially according to whether or not they are taken to be "on topic" by the participants. Knowledge of deviation or convergence with a given topic is based on knowledge of a lexicon; i.e., according to the relationships between and the definitions of words. But, over the course of the lifetime of a group, new (e.g., slang) words are coined, some words gain new meanings and others lose their currency, connotations, or the controversy that surrounds them. Thus, conversation both depends upon and changes the lexicon or "thesaurus" of a group.

Introduction: Conversation, Medium, and Discourse

To design a tool, like a browser, for VLSC it is necessary to articulate a theory of conversation, or discourse -- an epistemology -- that describes and explains the particularities of VLSC. This chapter proposes an epistemology[247] of VLSC. Specifically I pose and respond to the following questions: What is a knowledge of conversation and what sorts of knowledge are produced through conversation, especially VLSC?

VLSC is a medium; i.e., it mediates people. As such, it connects some people together and separates others. Like previous media it functions as a substrate in which and on which groups of people constitute themselves. There is a long history of the use of media as technologies of the self,[248] as reflective and communicative media for the construction of social, psychological, economic, and political self-governing people and peoples. Diaries have been used for millennia by individuals as a medium for self-reflection, for writing down and shaping one's image of one's self.[249] Diaries are a medium that functions as a technology of the self where "self" is understood to be the self of an individual. The oral storytelling practices of folktales function in an analogous manner for the formation and description of a slightly larger self, a self of small groups of people. Oral storytelling of folktales is a means for articulating the values and identity of small, tightly knit clusters of people. The facilitation of the production of larger selves, of the selves of self-governing nations, for instance, requires a different kind of medium. Scholars such as Benedict Anderson have shown how the mass production capabilities of high-speed printing presses made possible the media of novels and newspaper stories that were essential to the formation of the modern nation-state.[250] Very large-scale conversations do and can function as the substrate for new kinds of selves, new sorts of groups of people, that are as yet unnamed.[251] These new groups of people can be transnational or international in scope.

This conceptualization of VLSC -- as the substrate and catalyst of community -- is concordant with a large amount of work in sociolinguistics and the sociology of knowledge. Roughly speaking, what characterizes many of these sociolinguistic and sociological approach to conversation and discourse is this: through the

---

[247] *Epistemology: the study of the origin, nature, and limits of human knowledge.* "Epistemology," In *Encyclopedia Britannica*, see http://search.britannica.com/bcom/eb/article/0/0,5716,33390+1,00.html

[248] Michel Foucault, "Technologies of the Self," in *Ethics: Subjectivity and Truth; Essential Works of Foucault 1954-1984, Volume I,* edited by Paul Rabinow (New York: The New Press, 1997).

[249] *Ibid.*

[250] Benedict Anderson, *Imagined communities: reflections on the origin and spread of nationalism* (London: Verso, 1983).

[251] Gilles Deleuze and Félix Guartarri, *A thousand plateaus: capitalism and schizophrenia*, translated by Brian Massumi (London: Athlone Press, 1988), 469-471.

production and reproduction of a way of speaking and/or writing about certain, pivotal subjects a group is formed and distinguishes itself from other groups. Thus, chemists in the eighteenth century distinguished themselves from alchemists by developing a new discourse that we now recognize as the science of chemistry. Rather than talking about water as an essential element, chemists talk of the combination of hydrogen and oxygen. So, a new way of speaking and writing simultaneously produces a new group (e.g., chemists) and unravels or divides itself from a preexisting group (e.g., alchemists).

A way of speaking and writing (re)produces limits and possibilities for the way a subject can be spoken and/or written about and, simultaneously, (re)produces a social structure (e.g., a group or community). This way of thinking about the process and product of verbal interaction is well-known in, for instance, conversational analysis.[252] This way of describing the product or production of written and conversational forms has been termed "a discourse" by various continental theorists:

... continental discourse theorists such as Foucault, Lyotard, Donzelot, Pêcheux, and De Certeau tend to use the term "discourse" to refer to relatively well-bounded areas of social knowledge. So, at any given historical conjuncture, it is only possible to write, speak, or think about a given social object (madness, for example) in specific ways and not in others. "A discourse" would then be whatever constrains -- but also enables -- writing, speaking, and thinking within such specific historical limits. Thus while a discourse can be thought of as linguistic in one sense, it also has to be treated in terms of the conditions of possibility of knowing a specific social object.[253]

From this continental perspective it is, therefore, possible to talk about, for instance, "the discourse of chemistry." This usage of the term "discourse" (i.e., the use of the term discourse preceded by a definite or indefinite article like "the" or "a") is sometimes at odds or appears more-or-less incomprehensible to practitioners of other sorts of Anglo or American forms of discourse analysis.[254]

---

[252] *Now if we can take it that to some extent "hotrodders" is a category that is by and large employed by kids to characterize themselves, and whose use, to some considerable extent, they enforce, and whose properties they enforce, and obviously it's, at least to some extent, a category that rebellious persons can use, then at least one of the initial questions we might ask is: Why should it be the case that at least some people who go about doing kinds of rebellion, do it by formulating themselves as a particular type? That is, why do they set up a type? Why don't they try to make themselves observable as "individuals," so to speak?* Harvey Sacks, "'Hotrodders' as a Revolutionary Category," in *Lectures on Conversation, Volume 1*, edited by Gail Jefferson (Cambridge, MA: Blackwell, 1992), 172.

[253] A. McHoul, "Discourse," In R.E. Asher (editor-in-chief) *The Encyclopedia of Language and Linguistics* (New York: Pergamon Press, 1994), 944-945.

[254] Alastair Pennycock, "Incommensurable Discourses," *Journal of Applied Linguistics,* Volume 15, Number 2, 1994.

In her book that compares and contrasts six different Anglo-American approaches to discourse analysis (speech act theory, pragmatics, ethnomethodology, interactional sociolinguistics, ethnography of communication, and variation theory), Deborah Schiffrin states

Discourse has often been viewed in two different ways: a structure, i.e., a unit of language that is larger than the sentence; and the realization of functions, i.e. as the use of language for social, expressive, and referential purposes.[255]

In other words, from an Anglo-American perspective, "discourse" is a name for a sequence of sentences (a structure) or a certain kind of language use (a function). But, from a continental perspective, "discourse" is either the result of language use or the background conditions or context for a given sequence of sentences. Borrowing the trope of "figure/ground" from art history, one might say the difference between scholarly approaches to discourse analysis arise from the use of the term "discourse" to describe *figure* versus use of the term "discourse" to describe *ground*. Or, alternatively, the conflict involves the use of "discourse" as a name for *text* versus the use of "discourse" as a name for *context*.

Rather than sort out this knotty conflation and conflict of terminology, I will try to find a way around it. From a continental perspective one might talk about how a VLSC produces or reproduces a given or new discourse. From an Anglo-American perspective one might say that a VLSC is a discourse. Instead, I will simply state that a VLSC produces, reproduces, and relies on a set of social and semantic relationships. In the language of mathematics, one might say that there exists a mutually recursive relationship between a VLSC and a set of social and semantic networks. Or, one might say, the coherence of a VLSC depends upon social and semantic background knowledge, but, this background knowledge is also, at least partly, a product of the VLSC.

Three Dimensions of Commonsense, Conversational Knowledge

For conversations of a smaller scale (i.e., smaller than VLSC) it is possible to see when the background knowledge of a conversation is being abused or flouted. Commonsense, conversational, background knowledge can be described in a variety of ways; e.g., as a set of common associations and common terms, as a set of social and semantic networks, or -- as will be elaborated below -- as a set of meta-functions named the *interpersonal*, the *textual*, and the *ideational* meta-functions by Michael Halliday.[256]

Divergences or differences of routine, conversational, background knowledge can produce misunderstandings and conflict, but they can also produce comedy. Consider the following one-liner from comedian Stephen Wright:

---

[255] Deborah Schiffrin, *Approaches to Discourse* (Cambridge, MA: Blackwell, 1994), 339.
[256] Michael A. K. Halliday. *An Introduction to Functional Grammar, Second Edition* (London: Edward Arnold, 1994), 179.

I was driving down the highway one morning and I saw a billboard advertising a restaurant that said "Breakfast any time" so I stopped and ordered French toast in the Renaissance.

The social coherence of a group underwrites conversation and depends upon a number of things. Semantics is just one of these things, but Wright's joke illustrates how the production of common terms -- a shared semantics -- is important to conversation.

If the terms of conversation, but not the conventional turn-taking "rules" are followed, another sort of nonsense is produced. Lewis Carroll illustrates the "rules" of riddles when he has the characters of *Wonderland* violate them.

"Have you guessed the riddle yet?" the Hatter said, turning to Alice again.
"No, I give it up," Alice replied: "what's the answer?"
"I haven't the slightest idea," said the Hatter.
"Nor I," said the March Hare.
Alice sighed wearily. "I think you might do something better with the time," she said, "than wasting it in asking riddles that have no answers."[257]

The common terms and rules of conversation are tightly coupled in the production of the cohesion of a conversation. When the cohesion is deliberately undone, the conversation is unhinged as this snippet from Eugene Ionesco's absurdist play *The Bald Soprano* illustrates. Suddenly, in this sequence dialogue, all of the people being discussed are named Bobby Watson:

Mrs. Smith: *"But who would take care of the children? You know very well that they have a boy and a girl. What are their names?"*
Mr. Smith: *"Bobby and Bobby like their parents. Bobby Watson's uncle, old Bobby Watson, is a rich man and very fond of the boy. He might very well pay for Bobby's education."*
Mrs. Smith: *"That would be proper. And Bobby Watson's aunt, old Bobby Watson, might very well, in her turn, pay for the education of Bobby Watson, Bobby Watson's daughter. That way Bobby, Bobby Watson's mother, could remarry. Has she anyone in mind?"*
Mr. Smith: *"Yes, a cousin of Bobby Watson's."*
Mrs. Smith: *"Who? Bobby Watson?"*
Mr. Smith: *"Which Bobby Watson do you mean?"*
Mrs. Smith: *"Why, Bobby Watson, the son of old Bobby Watson, the late Bobby Watson's other uncle."*[258]

---

[257] Lewis Carroll, "Alice's Adventures in Wonderland," in *The Annotated Alice* (New York: Clarckson N. Potter, Inc., Pub, 1960), 97.
[258] Eugene Ionesco, *The bald soprano and other plays*, translated by Donald M. Allen (New York: Grove Press, 1958).

