

Opening Up Video Databases to Collaborative Commentary

Brian MacWhinney, CMU, macw@mac.com

To appear in: Goldman, R., Pea, R., Barron, B., & Derry, S. (Eds.) Video Research in the Learning Sciences.

ABSTRACT

The TalkBank Project has constructed a web-accessible database for spoken language interactions with transcripts that are linked on the level of the sentence to both audio and video materials. This database includes several large corpora documenting learning in classrooms, tutorial sessions, meetings, and the home. Now that these data are publicly available, we can begin to build tools to facilitate a new process called collaborative commentary. We can define collaborative commentary as the involvement of a research community in the interpretive annotation of electronic records. The goal of this process is the evaluation of competing theoretical claims. The process requires commentators to link their comments and related evidentiary materials to specific segments of either transcripts or electronic media. We will examine current work in the construction of technical methods for facilitating collaborative commentary through browser technology. We will look at seven spoken language database projects that have reached a level of web-based publication that makes them promising as targets for collaborative commentary. For each database, we will consider how collaborative commentary can advance the relevant research agendas.

The TalkBank Project (<http://talkbank.org>) is an international collaborative effort that is building a web-accessible database for spoken language interactions with transcripts linked to both audio and video materials. The project has targeted 12 research areas for development. Of these, one of the most important areas involves the study of learning in classrooms, tutorial sessions, meetings, and the home. For this area, we now have published 12 video corpora linked to transcripts. Now that these data are publicly available, we can begin to build tools to facilitate a new process called collaborative commentary.

We can define collaborative commentary as the involvement of a research community in the interpretive annotation of electronic records. The goal of this process is the evaluation of competing theoretical claims. To achieve clear connections between data and theory, commentators need to link their comments and other evidentiary materials to specific segments of either transcripts or electronic media. This means that, in order to even think about engaging in collaborative commentary, researchers must first have web access to corpora of transcripts of learning sessions linked to video data.

The TalkBank database is designed to provide the raw material that can stimulate and support collaborative commentary. In this paper, I will first discuss the ongoing construction of this database. Then, I will summarize current work in the construction of technical methods for facilitating collaborative commentary on the database. The focus of this analysis will be on the use of video data for the study of learning in the contexts of the class, tutorial sessions, meetings, and the home. However, it is important to understand the extent to which work in these areas is embedded in a large

interdisciplinary context that includes the analysis of video in areas as diverse as field linguistics, emergency medicine, and psychotherapy.

Data Sharing

TalkBank is an interdisciplinary research project funded by the National Science Foundation (BCS-998009). The goal of the project is to support data sharing and direct, community-wide access to naturalistic recordings and transcripts of human and animal communication. The concept emerged from two ongoing initiatives that had already proven important to their respective user communities. The first is the Linguistic Data Consortium (<http://ldc.upenn.edu>) that has published some 288 large corpora over the past decade. The second is the CHILDES system (<http://childes.psy.cmu.edu>) that has constructed a database of 150 corpora of parent-child interactions in 20 languages. The data-sharing model for the new TalkBank project is based on the model from the CHILDES project {MacWhinney, 2000 #9098}.

Having reviewed best practice in 12 very different research areas studying communicative interaction, Talkbank has identified these 7 shared needs:

1. guidelines for ethical sharing of data,
2. metadata and infrastructure for identifying available data,
3. common, well-specified formats for text, audio and video,
4. tools for time aligned transcription and annotation,
5. a common interchange format for annotations,
6. network based infrastructure to support efficient (real time) collaboration, and

7. dissemination of shared data, tools, standards and best practices to the research community.

Six Interest Areas

TalkBank has constructed databases in 12 different areas. Here, we will first review data focusing on learning contexts and then briefly mention databases in five additional related areas.