Obviously, writers and comics know and bend the common terms and rules of conversation in order to produce these sorts of effects. Using insights of this sort, scholars like Roman Jakobson have been able to explain the linguistic workings of avant-garde artistic literature[259], but the wittiness of more common performances also often depends upon an explicit understanding of how conversation engenders social cohesion and/or how the norms can be manipulated to reveal or break the underpinnings of social cohesion.[260] It is

---

[259] Roman Jakobsen, *Verbal art, verbal sign, verbal time*, edited by Krystyna Pomorska and Stephen Rudy (Oxford, UK: Blackwell, 1985).

[260] In linguistics there exists a principle called Ziff's Law; P. Ziff, *Semantic Analysis* (Ithaca, NY: Cornell University Press, 1960). Ziff's Law is the observation that any arbitrary string can be interpreted as a proper name. This is often mentioned as a serious difficulty for the construction of computer programs to parse natural language texts. However, it is also the main observation underlying Bud Abbott's and Lou Costello's famous "Who's on first?" skit first performed on the *Kate Smith Radio Hour* in 1936.

Costello: Look Abbott, if you're the coach, you must know all the players.
Abbott: I certainly do.
Costello: Well you know I've not met the guys. So you'll have to tell me their names, and then I'll know who's playing on the team.
Abbott: Oh, I'll tell you their names, but you know it seems to me they give these ball players now-a-days very peculiar names.
Costello: You mean funny names?
Abbott: Strange names, pet names...like Dizzy Dean...
Costello: His brother Daffy
Abbott: Daffy Dean...
Costello: And their French cousin.
Abbott: French?
Costello: Goofé
Abbott: Goofé Dean. Well, let's see, we have on the bags, Who's on first, What's on second, I Don't Know is on third...
Costello: That's what I want to find out.
Abbott: I say Who's on first, What's on second, I Don't Know's on third.
Costello: Are you the manager?
Abbott: Yes.
Costello: You gonna be the coach too?
Abbott: Yes.
Costello: And you don't know the fellows' names.
Abbott: Well I should.
Costello: Well then who's on first?
Abbott: Yes.
Costello: I mean the fellow's name.
Abbott: Who.
Costello: The guy on first.
Abbott: Who.
Costello: The first baseman.
Abbott: Who.
Costello: The guy playing...
Abbott: Who is on first!

equally as obvious that anyone who finds these manipulations funny or absurd has a set of well-developed intuitions about the rules and terms of conversation: the commonsense knowledge of conversation.

Each of the comedic examples above illustrates a different meta-function of language. According to Michael Halliday[261], language has at least three meta-functions: (1) *ideational*: language can represent ideas; (2) *interpersonal*: language functions as a medium of exchange between people; and, (3) *textual*: language functions to organize, structure, and hold itself together; this function allows the various devices of cohesion, including citation, ellipsis, anaphoric reference, etc. to be used. The Steven Wright joke shows what can happen when the ideational meta-function breaks down. The selection from *Alice in Wonderland* illustrates the breakdown of the interpersonal meta-function. And, in Ionesco's dialogue the textual meta-function is thwarted by a breakdown of lexical cohesion. The point of these examples is simply to give examples of what might be considered the three different dimensions of commonsense knowledge about conversations that must be in place for a conversation -- and so, transitively, a group of interlocutors -- to hold together.

When one or all of these sorts of conversational background knowledge fall apart, the result can be funny.[262] But, by only citing the absurd and the comedic it is difficult to picture what can be lost if the terms or rules of conversation are questioned or broken. While these questions and breaks can be funny, they can also arouse anger or mistrust.

Harold Garfinkel asked his students to document this, the breakdown of common terms assumed in conversation; i.e., to document the breakdown of the *ideational* meta-function. In the course of everyday conversation, Garfinkel's students questioned the assumed, common terms. The results were graphic. In the following accounts Garfinkel's students play the role of the so-called "experimenter" (E).

The subject was telling the experimenter, a member of the subject's car pool, about having had a flat tire while going to work the previous day.
(S) I had a flat tire.
(E) What do you mean, you had a flat tire?

Costello: I'm asking you who's on first.
Abbott: That's the man's name.

...

[261] Michael A. K. Halliday. *An Introduction to Functional Grammar, Second Edition* (London: Edward Arnold, 1994), 179.

[262] The violation of these sorts of commonsense knowledge can be seen as funny as can the violation of a large variety of everyday expectations. See Sigmund Freud, *Jokes and their relation to the unconscious*, translated and edited by James Strachey (New York: Norton, 1960).

She appeared momentarily stunned. Then she answered in a hostile way: "What do you mean, 'What do you mean?' A flat tire is a flat tire. That is what I meant. Nothing special. What a crazy question!"

...

"On Friday night my husband and I were watching television. My husband remarked that he was tired. I asked, 'How are you tired? Physically, mentally, or just bored?'"
(S) I don't know, I guess physically, mainly.
(E) You mean that your muscles ache or your bones?
(S) I guess so. Don't be so technical.
(After more watching)
(S) All these old movies have the same kind of old iron bedstead in them.
(E) What do you mean? Do you mean all old movies, or some of them, or just ones you have seen?
(S) What's the matter with you? You know what I mean.
(E) I wish you would be more specific.
(S) You know what I mean! Drop dead!

...

The victim waved his hand cheerily.
(S) How are you?
(E) How am I in regard to what? My health, my finances, my school work, my peace of mind, my ...?
(S) (Red in the face and suddenly out of control.) Look! I was just trying to be polite. Frankly, I don't give a damn how you are.[263]

These examples make the risks clear. By questioning the common terms of conversation the students threaten the existence, or at least the smooth functioning, of various small groups of people: the car pool, the marriage, the friendship.

Questioning the common terms -- the ideational meta-function of language -- has risks. Analogously, there are risks to questioning the textual and interpersonal metafunctions. Using an ethnographic methodology, John Gumperz and his colleagues have documented how the textual and interpersonal meta-functions of language can break down in cross-cultural conversational situations. Consider the following utterances spoken by a Malaysian-born Indian immigrant in a London Adult Education class discussion about mortgages:

---

[263] Harold Garfinkel, *Studies in Ethnomethodology* (Cambridge, MA: Blackwell Publishers Inc., 1967), 43-44.

Mortgages. If you are to buy a house. Who can get and who cannot get. What assumptions we made, what? If you work. If you don't work, can you get a mortgage?[264]

Gumperz et al. comment on this example: *...the difficulties here are in following the connections that are being made, and consequently in understanding the intention of the final questions. [The] example starts with a string of noun phrases that appear to announce the intended topics. Is the final question intended to elicit a review of the assumptions made at another time, or is it the commencement of discussion of the topic of "who can get and who cannot get" a mortgage?*[265]

Of course, the "difficulties" that Gumperz et al. mention are their difficulties, not the difficulties of the speaker or her audience who are also, largely, English-speaking Indian immigrants. Gumperz et al. show how the structures and resources of grammar, prosody, and intonation of Hindi, Urdu, Punjabi, Gujerati, and Marathi are employed to join together multiple sentences when speakers of these North Indian languages speak English. Thus, the difficulties in resolving the cohesion are mostly difficulties for native English speakers, not English-speaking Indians and Pakastanis. Consequently, even in a situation where the language being spoken is English and everyone in the situation is perfectly fluent in English, cross-cultural ties cannot, at times, be created because the textual and interpersonal meta-functions are produced very differently by members of different cultural groups.

This can have grave repercussions in legal, medical, and employment situations. In such situations bilinguals are sometimes thought to not be telling the truth because their testimony seems to be self-contradictory when interpreted by monolinguals; or, the bilingual does not receive the medical care they need because the doctor doesn't understand them; or, the bilingual doesn't get the job because the monolingual thinks the bilingual is hard to understand. In other words, in such situations -- unlike the example situations of Garfinkel -- the social fabric of a group is not ripped; rather, the group or social relationship is never threaded together or is clipped off right from the start.

Obviously the ideational, interpersonal, and textual relations established through inter- and intra-cultural conversational interactions do not remain static. Some groups become closer knit over time. Others fall apart. Intercultural, multilingual interactions can produce creoles and new forms of intelligibility; or, unfortunately, such situations can deteriorate through repeated miscommunication and so cross-cultural conversation can become more and more difficult. To understand these shifts it is necessary to understand how a series of conversational

---

[264] John J. Gumperz, Gurinder Aulakh, and Hannah Kaltman, "Thematic structure and progression in discourse," In *Language and Social Identity*, edited by John J. Gumperz (Cambridge, UK: Cambridge University Press, 1982), 25.
[265] *Ibid.*

interactions "add up" and, thereby, influence the performance of the meta-functions of language. For instance, how can good "first impressions" make interactions thereafter easy? Or, how can a set of misunderstandings lead to diminished rapport between people who have gotten along for years?

The Micro-Macro Divide

It is quite easy to roughly characterize the difficulties of visualizing VLSC as a substrate and catalyst for community. It is a "chicken and egg" problem. The communities of VLSC do not pre-exist the VLSC except in some very vague manner. The texture and ideas of online communities come through collective actions and individual interactions, but it is difficult to see how a multiplicity of such (inter)actions might add up to, for instance, a coherent conversation, or, what in continental theories might be called a *discourse*. The difficulty is what is often referred to as the *micro-macro problem* by social theorists: how can a large number of individual interactions add up to a larger social or political force and, vice versa, how does a larger social force act on small-scale, even intimate interactions?

Social theory has been in general terms concerned with different levels of analysis. In economic theory we are familiar with the idea of micro and macro economics to describe these different levels. Micro economics is concerned with the economic activity of individual economic units such as the household. Macro economics considers the behavior of the economy as a whole. Political science and sociology also work with such a distinction. In commonsense terms the micro level is the level of everyday interaction typically involving face-to-face negotiation between individuals. By contrast the macro level refers to the global structure of societies, and the analysis of major institutions such as the interface between the economy and politics; it also deals with large-scale collective action such as global social movements. The majority of social theorists recognize implicitly some form of this distinction, and various social theories have attempted to explicate the relationship between the micro and macro levels.[266]

Just as there is a micro/macro divide in economics and sociology research, there is also a micro/macro divide in discourse and conversation analysis work. The great majority of work done on conversation, by linguists and sociologists, are micro-analyses of interactions between a small group of people. For example, work in conversation analysis often examines interactions between two or three people.[267] Larger-scale work includes analyses of individual classrooms or small group interactions involving ten, twenty, or thirty people.[268] But large-scale work

---

[266] Bryan S. Turner, "The Micro-Macro Problem," in *The Blackwell Companion to Social Theory*, edited by Bryan S. Turner (Malden, MA: Blackwell, 1996), 222.

[267] E.g., Harvey Sacks, *Lectures on Conversation, Volumes 1 and 2*, edited by Gail Jefferson (Cambridge, MA: Blackwell, 1992).

[268] E.g., J.McH. Sinclair and M. Coulthard, *Towards an Analysis of Discourse: The English Used by Teachers and Pupils* (London: Oxford University Press, 1975).

in examining interactions between hundreds or thousands of people, for instance, in online newsgroups or interchanges in scientific literatures, usually effaces so many of the rich language details that micro-analyses take particular care with that these large-scale investigations are a completely different species of work. These studies are, in other words, macro-analyses and it is difficult to see whether or how they complement the work of micro-analyses. Historically, the most expedient thing to do has been to choose either a micro- or a macro-analysis methodology and then ignore the results of the other. However, this is not an option for VLSC because it is large-scale, thus macro, in size, but its rich details are what makes it a conversation rather than just, for example, an "information superhighway."

For instance, recent work by Steve Whittaker, Loren Terveen, Will Hill, and Lynn Cherny on "The dynamics of massive interaction" analyzes the headers (i.e., the *to:*, *from:*, *references:*, etc. slots) of several million email messages to investigate online conversational dynamics, but they do this analysis by completely ignoring the contents of the messages.[269] Arguably, this sort of methodology -- like a lot of work in sociology on social networks[270] and co-citation analysis[271] -- is an exploration of some of the interpersonal dimensions of the medium of VLSC, but it leaves untouched the textual and ideational relations established or broken by VLSCs. Since the production and reproduction of social groups through VLSC is a function of at least all three of these aspects of language (the interpersonal, the textual, and the ideational), a strictly social network-based examination (whose responding to whom) is not sufficient as a complement to detailed micro-analysis work.

Conversely, much other large-scale work has been done on text corpora that reveals recurrent patterns of ideational and textual relations, but that ignores how a series of texts can produce or reinforce a social network, a set of interpersonal relationships. For example, corpus-based, computational linguistics work has developed technologies for automatically compiling rough-draft thesauri given a large archive of texts[272]; or, given an archive of tagged and bracketed texts,

---

[269] Steve Whittaker, Loren Terveen, Will Hill, and Lynn Cherny, "The dynamics of massive interaction," In the *Proceedings of the International Conference on Computer-Supported Cooperative Work* (Seattle, WA: Association for Computing Machinery, 1998), 257-264.

[270] E.g., Stanley Wasserman and Joseph Galaskiewicz (editors) *Advances in Social Network Analysis: Research in the Social and Behavioral Sciences* (Thousand Oaks, CA: Sage Publications, 1994).

[271] E.g., E. Garfield. *Citation Indexing: Its Theory and Applications in Science, Technology and Humanities* (New York: John Wiley, 1979).

[272] For contemporary work see Gregory Grefenstette, *Explorations in Automatic Thesaurus Discovery* (Kluwer Academic Publishers: Boston, 1994); Marti Hearst, "Automated Discovery of WordNet Relations," in *WordNet: An Electronic Lexical Database*, edited by Christiane Fellbaum (Cambridge, MA: MIT Press, 1998); Sanda Harabagiu and Dan Moldovan, "Enriching the WordNet Taxonomy with Contextual Knowledge Acquired from Text," in *Natural Language Processing and Knowledge*

machines have been developed to automatically generate a grammar and a parser;[273] or, given a set of texts that mention many of the same people or places, some newly developed machines can now automatically hyperlink the texts so that entities in one text are automatically connected to mentions of the same entities in other texts.[274] Many of these same techniques have been taken up by sociologists of science working in the area of actor-network theory.[275] For example, Geneviéve Teil and Bruno Latour describe a machine that uses measurements of conditional probability and mutual information to automatically compile a rough-draft thesaurus from a corpus of scientific abstracts.[276]

The difficulty with visualizing the conditions and productions of VLSC is therefore the following. Even though it would be ideal to simply "scale-up" the methodologies of the micro-analysis of conversations and discourse, such methodologies can no more be "scaled-up" than the rich insights into bird flight gathered by a keen-eyed ornithologist can be "scaled-up" to analyze the dynamics of jet airplane travel. This is because these micro-scale analyses require too much of the analyst. Often a micro-analysis of a conversation demands that the analyst identify the intentions of the participants; this is nontrivial if not impossible to do.

Moreover, previous attempts create theoretical tools for the examination of large text corpora have often neglected one or another meta-function of language (e.g., the interpersonal, the textual, the ideational) that is clearly important for conversational interaction. These large-scale theories of language are not adequate as large-scale theories of conversation because they leave too many details out. Large-scale work tends to fall into either (a) a social network type of

---

Representation: Language for Knowledge and Knowledge for Language, edited by S. Shapiro and L. Iwanska (Cambridge, MA: AAAI/MIT Press, 1999). For the history of this field see, for example, Dagobert Soergel, *Indexing Languages and Thesauri: Construction and Maintenance* (Los Angeles: Melville Pub. Co., a division of John Wiley, 1974).

[273] See, for example, D. Magerman, *Natural language parsing as statistical pattern matching*, Ph.D. thesis (Stanford, CA: Stanford University, 1994).

[274] See Amit Bagga, Breck Baldwin, and Sara Shelton (editors) *Coreference and Its Applications, Proceedings of the Workshop* (New Brunswick, NJ: Association for Computational Linguistics, 22 June 1999); see also, Stephen Green, *Automatically generating hypertext by computing semantic similarity*, Ph.D. thesis (Toronto, Canada: Department of Computer Science, University of Toronto, August 1997).

[275] John Law and John Hassard (editors) *Actor network theory and after* (Oxford, UK; Malden, MA: Blackwell/The Sociological Review, 1999).

[276] Geneviéve Teil and Bruno Latour, "The Hume Machine: Can Association Networks do more than Formal Rules?," in *Stanford Humanities Review 4.2* (1995): 47-65. Teil's and Latour's work is one of the latest outgrowths of a long line of such computerized text analysis work (on co-word analysis) conducted at the Centre de Sociologie de l'Innovation, Ecole des Mines de Paris. See also, for example, Michel Callon, John Law, and Arie Rip (editors) *Mapping the Dynamics of Science and Technology*, London: Macmillan, 1986.

work, that usually leaves out a lot about the contents of the text or talk exchanged; or, (b) a corpus-based linguistics style of work that tends to overlook too much of the interpersonal, or social structure of the texts examined. If, however, these two different styles of macro-analysis could be combined, then a richer picture of the combined social and semantic (re)productions of VLSC could be painted.

Thesaurus as Fulcrum

If these different sorts of macro-analysis (social network-based and computational linguistics-based) are to be connected together a linkage point must be found. By looking at the history of discourse analysis a linkage point can be found: it is the *thesaurus*.

In the first essay in linguistics to mention discourse analysis, Zelig Harris provided a rough-draft version of this linkage point. In his 1952 journal article, Harris explains that the key to discourse analysis is to find corpus-specific equivalencies:

Suppose our text contains the following four sentences: The trees turn here about the middle of autumn; The trees turn here about the end of October; The first frost comes after the middle of autumn; We start heating after the end of October. *Then we may say that* the middle of autumn *and* the end of October *are equivalent because they occur in the same environment (*The trees turn here about *--), and that this equivalence is carried over into the later two sentences. On that basis, we may say further that* The first frost comes *and* We start heating *occur in equivalent environments.*[277]

In the Anglo-American traditions of discourse analysis no one has followed up on Harris' work.[278] However, the description provided by Harris on how to find "equivalencies" is a sort of technically unnuanced description of the sort of work that some researchers in contemporary, corpus-based, computational linguistics have undertaken to automatically generate or extend thesauri. Harris' insight

---

[277] Zelig Harris, "Discourse Analysis," *Language*, 28 (1952): 1-30 and 474-494, 6.

[278] *...with Chomsky's appropriation of the notion of transformations as an intrasentential feature, and with the overwhelming dominance of linguistics by the transformational-generative movement which Chomsky came to lead, Harris' early attempt with longer stretches of texts was not followed up, and the models of discourse analysis described below [discourse analysis as influenced by Michael Halliday and conversation analysis as influenced by Harold Garfinkel) cannot be seen as direct developments of Harris's model.* Kirsten Malmkjaer "Discourse and conversational analysis" In Kirsten Malmkjaer (editor) *The Linguistics Encyclopedia* (New York: Routledge, 1991), 100-101. However, Michel Pêcheux and his colleagues and students in France did attempt to use Harris' (or at least Harris-like) insights to examine differences and similarities between specific discourses. See, for instance, Michel Pêcheux, *Automatic discourse analysis,* edited by Tony Hak and Niels Helsloot and translated by David Macey (Amsterdam ; Atlanta, GA: Rodopi, 1995).

about what he called his "distributional analysis of discourse" was that regularities within a given discourse, rendered as "equivalencies," could be descriptive of the cultural specificities produced and reproduced within a given discourse. By blending the technology of contemporary corpus-based linguistics with Harris' insight, it is possible to use this insight as a pivot point through which different kinds of macro-analysis connect together with the concerns of micro-analysis of conversation.