Learning Contexts

For years, educators have been relying heavily on video data to study classroom discourse. For example, at the 2002 meeting of the American Educational Research Association (AERA), there were 44 scientific panels and symposia that relied on analysis of classroom video. However, virtually none of these data are publicly available, largely because of technical limitations and the failure to link digitized video to transcripts. To demonstrate the use of TalkBank methods for this area, we organized a meeting on classroom discourse in 1999 and then constructed an initial database of materials from story-telling, math education, dyadic tutoring sessions, college lectures, and bilingual classrooms. The materials in this collection include:

1. Studies of problem-based learning in medical school education from Tim Koschmann and Curtis LeBaron.
2. Discussions of the interpretations of graphs from a special issue of the Journal of the Learning Sciences, edited by Anna Sfard and Kay McClain.
3. A college lecture on research methods in psychology from Brian MacWhinney.
4. A lecture on map reading to a 6th grade class contributed by Wolf-Michael Roth.

5. A discussion of a unit on camels in a 5th grade class contributed by Rosalind Horowitz.
6. A comparison of classroom, business, and meeting contexts contributed by Reed Stevens.
7. Discussions of international education contexts contributed by Erica McClave.
8. Dyadic tutorial sessions on the f-ratio in the analysis of variance from Carl Frederiksen.
9. Dyadic tutorial sessions on Psychology Research Methods from Natalie Person and Arthur Graesser.
10. Dyadic tutorial sessions on how to play a video game from Nikolinka Collier.
11. Videos of 7th grade children in Dresden, Germany learning English, French, and Czech from Angelika Kubanek-German.

Seven other video datasets have not yet been added to the database. These include interactions in a children's museum, sessions from bilingual classrooms in Alaska and New York, discussions of gravity and friction, and examinations of the motives of characters in a novel. In addition, the CHILDES database contains several major video studies of more informal learning in the home in English, German, Japanese, Spanish, Cantonese, and Thai.

In an early attempt to promote collaborative commentary, TalkBank promoted the creation of a CD-ROM for a special issue in the Journal of the Learning Sciences {Sfard, 2002 #9610}. This CD-ROM contains articles commenting on two lessons on graphs in a 7th grade classroom. The PDF files for these articles contain links that replay the relevant

video. In addition, there is a demonstration transcript that serves as a compendium of commentary on particular analyses.

We are now working to provide streaming video access to three large longitudinal classroom corpora. One corpus, from Carolyn Maher and her associates at Rutgers, contains 3000 hours recorded over a span of 12 years tracking the math learning of a group of 15 students. Another, from Rich Lehrer and Carmen Curtis, records a year's worth of integrative geometry lessons from a 3rd grade classroom. The third, from Juliet Langman, compares alternative formats for bilingual classrooms.

The next few sections review TalkBank corpora outside of learning contexts. It is useful to review this additional material in order to get a sense of the scope of the project and the nature of material from related areas.

Conversation Analysis

The TalkBank database for adult conversational interaction includes transcripts in CA notation for a subset of the CallFriend phone conversations available through the LDC, the Santa Barbara Corpus of Spoken American English (SBCSAE), recorded phone calls from the Nixon Whitehouse, European political television programs, informal interview materials from a special issue of the Journal of Communication, and a variety of classic materials from the field of Conversation Analysis (CA). Materials in this collection were originally transcribed in a variety of alternative forms of CA format. Recently, we have succeeded in incorporating the full set of CA conventions within the well-defined TalkBank XML Schema. This means that we can not only edit and display CA transcripts, but also provide analysis through programs, both over the web and on the desktop.

Meetings

One subcomponent of the database focuses on conversations during formal meetings. Here, we have two videotapes of dissertation defense reviews and a series of multi-party university work-group meetings, initially transcribed in ISL format. The single largest audio collection in TalkBank includes the oral arguments of the last 30 years of the Supreme Court of the United States (SCOTUS). This enormous dataset is currently being formatted into the TalkBank Schema and linked to the digitized media.

Gesture

There is an increasing interest in corpus-based research on sign language and gesture. With appropriate tools, researchers can create corpora containing grammatical information, discourse structure, facial expression, along with gesture. The resulting corpora can be used to test hypotheses concerning the relationship of the paralinguistic aspects of communication to speech and to meaning. To begin to address this need for a multi-modal corpus, the LDC developed FORM, a non-semantic, geometrically-based annotation scheme. FORM provides a well specified system for coding gestural patterns and kinematics from videos. FORM stores this information using Annotation Graphs (AG), allowing for easy integration of gesture information with other types of communication information. Intensive work with FORM has so far produced 30 minutes of annotations for videos of Brian MacWhinney teaching at CMU. Five other gestural corpora are also available.

Bilingualism and Second Language

Data in this collection include the European Science Foundation (ESF) study of guest workers in seven countries, five code-switching corpora, Chinese immigrants in Hungary, early second language teaching in Dresden, and the Southampton FLLOC database <http://www.flloc.soton.ac.uk> of English learners of French.

Sociolinguistics

Sociolinguistics in the variationist tradition rely on quantitative methods to analyze variation in vernacular speech. TalkBank has contributed data, tools, standards and examples of best practice to enrich corpus creation, analysis and sharing among sociolinguists. These contributions flow directly from TalkBank sponsored research on Annotation Graphs and associated tools and from TalkBank's support of the LDC project on Data and Annotations for Sociolinguistics (DASL). Specific outcomes of this collaboration include: the SLx corpus of classic sociolinguistic interviews from a wide variety of regional and social dialects and speaking styles. SLx surveys more than 150 sociolinguistic variables giving examples of each from the interviews in which they occur.

Ethics for Data-Sharing

Public sharing of data over the web brings with it a variety of challenges regarding participant rights and professional ethics. These issues have been an ongoing topic of discussion within the TalkBank communities. The current result of this process is a set of ethical and practical guidelines adopted for all TalkBank data sets, described at

<http://talkbank.org/share> and available for use beyond the TalkBank project. The centerpiece of this approach is the idea that participants can opt to provide releases for the use of their data at any one of eight different levels. The lowest level of protection allows for full web access to transcripts and video with no attempt at anonymization. Higher levels of protection add anonymization in transcripts and media, password protection, and finally no access but only archiving for the future. The choice of an appropriate level for a given dataset is decided first by the Human Subjects review process at each institution and then by the participants themselves. In addition, TalkBank discourages any use of the data that is critical of the performance or motives of individuals recorded in the interactions. Groups that require further privacy and respect considerations include indigenous groups, speakers of endangered languages, clinical subjects, subjects in psychiatric treatment, and classroom teachers.

Infrastructure for Annotation

Researchers working with video records from learning contexts need support for a wide range of transcription, editing, and analysis functions. The TalkBank project has built several tools to support these functions plus infrastructure to support further development of annotation tools for new disciplines.

One set of tools developed within the AG framework (<http://agtk.sourceforge.net>) {Bird, 2001 #9746} is ideal for datasets in which annotations are fully linked to media. Using the Annotation Graph ToolKit (AGTK) we have developed four useful sample applications: MultiTrans, TableTrans, TreeTrans, and InterTrans {Bird, 2002 #9745}. With TableTrans (Figure 1), we have annotated a corpus of vervet monkey vocalizations

(80 open-reel tapes) for variables such as call type, caller and recipient. TableTrans has also been productively used with sociolinguistic and classroom discourse data.