Consider the following conversational exchange:

A: What sorts of fruit do you like?
B: Oh, apples and bananas.
A: What about strawberries? Do you buy them when they're in season?
B: No, I don't really like berries.

To find the lexical cohesion between the statements in this exchange it is necessary to know that apples, bananas, strawberries, and berries in general are all kinds of fruit. Knowing this it is possible to say that this short sequence concerns fruit, but it is also possible to say that A and B have had a verbal exchange concerning fruit. In short, thesaurus-like knowledge about fruit provides a means for more specifically describing interpersonal and textual relations of the conversation.

Within the micro-analysis specialty of conversation analysis, scholars have noted the key role that thesaurus-like knowledge of categories plays in the construction of coherent sequences of dialogue. For instance, in elaborating his theory of categories, the inventor of conversation analysis, Harvey Sacks, provided the following definition and example:

Sacks refers to activities which imply identities as category-bound *activities (CBAs). His definition is as follows.* Category-bound activities*: "many activities are taken by Members to be done by some particular or several particular categories of Members where the categories are categories from membership categorization devices."*[279] *CBAs explain why, if the story read "The X cried. The Y picked it up," we might have guessed that X was a baby and Y was a mommy. Crying, after all, is something that babies do and picking up (at least in the possibly sexist 1960s) is something that mothers did.*[280]

Similar observations about the key role of semantic and pragmatic associations for given terms in the construction of the coherence and cohesion of texts was realized within some work in computational linguistics.[281] However, this

---

[279] Sacks, 249.
[280] David Silverman, *Harvey Sacks: Social Science and Conversation Analysis* (New York, Oxford University Press, 1998), 83.
[281] Cf., Jaime G. Carbonell, "Towards a Process Model of Human Personality Traits," *Artificial Intelligence* 15 (1980): 49-74.

computational linguistics work, like most other non-Harris-like, computational work on discourse analysis has been -- for all practical purposes -- a methodology of micro-analysis of conversation and discourse.[282]

What has been left unexplored is the fact that there now exist empirical methods applicable to large-scale corpora that can provide a means for documenting the emergence categories of terms, what Harris called "equivalencies" between terms. But, it is not the case that these new techniques from corpus-based linguistics can automatically bridge the theoretical chasms dividing micro from macro conversational analysis and social network versus computational linguistic macro analyses. One more theoretical insight is necessary: even as much as the textual and interpersonal relations are influenced by the ideational relations (i.e., the semantic links articulated in thesaurus-like compilations), the inverse is also true. In other words, the social and semantic aspects of VLSC are related in a mutually recursive manner: ideational → textual → interpersonal → ideational.

From the word usages (what Saussure would call *parole* or what Chomsky would call *performance*) in a corpus of texts, a set of equivalencies and thus a rough idea of semantic relations between terms can be derived with the procedures of corpus-based, computational linguistics. These "equivalencies" can be compiled as a kind of rough draft thesaurus. The categories and equivalencies in the thesaurus have, in turn, an influence of how cohesion (i.e. textual) and social (i.e., interpersonal) relations are labeled. By looking at which terms are important to a conversation (i.e., which terms label a large number of social and cohesion relations present in a corpus of texts (e.g., an archive of email messages), one can get a feel for which parts of the rough draft thesaurus are important. The ways in which these highlighted elements of the rough draft thesaurus are "spoken about" by members of the conversation provides a means for characterizing the conversation as a whole. Thus, for instance, a conversation that is dominated by terms like "hydrogen," "oxygen," etc. might be characterizes as conversation about chemistry rather than as a conversation about fashion design.

As conversations and so cultures evolve, so do the "equivalencies" or thesauri that can be derived from them. This is true too of more official, hand-compiled reference works.[283]

---

[282] Most "discourse analysis systems" that have been built in the fields of artificial intelligence and computational linguistics have been very elaborate productions constructed to illustrate the analysis of interchanges that can transcribed into one or two pages of text. E.g., James F. Allen, B. Miller, E. Ringger and T. Sikorski, "A Robust System for Natural Spoken Dialogue," in the *Proceedings of the 34th Meeting of the Association for Computational Linguistics*, 1996.

[283] The literary theorist Roland Barthes speaks of the contents of reference books, like thesauri, as "cultural codes" central to the process of reading. *The **cultural** codes, which are extremely numerous and heterogeneous, to a very large degree subsume all the other categories. They speak the familiar "truths" of the existing cultural order,*

Conclusions

The theory of conversation and discourse presented here is a kind of meta-theory which provides a means to bridge the gaps that exist between different kinds macro-analyses of conversation and certain types of micro-analysis of conversation. By concentrating on approaches that stress sociolinguistic and/or the sociology of knowledge (rather than more psychological or intention-based models which rely on detailed micro-understandings of participants' states of mind) it is shown that thesaurus-like knowledge of terms and categories of terms provide a means to link the various dimensions of conversational language together (the ideational, the interpersonal, and the textual); and, furthermore, this thesaurus-like knowledge also provides a way of linking micro with macro analyses because, in corpus-based, computational linguistics work, it is now well known how to create a rough-draft thesaurus from a large body of texts. Insights into the relative interdependence between the different meta-functions of conversational language and the ways in which they can be cross-indexed through thesaurus-like knowledge have been incorporated into the design of the Conversation Map system.

---

repeat what has "always been already read, seen, done experienced." ... Barthes underscores the discursive basis of the "reality" to which cultural codes refer by equating it with "the set of seven or eight handbooks accessible to a diligent student in the classical bourgeois educational system," Kaja Silverman, *The Subject of Semiotics* (New York: Oxford University Press, 1983), 241 & 274.

CONCLUSIONS

Two pieces of work have been accomplished for this dissertation. Firstly, a proof-of-concept VLSC browser was designed, implemented, and demonstrated. This VLSC browser, the *Conversation Map* system, provides a means to explore and navigate the social and semantic structures of VLSCs (e.g., Usenet newsgroups). Secondly, the thinking that went into the design of the Conversation Map system was generalized and articulated as an aesthetics, ethics, and epistemology of design for VLSC. This second, theoretical portion of the thesis provides a means to describe the emergent phenomenon of VLSC and a vocabulary for critiquing software designed for VLSC and computer-mediated conversation in general.

With the graphical interface of the Conversation Map system one can browse a set of Usenet newsgroup articles according to who is "talking" to whom, what participants are "talking" about, and the central terms and possible emergent metaphors of the conversation. To allow this combination of social and semantic navigation the Conversation Map system computes a social network corresponding to who is replying to (or citing) whose messages. The Conversation Map system also parses and analyzes the contents of the newsgroup articles to calculate a semantic network that highlights frequently used terms that are similar to one another in a VLSC. The design and implementation of these two functionalities required research and development of three new techniques of potential interest for computational linguistics and sociology:

(1) **Inter-message lexical cohesion**: An algorithm has been created for the computation of *lexical ties* between email messages that takes into account information about both threading and quotation in the messages (thereby producing an approximation of the themes of conversation). The computation of lexical ties between messages with complex, inter-message threading and quotation structure is novel in computational linguistics.

(2) **Social cohesion**: A computable definition of *social cohesion* has been developed. Social cohesion is defined to be an overlay of lexical cohesion information on top of a social network (thus creating a rough description of what the conversations between participants are about). This computable definition of social cohesion for the analysis of online VLSC should be interesting to compare with other definitions of social cohesion that sociologists have developed. In particular, most work in social network analysis is not integrated with a sophisticated linguistic analysis of inter-personal communications (like the inter-message lexical tie analysis described above) and so this definition of social cohesion may prove to be an interesting new tool for the sociology of computer-mediated communication.

(3) **Social network-based sorting**: A means has been designed and implemented to use the computed social networks to automatically weight the importance of term entries in computed rough-draft thesauri. The terms

used more frequently as themes of conversation are weighted more heavily than those which are only mentioned but not taken up as themes (thus providing some means to highlight the possible semantics of terms that are most crucial for the cohesion of a group's conversation). No other corpus-based computational linguistics algorithms for automatic thesaurus generation have employed a measure of social relevance (like the social networks used here) to order the results of the algorithm.

The three text analysis procedures are of theoretical interest because they extend and improve upon sociology and linguistics by interfacing and extending techniques from both domains.

With the results of these three text analysis procedures a researcher can begin to look for answers to questions concerning existing and emergent groups such as television audiences, social movements, user groups, and online, citizen diplomats. Are the group members of a VLSC creating a new way of talking about themselves? Are they creating a new way of talking about certain subjects; e.g., political issues or characters in a television show? Are they a cohesive group? Do they have leaders? Or, does everyone contribute equally to the social structure of the group?

These are the tools and questions of social science and linguistics. In an effort to improve the Conversation Map system for social scientists and linguists a participatory design process is now underway to write a new version of the system in collaboration with them.

Even so, the larger goal of this thesis is to contribute towards the transformation of existing VLSCs into a set of global conversations. Consequently, the Conversation Map system has been developed as a VLSC browser that could, at least in principle, be used by anyone participating in a VLSC. It is not meant to be exclusively a tool for social scientists and linguists. In addition to the support of a researcher's inquiries, a VLSC browser should also give the newcomer or the VLSC regular the power to pursue their own questions and reflections concerning a VLSC. For instance, a newcomer might ask questions such as these: What have people been talking about? Is this a group where people talk to one another? Or, is it one where people talk at one another? Regular participants have a set of questions unlike those of the researchers or newcomers: Am I well known? Am I well connected? Are my friends well connected? Which themes of discussion are divisive? Which themes bring the group together?