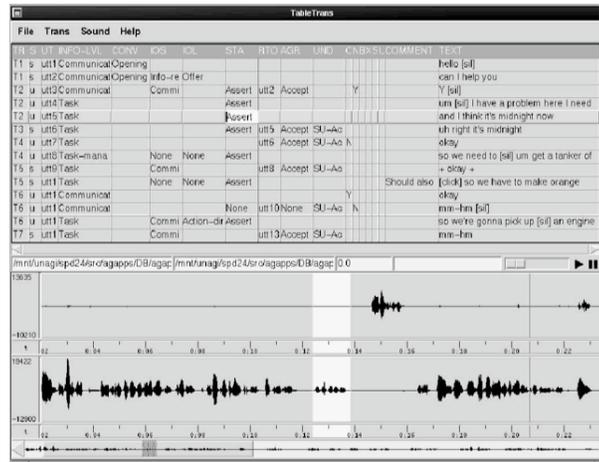


Figure 1: A screen from TableTrans

A second set of tools uses the CHAT format that was originally designed for the transcription of child language corpora. This format provides an extensive set of codes for transcribing all manner of conversational structures, along with a variety of morphosyntactic and discourse coding systems. Files in the CHAT format can be created using the CLAN editor and analyzed by a set of CLAN search and concordance commands. The editor now also supports transcription in Conversation Analysis (CA) format.

TableTrans allows the user to create metadata specific to a file by entering information into dialog boxes. In CLAN, metadata is entered directly into @ID fields in the transcript. The harvesting and publication of metadata is done in conformity with Open Language Archives Community (OLAC: <http://www.linguistlist.org/olac>) standards.

TalkBank has supported the development of several methods for linking transcripts to media, including TransAna (<http://www2.wcer.wisc.edu/Transana/>) for video transcription and Transcriber for audio transcription (<http://www.etca.fr/CTA/gip/Projets/Transcriber>). In addition, the CLAN program supports direct linkage of CHAT files to either audio or video through Transcriber-style linkage, simulation of foot pedal control, numerical editing of time points, or highlighting of segments from a waveform.

Interchange Formats for Data Annotation

The greatest challenge facing TalkBank has been the need to bring hundreds of corpora created in diverse ways into conformity with a common standard. The first step in this direction involved the specification of a proper XML Schema for the CHAT transcription system. The system¹ involves three major steps. In the first step, the ANTLR parser generator creates a parse tree that is converted to a JAXB tree that is then serialized into XML. JAXB is Sun's data binding framework that generates Java code for specialized DOM construction, validation, and serialization. In the second step, XSLT outputs CHAT. In the third step, a modified version of Unix DIFF compares the original CHAT with the converted CHAT and reports differences for correction. Once a corpus passes through this process with no errors, it is included in TalkBank. CHAT versions are zipped so that users can download complete datasets and the XML versions are shipped to the server (<http://xml.talkbank.org>) to support online transcript and media browsing. Getting the CHILDES database through this complete process required a full year of work correcting deviations from the schema; we are now about halfway through the

process of converting all of the other TalkBank materials, including CA transcripts, to XML. The tools we have constructed during this process will be reusable as we continue to reformat additional corpora.

Browsable Transcripts

Once the transcripts are in the TalkBank XML format, they can easily be rendered as HTML pages. We have constructed a Java application called the TalkBank viewer that allows users to view transcripts over the web and play the underlying audio or video by clicking on utterances. After clicking on the first utterance, the user can decide to either stop playback or to allow the viewer to play back the whole transcript by highlighting each utterance as it is played. The transcript can be scrolled and paged to allow for full control of playback over the web. The playback application is installed by clicking on a link at the TalkBank page that then runs Java WebStart.

The playback facility relies on the use of hinted video and audio from Apple's QuickTime Streaming Server. Figure 2 illustrates a typical page for browsing and Figure 3 shows the streaming QuickTime video playback with the transcript echoed below. For audio-only corpora, the QuickTime window displays the utterance and plays the sound without any video.

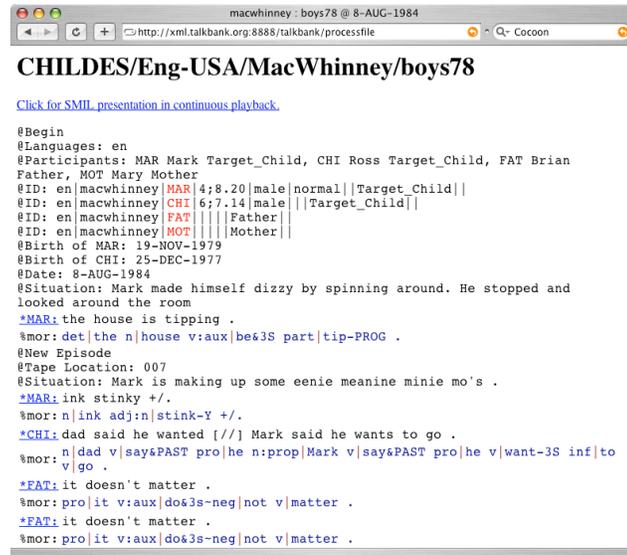


Figure 2: A browsable transcript with hyperlinks



Figure 3: QuickTime playback from a Thai interaction

Collaborative Commentary

Earlier, we defined collaborative commentary² as the involvement of a research community in the interpretive annotation of electronic records. The goal of this process is the evaluation of competing theoretical claims. The process requires commentators to link their comments and related evidentiary materials to specific segments of either transcripts or electronic media. Now that we have access to browsable corpora representing various learning contexts, we can begin to think about how to implement the process of collaborative commentary.

To illustrate the goals of this process, consider the following real-life example. Brian MacWhinney, a researcher in the field of child language, wants to explore evidence for the neo-Vygotskian claim {Nelson, 1998 #9259} that word meanings are shaped through communicative interactions. While browsing through online media at the CHILDES childes.psy.cmu.edu site, he locates several instances of videos of mother-child book reading in the Julie and Rollins corpora. In these interactions, mothers help children turn the pages and name the animals or objects in the pictures. In some cases, children call the pictures by the wrong name. In accord with his theoretical views on the logical problem of language acquisition, MacWhinney {, in press #9536} believes that mothers will use these errors as opportunities to provide corrective positive feedback. For example, if the child calls a bear a “doggie,” the mother should respond, “no, that’s a bear, not a doggie.” MacWhinney would be particularly happy if the child subsequently engaged in self-

correction and said “a bear.” Pursuing this idea, he locates 26 segments in these corpora that are relevant to the position he is advocating.

To initiate the process of collaborative commentary, MacWhinney writes up a short summary of his analysis. He wants to make this analysis available in three ways. First, he wants to post his claim to some discipline-based commentary space on the web. Second, he wants to make sure others who view the relevant segments from the Julie and Rollins corpora are able to see that he has provided detailed interpretive commentary regarding at least 26 specific segments. Finally, he hopes to receive feedback from other researchers regarding his interpretations and arguments.

An Initial Vision

Let us consider for a moment what type of browser display would minimally support MacWhinney’s basic goals. Figure 4 presents an initial framework.



Figure 4: A CHILDES page with commentary links

This figure shows an HTML page displayed dynamically by the CHILDES and TalkBank servers from underlying XML. Thumbnails from the video are used to provide a sense of the nature of the interaction at the relevant point. To the right of the thumbnail is a pointer to the commentary. By clicking on this field, a separate commentary window should open up. This commentary window should allow the research to create the following elements:

1. A brief summary of the claim or analysis relevant to the current utterance or utterance sequence.
2. Typing of the claims and analysis into specific categories.
3. Explanations of the evidentiary role of the texts and media being referenced.
4. Links to other texts or claims that are relevant to the current claim.
5. Links to external web content, including material (HTML, PDF, Word) that presents the proposed analysis more fully.
6. Embedded HTML code.