Note that even though each of these three sets of questions is different (i.e., the questions of the researcher versus those of the newcomer versus those of the regular participant), they are all questions of mapping, summary, or overview. The goal is to find the broad outlines and emergent properties of VLSCs as compared to, for instance, a discrete sentence or phrase within the archive of a

conversation that might contain the answer to a given question. In other words, all of these questions are questions of mapping, summarization, browsing, and navigation rather than, for example, questions of information retrieval.

While it may be possible to evaluate with simple metrics the interface and functionalities of software for browsing, mapping, summarization, and navigation, it is clearly the case that the questions and activities supported by such software are intertwined with a many large, diverse, complicated, and probably conflicting goals. For instance, today it is the case that the design, implementation, marketing and distribution of browser software is a multi-billion dollar business (witness Netscape or the recent court battles between the Justice Department and Microsoft). Nevertheless, no simple metric of what makes a better browser underlies the tallying of points in the current browser wars.

Part of the philosophical work of this thesis has been to argue, firstly, that software aimed at improving VLSC should improve what is absolutely essentially to any good conversation, discussion, or debate; i.e., the exercise of free speech. Secondly, it was argued that, given the right perspective, it is possible to see how the technical vocabulary of software design and computation is concordant or sometimes in conflict with the philosophical frameworks that support free speech (e.g., the language of the U.S. constitution). By recognizing free speech as a top-level goal for VLSC software design and by investigating how the technical vocabulary of computation and software meets or misses the philosophical analysis and articulation of free speech, it is possible to begin to outline an ethics of software design for VLSC. While this sort of philosophical work will not yield a simple metric to better measure browsers, it is essential work if browsing and navigation software is to play a significant role in how we orient and position ourselves in the larger, electronic worlds of the net.

A similar combination of philosophical and technical work has to be done to reason about the design aesthetics of VLSC software. Especially in a setting like a VLSC which -- at least as it is now practiced on the Internet in, for instance, Usenet newsgroups – is usually an international exchange involving people from many cultures, and so designers cannot assume that everyone in the group of potential, software operators (or "users") will possess a common knowledge, commonsense, or common set of intuitions. In a culturally diverse environment, one cannot presume a given commonsense or a single sense of what is intuitive and what is not. Instead, software design must provide for the possibility of new intuitions and the production of new commonsense knowledge, new common grounds, and novel metaphors.

This is a particularly difficult provision to design for. On one hand, designers know that commonsense metaphors – like the so-called "desktop metaphor" used in the design of interfaces for computer operating systems – can be extremely useful for designing "user friendly" software. On the other hand, what might be intuitive or commonsensical to a member of one culture might be

bizarre, frightening, or just strange to a member of another. Because so much software design has been done within a single culture or a hegemonic culture, most contemporary software design vocabularies for describing intuitive interfaces and good design are either just wrong or overly constrained for the articulation of an aesthetics of software design for domains like VLSC. While work for this thesis includes one instance of a possible alternative – i.e., the implemented interface for the Conversation Map system – this alternative is only one of many possibilities. Much of the philosophical work done in this thesis towards a design aesthetics for VLSC is a review and critique of the inadequacies of current aesthetics. This critique is a necessary, preliminary step to any large progress in this domain.

By coining a new name for the emerging domain of VLSC and then by showing some of its dimensions in existing, everyday practices of, for instance, Usenet newsgroup email exchanges, this thesis reveals some of the inadequacies of today's net-based communications and communities and, simultaneously, provides a means to imagine and design for a better future. Today it is possible to design for an individual client, a larger, corporate client, or even a nation (if one is designing, for instance, stamps or money or even a constitution). But, if VLSC is bringing people together in new, unprecedented relationships and configurations, one can imagine a day when it will be possible to design for a VLSC, just as it is possible to design for other sorts of groups today.

Much of the bulk of this thesis has been spent in forging and then bringing together the appropriate theory and software to establish the existence of VLSC as a legitimate domain of inquiry and interaction. In the near future, this work will be extended through a participatory design practice in collaboration with VLSC participants and observers.

Future Work

The toolkit assembled for the Conversation Map system is arguably an interesting assemblage, but clearly just one of many possibilities. It should be possible to disassemble the underlying text analysis machinery and interface components of the Conversation Map system into a sort of "construction kit" that would allow different people to see different aspects of a conversation with it. For example, in English many nouns get "verbed." I.e., over time some nouns become acceptable verbs. One might, therefore, want to see the system produce a semantic network for nouns (as it now does), but also one for verbs (which it could do with some very small changes to the underlying code).

A second example of the sorts of options or adjustable parameters that could be built into the system concerns the calculation of discussion themes. Some conversations are very tightly reasoned and so the sorts of inference necessary to following the thread of the conversation is quite controlled. But, other conversations are more open-ended and accepting of tangential contributions.

Consequently, it might be useful to have an adjustable "lever" to control how expansive the system should be in attempting to connect two messages together by a theme of discussion.

Given the discussion of an epistemology of VLSC presented in the last chapter of this thesis it is possible to begin to imagine how one might experiment with a mutually recursive definition of the semantics used in a conversation and the social cohesion established by the conversation. Over the course of a discussion through an aggregation of language use, it is possible to see – even with the current Conversation Map system – how a certain term comes to have a meaning quite similar to another term which may, initially, have had a completely unrelated set of connotations. E.g., in American slang, the adjective "cool" acquired the meaning of what a previous generation might have called "hot." Similarly, the word "bad" came to mean "good," etc. One might imagine a slight modification of the Conversation Map system that works through a series of weeks, months or years of a given VLSC using the thesaurus generated from the previous time period as the input thesaurus (which in the current system is just always the WordNet thesaurus) used in the computation of discussion themes and social cohesion of the next time period. With the output of a machine like this one could then start to ask whether, for example, the term "cool" today plays the same kind of role in establishing social relations for the group as the term "hot" did yesterday? And, simultaneously, one might ask, does the establishment of a particular configuration of social relationships between participants have an effect on which terms end up as the pivotal ones for the group? What about the configuration of successive generations of a group makes synonyms out of antonyms?

By redesigning the Conversation Map system so that it can be taken apart and reconfigured, or even taken apart and used in part or whole within another piece of software (e.g., within a search engine to help one find the newsgroup with a desired social configuration and lexicon), a participatory design practice can be pursued that allows non-software-designer collaborators to be tool builders. This approach – which might be called an open-source tactical approach to design – will eventually be what it means to navigate through and do design for VLSCs. This approach will be a convergence of design, navigation, self-governance, and free speech.

REFERENCES

Robert P. Abelson and J.D. Carroll, "Computer Simulation of Individual Belief Systems," *American Behavior Scientist*, 8 (1965): 24-30.

Philip E. Agre. *Computation and Human Experience* (New York, NY: Cambridge University Press, 1997).

Phil Agre. "Growing a Democratic Culture: John Commons on the Wiring of Civil Society" presented at *Media in Transition* (Cambridge, MA: MIT, October 1999). See http://dlis.gseis.ucla.edu/people/pagre/commons.html

Philip E. Agre and David Chapman, "Pengi: An Implementation of a Theory of Activity," in *Proceedings of the Fifth National Conference on Artificial Intelligence* (Seattle, WA: Morgan Kaufmann, 1987), 268-272.

James F. Allen, B. Miller, E. Ringger and T. Sikorski, "A Robust System for Natural Spoken Dialogue," in the *Proceedings of the 34th Meeting of the Association for Computational Linguistics*, 1996.

Louis Althusser, "Ideology and Ideological State Apparatuses (Notes towards an Investigation)," in *Mapping Ideology*, edited by Slavoj Zizek (New York: Verso, 1994).

Benedict Anderson, *Imagined communities: reflections on the origin and spread of nationalism* (London: Verso, 1983).

Amit Bagga and Breck Baldwin "Entity-Based Cross-Document Coreferencing Using the Vector Space Model" In *Proceedings of ACL-COLING'98* (Montreal, Canada, June 1998).

Amit Bagga, Breck Baldwin, and Sara Shelton (editors) *Coreference and Its Applications, Proceedings of the Workshop*, New Brunswick, NJ: Association for Computational Linguistics, 22 June 1999.

Jeremy Bentham, *The Works of Jeremy Bentham* (London: W. Tait, Simpkin, Marshall, 1843).

Tim Berners-Lee. *Weaving the Web: The original design and ultimate destiny of the World Wide Web by its inventor* (San Francisco: HarperSanFrancisco, 1999).

Michael Berthold, Fay Sudweeks, Sid Newton, Richard Coyne. "It makes sense: Using an autoassociative neural network to explore typicality in computer mediated discussions" In *Network and Netplay: Virtual Groups on the Internet,* edited by F. Sudweeks, M. McLaughlin, and S. Rafaeli (Cambridge, MA: AAAI/MIT Press, 1998).

Michael L. Best. "Corporal ecologies and population fitness on the net." *Journal of Artificial Life*, 3(4), 1998.

Margaret A. Boden, *Artificial intelligence and natural man* (New York: Basic Books, 1977).

A.H. Bond and Les Gasser (eds.), *Readings in Distributed Artificial Intelligence* (Los Altos, CA: Morgan Kaufmann Publishers, 1988).

Bertolt Brecht, *Brecht on theatre: The development of an aesthetic*, edited and translated by John Willett (New York: Hill and Wang, 1964).

Andrew Brook, *Kant and the Mind* (Cambridge, UK: Cambridge University Press, 1994).

Rodney Brooks, "Intelligence Without Representation," *Artificial Intelligence* 47 (1991): 139-160.

Rodney Brooks and Pattie Maes (eds.) *Artificial life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems* (Cambridge, MA: MIT Press, 1994).

Benjamin Buchloh, "From Faktura to Factography," *October*, 30 (Fall 1984).

Michel Callon, John Law, Arie Rip (editors) *Mapping the Dynamics of Science: Sociology in the Real World* (London: Macmillan Press, Ltd., 1986).

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

Jaime G. Carbonell, *Subjective Understanding: Computer Models of Belief Systems*, Ph.D. Thesis (New Haven, CT: Yale University, 1979).

Jaime G. Carbonell, "Towards a Process Model of Human Personality Traits," *Artificial Intelligence* 15 (1980): 49-74.

Lewis Carroll, "Alice's Adventures in Wonderland," in *The Annotated Alice* (New York: Clarckson N. Potter, Inc., Pub, 1960).

Howard Caygill, *A Kant dictionary* (Cambridge, MA, Blackwell Reference, 1995), 388.

Michel De Certeau, *The Practice of Everyday Life*, translated by Steven Rendall (Berkeley, CA: University of California Press, 1984).

Eugene Charniak, *Statistical Language Learning* (Cambridge, MA: MIT Press, 1993).

Alonzo Church, "The Calculi of Lambda-Conversation," *Annals of Mathematical Studies, Vol. 6*, Princeton University Press, 1941.

William J. Clancey. *Situated cognition: On human knowledge and computer representations* (New York, NY: Cambridge University Press, 1997).

W. Clinger, J. Rees, H. Abelson, R.K. Dybvig, C.T. Haynes, G.J. Rozas, N.I. Adams IV, D.P. Friedman, E. Kohlbecker, G.L. Steele Jr., D.H. Bartley, R. Halstead, D. Oxley, G.J. Sussman, G. Brooks, C. Hanson, K.M. Pitman, M. Wand, *Revised(4) Report on the Algorithmic Language Scheme*, http://www-swiss.ai.mit.edu/~jaffer/r4rs_toc.html

Kenneth Mark Colby, "Modeling a Paranoid Mind," *Behavioral and Brain Sciences*, 4 (1981): 515-534.

A.M. Collins and E.F. Loftus "A Spreading Activation Theory of Semantic Processing," *Psychological Review, 82*, 1975: 407-428.

Richard Coyne, *Designing Information Technology in the Postmodern Age* (Cambridge, MA: MIT Press, 1995).

Simon Critchley, *The Ethics of Deconstruction: Derrida and Levinas* (Cambridge, MA: Blackwell, 1992).

Manuel DeLanda *War in the Age of Intelligent Machines* (New York: Zone Books, 1991).

Gilles Deleuze, *Kant's Critical Philosophy: The Doctrine of the Faculties*, translated by Hugh Tomlinson and Barbara Habberjam (Minneapolis, MN: University of Minnesota Press, 1984).

Gilles Deleuze, *Cinema 1: The movement-image*, translated by Hugh Tomlinson and Barbara Habberjam (Minneapolis: University of Minnesota Press, 1986).

Gilles Deleuze, *Cinema 2: The time-image*, translated by Hugh Tomlinson and Barbara Habberjam (Minneapolis: University of Minnesota Press, 1986).

Gilles Deleuze *Logic of Sense,* translated by Mark Lester and edited by Constantin V. Boundas (New York: Columbia University Press, 1990).

Gilles Deleuze and Félix Guattari, *Anti-Oedipus: capitalism and schizophrenia*, translated by Robert Hurley, Mark Seem, and Helen R. Lane (Minneapolis: University of Minnesota Press, 1983).

Gilles Deleuze and Félix Guartarri, *A thousand plateaus: capitalism and schizophrenia*, translated by Brian Massumi (London: Athlone Press, 1988).

Jacques Derrida, *Positions* (Chicago: University of Chicago Press, 1981).

Jacques Derrida, "Signature Event Context," in *Limited Inc* (Evanston, IL: Northwestern University Press, 1988).

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

Christine Doran, *Incorporating Punctuation into the Sentence Grammar: A Lexicalized Tree Adjoining Grammar Perspective*, Ph.D. Dissertation and Technical Report IRCS-98-24 (Philadelphia, PA: University of Pennsylvania, 1998).

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

Claudia Dreifus, "A Conversation with Dr. Marvin Minsky: Why Isn't Artificial Intelligence More Like the Real Thing?," *New York Times*, July 28, 1998; or see http://www.nytimes.com/library/national/science/072898sci-minsky.html

Hubert L. Dreyfus, *What computers can't do; a critique of artificial reason. 1st edition* (New York: Harper and Row, 1972).

Hubert L. Dreyfus, *What Computers Still Can't Do* (Cambridge, MA: MIT Press, 1992).

Hubert L. Dreyfus with Harrison Hall (eds.) *Husserl, intentionality, and cognitive science* (Cambridge, MA: MIT Press, 1982).

Joseph Dumit, "Symptomatic, Ill and Structurally Damned: Notes on Liminal Creativity and Social Movements," presented at the *American Anthropological Association*, Philadelphia, PA, 1998. See http://symptom.mit.edu/information/Dumit-Liminality.PDF.

Joseph Dumit *Whose Brain Is This? PET Scans and Personhood in Biomedical America* (forthcoming; see http://symptom.mit.edu).

Paul Edwards (editor-in-chief) "Dialectics," in *Encyclopedia of Philosophy, Volume 2* (New York: The Macmillan Company and The Free Press, 1967).

Peter Eisenman "Diagram: An Original Scene of Writing" in *Architecture New York* (23:27, 1998).

Sergei Eisenstein, *Film form; essays in film theory*, ed. and tr. by Jay Leyda (New York, Harcourt, Brace, 1949).

Thomas Erickson, David N. Smith, Wendy A. Kellogg, Mark Laff, John T. Richards, Erin Bradner "Socially Translucent Systems: Social Proxies, Persistent Conversation, and the Design of 'Babble'" in *Proceedings of CHI'99* (Pittsburg, PA, May 1999).

Victor Erlich, *Russian formalism: history, doctrine, 3rd edition* (New Haven: Yale University Press, 1965).

Ronald Fagin, J.Y. Halpern, Y. Moses and M.Y. Vardi, *Reasoning About Knowledge* (Cambridge, MA: MIT Press, 1995).

Edward A. Feigenbaum and Pamela McCorduck, *The fifth generation: artificial intelligence and Japan's computer challenge to the world* (Reading, MA: Addison-Wesley, 1983).

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

Thomas Flynn, "Foucault as Parrhesiast: His Last Course at the Collège de France (1984)" in The Final Foucault, edited by James Bernauer and David Rasmussen (Cambridge, MA: MIT Press, 1987).

Michel Foucault, *Discipline and Punish: The Birth of the Prison*, translated by Alan Sheridan (New York: Vintage Books, 1979).

Michel Foucault, "Parrhesia and Community Life," In *Discourse and truth: the problematization of parrhesia*, edited by Joseph Pearson, Six lectures given at the University of California at Berkeley, Oct-Nov. 1983. The text was compiled from tape-recordings made of the lectures delivered, in English, by Michel Foucault at the University of California at Berkeley in the Fall Term of 1983, transcribed and edited in 1985. The lectures were given as part of Foucault's seminar, entitled "Discourse and Truth." Since Foucault did not write, correct, or edit any part of the text transcripts, they lack his imprimatur and do not present his own lecture notes. The transcripts are the notes of one of his auditors and are online here: http://perso.club-internet.fr/kmille/discourse.htm

Michel Foucault, "Governmentality," in *The Foucault Effect: Studies in Governmentality*, edited by Graham Burchell, Colin Gordon and Peter Miller (Chicago: University of Chicago Press, 1991).

Michel Foucault, *Ethics: Subjectivity and Truth; Essential Works of Foucault 1954-1984, Volume I,* edited by Paul Rabinow (New York: The New Press, 1997).

Roger Fowler, *Linguistic Criticism, Second Edition* (New York: Oxford University Press, 1996).

Sigmund Freud, "The 'Uncanny'," in *The Standard Edition of the Complete Psychological Works of Sigmund Freud, Volume XVII*, translated by James Strachey. (London: Hogarth Press, 1919).

Sigmund Freud, *Jokes and their relation to the unconscious*, translated and edited by James Strachey (New York: Norton, 1960).

Oscar H. Gandy, Jr., *The Panoptic Sort: A Political Economy of Personal Information* (San Francisco, CA: Westview Press, 1993).

Howard Gardner, *The Mind's New Science: A History of the Cognitive Revolution* (New York: Basic Books, 1985).

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

Harold Garfinkel, *Studies in Ethnomethodology* (Cambridge, MA: Blackwell Publishers Inc., 1967).

Les Gasser, "Social Conceptions of Knowledge and Action: Distributed Artificial Intelligence and Open Systems Semantics," *Artificial Intelligence* 47 (1991): 107-138.

Judith Genova, "Turing's Sexual Guessing Game" *Social Epistemology*, 8(4) (1994): 313-326.

James Gibson, *The Ecological Approach to Visual Perception* (New York: Houghton Mifflin, 1979).

Antonio Gramsci. *Selections from the Prison Notebooks* (London: Lawrence and Wishart, 1971).

Stephen Green, *Automatically generating hypertext by computing semantic similarity*, Ph.D. thesis (Toronto, Canada: Department of Computer Science, University of Toronto, August 1997).

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

A.J. Greimas, *On Meaning: Selected Writings in Semiotic Theory*, translated by P.J. Perron and F.H. Collins (Minneapolis: University of Minnesota Press, 1987).

Barbara Grosz and Candace Sidner, "Attention, intentions, and the structure of discourse," *Journal of Computational Linguistics* 12 (3) (1986): 175-204.

Elizabeth A. Grosz, *Jacques Lacan: a feminist introduction* (New York: Routledge, 1990).

Félix Guattari, "Pour une éthique des médias," *Le Monde*, Mecredi 6 novembre 1991.

Félix Guattari, *Chaosmosis: An Ethico-Aesthetic Paradigm*, translated by Paul Bains and Julian Pefanis (Bloomington, IN: Indiana University Press, 1995).

John J. Gumperz, Gurinder Aulakh, and Hannah Kaltman, "Thematic structure and progression in discourse," In *Language and Social Identity*, edited by John J. Gumperz (Cambridge, UK: Cambridge University Press, 1982).

Ken Haase, "Analogy in the Large," in *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-95)*, 1995

Stuart Hall, "The rediscovery of 'ideology': return of the repressed in media studies," in *Culture, Society, and the Media*, edited by Michael Gurevitch, Tony Bennett, James Curran, and Janet Woollacott (New York: Routledge, 1982).

Michael A. K. Halliday. *An Introduction to Functional Grammar, Second Edition* (London: Edward Arnold, 1994).

Michael A.K. Halliday and Ruqaiya Hasan *Cohesion in English* (New York: Longman, 1976).

Sanda Harabagiu. "WordNet-Based Inference of Textual Cohesion and Coherence" in the *Proceedings of FLAIRS-98* (Sanibel Island, FL, May 1998).

Sanda Harabagiu and Dan Moldovan, "Enriching the WordNet Taxonomy with Contextual Knowledge Acquired from Text," in *Natural Language Processing and Knowledge Representation: Language for Knowledge and Knowledge for Language*, edited by S. Shapiro and L. Iwanska (Cambridge, MA: AAAI/MIT Press, 1999).

Zelig Harris, "Discourse Analysis," *Language*, 28 (1952): 1-30 and 474-494.

Patrick Hayes and Kenneth Ford, "Turing Test Considered Harmful," in *Proceedings of the International Joint Conference on Artificial Intelligence* (Los Altos, CA: Morgan Kaufmann Publishers, 1995), 972-977 (976).

Marti Hearst, "Automated Discovery of WordNet Relations," in *WordNet: An Electronic Lexical Database*, edited by Christiane Fellbaum (Cambridge, MA: MIT Press, 1998).

Martin Heidegger, *The question concerning technology, and other essays*, translated by William Lovitt (New York: Garland Pub., 1977).

Steve Heims, *The Cybernetics Group* (Cambridge, MA: MIT Press, 1991).

James Herrick, "Empowerment Practice and Social Change: The Place for New Social Movement Theory," presented at *The New Social Movement and Community Organizing Conference*, University of Washington, Seattle, WA. November 1-3, 1995. See also http://www.interweb-tech.com/nsmnet/herrick.htm

Susan Herring, Deborah A. Johnson, Tamra DiBenedetto. "'This discussion is going too far!': Male resistance to female participation on the Internet," in *Gender Articulated: Language and the Socially Constructed Self* edited by K. Hall and M. Bucholtz (New York: Routledge, 1995).

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).

Andrew Hodges, *Alan Turing: The Enigma* (New York: Simon and Schuster, 1983).

Ernst Theodor Amadeus Hoffmann, "The Sandman," in *Hoffmann's strange stories*, translated by L. Burnham (Boston: Burnham Brothers, 1855).

Patrick H. Hutton, "Foucault, Freud, and the Technologies of the Self," in *Technologies of the Self: A Seminar with Michel Foucault*, edited by Luther H. Martin, Huck Gutman and Patrick H. Hutton (Amherst, MA: University of Massachusetts Press, 1988).

Eugene Ionesco, *The bald soprano and other plays*, translated by Donald M. Allen (New York: Grove Press, 1958).

Hitoshi Isahara and Hiromi Ozaku. "Intelligent Network News Reader," in *Proceedings of IUI'97* (Orlando, FL: Association for Computing Machinery, 1997).

Roman Jakobsen, *Verbal art, verbal sign, verbal time*, edited by Krystyna Pomorska and Stephen Rudy (Oxford, UK: Blackwell, 1985).

Fredric Jameson, *The Prison-House of Language: a critical account of structuralism and Russian formalism* (Princeton, NJ: Princeton University Press, 1972).

Karen Sparck Jones, "Automatic summarizing: factors and directions," in *Advances in Automatic Text Summarization*, edited by Inderjeet Mani and Mark T. Maybury (Cambridge, MA: MIT Press, 1999).

Immanuel Kant, *Kritik der Urtheilskraft, in Immanuel Kant's gesammelte Schriften, Volume V* (Berlin: Bey Lagarde und Friederich ,1790), 165-485.

Daniel Karp, Yves Schabes, Martin Zaidel, and Dania Egedi. "A Freely Available Wide Coverage Morphological Analyzer for English" in *Proceedings of COLING-92*, (1992).

Julia Kristeva, *Strangers to Ourselves*, translated by Leon S. Roudiez (New York: Columbia University Press, 1991).

Morten Kyng, "Designing for Cooperation: Cooperating in Design," *Communications of the ACM*, 34(12), 1991, 65-73.

George Lakoff and Mark Johnson. *Metaphors We Live By* (Chicago: University of Chicago Press, 1980).

Bruno Latour and Geneviève Teil "The Hume Machine: Can association networks do more than formal rules" *Stanford Humanities Review (special issue on artificial intelligence)* 4.2 (1995): 47-65.

Brenda Laurel, *Computers as Theater* (Reading, MA: Addison-Wesley Publishing Company, 1991).

John Law and John Hassard (editors) *Actor network theory and after* (Oxford, UK; Malden, MA: Blackwell/The Sociological Review, 1999).

Wendy Lehnert, Claire Cardie, and Ellen Riloff. "Analyzing research papers using citation sentences. In *Proceedings of the 12th Annual Conference on Cognitive Science*, 1990.

Douglas Lenat and R. Guha, *Building large knowledge-based systems: representation and inference in the Cyc project* (Reading, MA: Addison-Wesley, 1990).

Andrew Leonard, *Bots: the origin of new species* (San Francisco: Hardwired, 1997).

Larry Lessig, "Open Code and Open Societies: Values of Internet Governance (Draft 3)," presented as the *1999 Sibley Lecture* (Athens, GA: University of Georgia, February 1999). See http://cyber.law.harvard.edu/works/lessig/kent.pdf

Diane J. Litman. "Cue Phrase Classification Using Machine Learning," *Journal of Artificial Intelligence Research*, 5:53-94, 1996.

Claude Lévi-Strauss. *The Savage Mind* (Chicago: University of Chicago Press, 1966).

Christina Lodder, *Russian Constructivism* (New Haven: Yale University Press, 1983).

Jean-François Lyotard, *The Inhuman: Reflections on Time*, translated by Geoffrey Bennington and Rachel Bowlby (Stanford, CA: Stanford University Press, 1991).

Jean-François Lyotard, *Lessons on the analytic of the sublime: Kant's Critique of Judgment*, translated by Elizabeth Rottenberg (Stanford, CA: Stanford University Press, 1994).

D. Magerman, *Natural language parsing as statistical pattern matching*, Ph.D. thesis (Stanford, CA: Stanford University, 1994).

Kirsten Malmkjaer "Discourse and conversational analysis" In Kirsten Malmkjaer (editor) *The Linguistics Encyclopedia* (New York: Routledge, 1991).

Inderjeet Mani and Mark T. Maybury (editors) *Advances in Automatic Text Summarization* (Cambridge, MA: MIT Press, 1999).

Daniel Marcu, *The Rhetorical Parsing, Summarization, and Generation of Natural Language Texts*, Ph.D. Thesis (Toronto: Department of Computer Science, University of Toronto, December 1997).

Daniel Marcu. "Improving summarization through rhetorical parsing tuning" in the *Sixth Workshop on Very Large Corpora* (Montreal, Canada, August 1998).

J. Martins, "Belief Revision" in *Encyclopedia of Artificial Intelligence, Second Edition,* editor-in-chief Stuart C. Shapiro (New York: Wiley, 1992).

Karl Marx, "Critique of Hegel's Dialectic and General Philosophy," in *Karl Marx: Early Writings* (New York: McGraw-Hill, 1983), 379-400.

Brian Massumi "The Autonomy of Affect" *Cultural Critique* (The Politics of Systems and Environments, Part II), 31 (Fall 1995).

John McCarthy, *Formalizing Common Sense: Papers by John McCarthy*, edited by Vladimir Lifschitz and John McCarthy (Oxford, UK: Intellect, 1998).

James L. McClelland and David E. Rumelhart (eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition; Volumes 1 and 2* (Cambridge, MA: MIT Press, 1986).

Pamela McCorduck, *Machines who think: a personal inquiry into the history and prospects of artificial intelligence* (San Francisco: W.H. Freeman, 1979).

A McHoul, "Discourse," In R.E. Asher (editor-in-chief) *The Encyclopedia of Language and Linguistics* (New York: Pergamon Press, 1994).

Alberto Melluci, "The Process of Collective Identity," in *Social movements and culture*, edited by Hank Johnston and Bert Klandermans (Minneapolis: University of Minnesota Press, 1995).

Christian Metz, *Psychoanalysis and Cinema* (London: Macmillan, 1983).

Marvin Minsky, "A Framework for Representing Knowledge" in *Mind Design: Philosophy, Psychology, Artificial Intelligence*, ed. by John Haugeland (Cambridge, MA: MIT Press, 1981).

Marvin Minsky, *The Society of Mind* (New York: Simon and Schuster, 1986).

Arnaldo Momigliano. "Freedom of Speech in Antiquity" In *Dictionary of the History of Ideas, Volume 2*, edited by Philip P. Wiener (New York: Scribner's Sons, 1973).

J. Moore and C. Paris, "Planning texts for advisory dialogues: Capturing intentional and rhetorical information," *Computational Linguistics 19*, 4 (1993), 651-694.

Nicholas Negroponte, *Being Digital* (New York: Knopf, 1995).

Otto Neumaier, "A Wittgensteinian View of Artificial Intelligence" in *Artificial Intelligence: The Case Against*, ed. by R. Born (London: Croom-Helm, 1987).

Alan Newell, J.C. Shaw, and Herbert A. Simon, "GPS, A Program That Simulates Human Thought" in *Computers and Thought*, ed. by Edward A. Feigenbaum and Julian Feldman (New York: McGraw-Hill, 1963), 279-293.

Virginia Nightingale, *Studying Audiences: The Shock of the Real* (New York, Routledge, 1996).

Peter Norvig, *Paradigms of Artificial Intelligence Programming: Case Studies in Common Lisp* (San Mateo, CA: Morgan Kaufmann Publishers, 1992).

Talcott Parsons, *Social Systems and the Evolution of Action Theory* (New York: Free Press, 1977).

Michel Pêcheux, *Automatic discourse analysis,* edited by Tony Hak and Niels Helsloot and translated by David Macey (Amsterdam; Atlanta, GA: Rodopi, 1995).

Charles Sanders Peirce, *Collected papers of Charles Sanders Peirce*, edited by Charles Hartshorne and Paul Weiss (Cambridge: Harvard University Press, 1931-1958).

Alastair Pennycock, "Incommensurable Discourses," *Journal of Applied Linguistics,* Volume 15, Number 2, 1994.

M.R. Quillian "Semantic Memory" In M. Minsky (editor) *Semantic Information Processing* (Cambridge, MA: MIT Press, 1968).

John Rajchman *Truth and Eros: Foucault, Lacan, and the Question of Ethics* (New York: Routledge, 1991).

John Rajchman, *Constructions* (Cambridge, MA: MIT Press, 1998).

Otto Rank, *The Double: A Psychoanalytic Study*, translated by Harry Tucker, Jr. (Chapel Hill, NC: University of North Carolina Press, 1971).

Earl Rennison, "Galaxies of News: An Approach to Visualizing and Understanding Expansive News Landscapes," in *Proceedings of UIST'94* (1994).

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

Phil Resnik and Judith L. Klavans (editors), *The balancing act: combining symbolic and statistical approaches to language* (Cambridge, MA: MIT Press, 1996).

Jeffrey C. Reynar and Adwait Ratnaparkhi. "A Maximum Entropy Approach to Identifying Sentence Boundaries." In *Proceedings of the Fifth Conference on Applied Natural Language Processing* (Washington, D.C., March 31-April 3, 1997).

D.M. Rind, I.S. Kohane, P. Szolovits, C. Safran, H.C. Chueh and G.O. Barnett, "Maintaining the Confidentiality of Medical Records Shared over the Internet and World Wide Web," *Annals in Internal Medicine* 1997;127(2):138-141.

R. Bruce Roberts and Ira P. Goldstein, "The FRL manual," (Cambridge, MA: Massachusetts Institute of Technology, 1977).

Carl Rogers, *On Becoming a Person* (Boston: Houghton Mifflin, 1961).

Warren Sack, "On the Computation of Point of View," in *Proceedings of the National Conference of Artificial Intelligence* (AAAI 94) (Seattle, WA: Morgan Kaufmann Publishers, July 31-August 4, 1994).

Warren Sack, "Painting Theory Machines," *Art and Design*, 48 (May 1996): 80-92.

Warren Sack, "Artificial Human Nature" *Design Issues*, 13 (Summer 1997), 55-64.

Warren Sack, "Artificial Intelligence and Aesthetics," in *The Encyclopedia of Aesthetics, Volume 1*, editor-in-chief Michael Kelly, editor-in-chief (New York: Oxford University Press, 1998).

Warren Sack "Diagrams of Social Cohesion," in *Descriptions of Demonstrated Systems, ACL'99* (University of Maryland, College Park: Association for Computational Linguistics, June 1999).

Warren Sack "Stories and Social Networks" in the *Proceedings of the American Association of Artificial Intelligence Workshop on Narrative Intelligence* (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: Association for Computing Machinery, January 2000)

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 2000).

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

Harvey Sacks, *Lectures on Conversation, Volumes 1 and 2*, edited by Gail Jefferson (Cambridge, MA: Blackwell, 1992).

Gerald Salton, James Allan, Chris Buckley, Amit Singhal. "Automatic Analysis, Theme Generation, and Summarization of Machine-Readable Texts," *Science*, Vol. 264, (3) (June 1994), 1421-1426.

Roger C. Schank and Robert P. Abelson, *Scripts, plans, goals and understanding: An inquiry into human knowledge structures* (New York: John Wiley and Sons, 1977).

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

Ben Schneiderman, *Designing the User Interface: Strategies for Effective Human-Computer Interaction* (Reading, MA: Addison-Wesley, 1987).

Stuart C. Shapiro (editor-in-chief) *Encyclopedia of Artificial Intelligence, Second Edition* (New York: Wiley, 1992).

Beerud Sheth *NEWT: A Learning Approach to Personalized Information Filtering,* Master's Thesis (Cambridge, MA: MIT Media Laboratory, 1993).

Viktor Shklovsky, "Art as Technique," in *Russian Formalist Criticism*, ed. by L.T. Lemon and M.J. Reis (Lincoln, NE: University of Nebraska Press, 1965).

David Silverman, *Harvey Sacks: Social Science and Conversation Analysis* (New York, Oxford University Press, 1998).

Kaja Silverman, *The Subject of Semiotics* (New York: Oxford University Press, 1983).

Michael Silverstein, "Shifters, Linguistic Categories, and Cultural Description," in *Meaning in Anthropology*, edited by K. H. Basso and H. A. Selby (Albuquerque: University of New Mexico Press, 1976), 11-56.

Herbert A. Simon, *The Sciences of the Artificial, Third Edition* (Cambridge, MA: MIT Press, 1996).

J.McH. Sinclair and M. Coulthard, *Towards an Analysis of Discourse: The English Used by Teachers and Pupils* (London: Oxford University Press, 1975).

B.F. Skinner, *Science and human behavior* (New York: Macmillan, 1953).

Marc Smith. "Netscan: Measuring and Mapping the Social Structure of Usenet" Presented at the *17th Annual International Sunbelt Social Network Conference* (San Diego, CA: February 13-16, 1997).

Dagobert Soergel, *Indexing Languages and Thesauri: Construction and Maintenance* (Los Angeles: Melville Pub. Co., a division of John Wiley, 1974).

Mark A. Stairmand. "Textual context analysis for information retrieval" in the *Proceedings of the 20th annual international ACM SIGIR conference on Research and development in information retrieval*, August 1997.

Lucy Suchman, *Plans and Situated Actions: The Problem of Human-Machine Communication* (New York: Cambridge University Press, 1986), vii-x.

Geneviéve Teil and Bruno Latour, "The Hume Machine: Can Association Networks do more than Formal Rules?," in *Stanford Humanities Review 4.2* (1995): 47-65.

Alan Turing, "Computing Machinery and Intelligence," *Mind*, Volume LIX, No. 236 (1950): 433-460.

Sherry Turkle, *The Second Self: Computers and the Human Spirit* (New York: Simon and Schuster, Inc., 1984).

Sherry Turkle, "Artificial Intelligence and Psychoanalysis: A New Alliance." *Daedalus*, 17(1) Winter 1988.

Bryan S. Turner, "The Micro-Macro Problem," in *The Blackwell Companion to Social Theory*, edited by Bryan S. Turner (Malden, MA: Blackwell, 1996).

Kiyotaka Uchimoto, Hiromi Ozaku, and Hitoshi Isahara. "A Method for Identifying Topic-Changing Articles in Discussion-type Newsgroups within the Intelligent Network News Reader HISHO," in *Proceedings of Natural Language Processing Pacific Rim Symposium* (Phuket, Thailand, December 2-4, 1997).

Victoria Vesna, "Database Aesthetics: Of Containers, Chronofiles, Time Capsules, Xanadu, Alexandria and the World Brain" *Journal of Artificial Intelligence and Society* (Fall 1999).

Anthony Vidler, *The architectural uncanny: essays in the modern unhomely* (Cambridge, MA: MIT Press, 1992).

Larry Wall, Tom Christiansen, and Randal L. Schwartz, *Programming in Perl, Second Edition* (Cambridge, MA: O'Reilly & Associates, Inc., 1996).

Stanley Wasserman and Joseph Galaskiewicz (editors) *Advances in Social Network Analysis: Research in the Social and Behavioral Sciences* (Thousand Oaks, CA: Sage Publications, 1994).

Joseph Weizenbaum, "ELIZA – A Computer Program for the Study of Natural Language Communication between Man and Machine," *Communications of the Association for Computing Machinery* (9), 1966: 36-45.

Barry Wellman, "Living networked in a wired world," *IEEE Intelligent Systems*, (January/February 1999).

Etienne Wenger, *Artificial Intelligence and Tutoring Systems: Computational and Cognitive Approaches to the Communication of Knowledge* (Los Altos, CA: Morgan Kaufmann Publishers, 1987).

Steve Whittaker, Loren Terveen, Will Hill, and Lynn Cherny, "The dynamics of massive interaction," In the *Proceedings of the International Conference on Computer-Supported Cooperative Work* (Seattle, WA: Association for Computing Machinery, 1998), 257-264.

Janyce M. Wiebe, "Tracking point of view in narrative," *Computational Linguistics 20*, 2 (1994), 233-287;

Janyce Wiebe, Graeme Hirst, & Diane Horton, "Language use in context," *Communications of the ACM (Special issue on natural language processing) 39*, 1 (1996), 102-111.

Norbert Wiener, *Cybernetics; or, Control and Communication in the Animal and the Machine* (New York: Wiley, 1955).

Norbert Wiener, *Cybernetics, science, and society: Ethics, aesthetics, literary criticism; Collected Works of Norbert Wiener, Volume 4*, edited by P. Masani (Cambridge, MA: MIT Press, 1985).

Yorick A. Wilks, Brian M. Slator and Louise M Guthrie, *Electric words: dictionaries, computers, and meanings* (Cambridge, MA: MIT Press, 1996).

Terry Winograd and Fernando Flores, *Understanding Computers and Cognition: A New Foundation for Design* (Norwood, NJ: Ablex Pub. Corp., 1986).

Ludwig Wittgenstein, *The Blue and Brown Books: Preliminary Studies for the "Philosophical Investigations"* (New York: Harper Torchbooks, 1958).

Krzysztof Wodiczko, *Critical vehicles: writings, projects, interviews* (Cambridge, MA: MIT Press, 1999).

P. Ziff, *Semantic Analysis* (Ithaca, NY: Cornell University Press, 1960).