Once this material has been entered into the commentary field, it can be redisplayed through the TalkBank Viewer³ facility, as illustrated in Figure 5. In this Figure, the QuickTime window echoes the comments that have been entered in the commentary database as continuous streaming playback progresses. It also provides color swatches that code for the nature of the commentary. If a given segment has more attached commentary than can be displayed in the reserved segment of the QuickTime window, the window will have a final line listing the number and size of the comments that cannot be displayed.

```

@Begin
@Languages: en
@Participants: MOT Mother, CHI David Target_Child
@ID: en|rollins|MOT||||Mother||
@ID: en|rollins|CHI|1;8.||||Target_Child||
@Activities: book
*MOT: <ahhah> [ > ] we open it up and there are a set of eyes
%spa: $DHA:YY $DHA:RP
*MOT: <it's a lock and see <book> [ > ] .
%spa: $DHA:ST
*MOT: <ahhah> [ > ] we open it up and there are a set of eyes
%spa: $DJF:ST $DHA:ST
*MOT: <the bear has a baby> [ > ] bottle .
%spa: $DHA:ST
*MOT: <yes # David has baby <bottles> [ > ] .
%spa: $DRP:ST
*MOT: <oh> [ > ] .
%spa: $DHA:MK
*MOT: <there's a mirror> [ > ] .
%spa: $DJF:ST
*MOT: <can David see <David> [ > ] .
%spa: $DHA:RQ
*CHI: 0 .
%act: CHI look-s at rattle in hand then puts rattle in mou

```



Figure 5: Transcript, QuickTime playback, and comments

Possible Implementations

How can we implement this vision? Initially, we might think of using “blogging” systems such as Blogger or mBlog. But these systems have no method for linking comments to lines in dynamic web pages that the CHILDES and TalkBank servers produce through on-the-fly generation of HTML from XML. Nor can they support an organized evidentiary database. A more promising alternative is provided by the W3C Annotea project www.w3.org/2001/Annotea that seeks to provide open-source code for building annotation servers. To explore this option, we set up an Annotea server using Zannot inside Zope, accessed by Annozilla inside Mozilla. Although we were able to publish simply commentaries to TalkBank pages using this method, we found that Annotea client development was not keeping pace with new versions of browsers. We

realized that it might make sense, of course, to take over the task of Annotea development ourselves. However, before moving in this direction, we wanted to survey other alternatives.

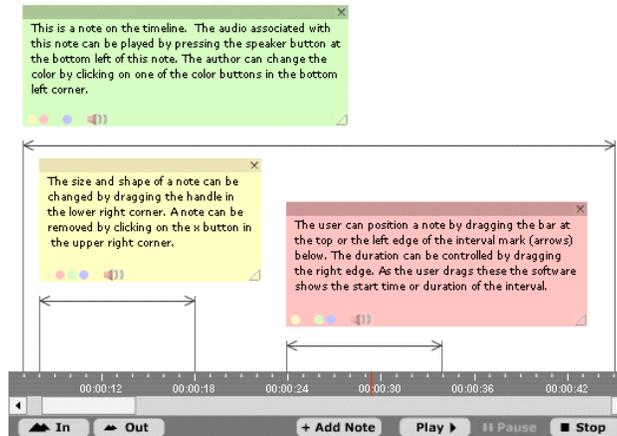


Figure 6: A Project Pad screen

At this point, we learned that Jonathan Smith of the Oyez project at Northwestern University www.oyez.org had built a tool called ProjectPad that showed promise of being able to implement our initial vision. Project Pad (Figure 6) is a program written in Java that controls browsers through Macromedia's Flash. Flash seems particularly well suited for the tight control of commentary entry, item selection, media scrolling, and media linkage that we need to support collaborative commentary. Moreover, the Java code can also interface well with the Java XML database that will store the commentary in terms of organized evidentiary types. Alternatively it could interface with a SQL database controlled by the AG-API {Ma, 2002 #9653}.

Naked Media

Because TalkBank transcripts are subject to ongoing modifications, reference to line numbers is not stable. A more reliable method links commentary to time points, as in the ProjectPad examples in Figure 6. The idea of linking commentary to media is also in accord with the theoretical emphasis in the Annotation Graph framework of Bird and Liberman {, 2001 #9745}. This framework relates all annotations to points in media. The ProjectPad method also opens up a more general possibility for multimedia databases that we will refer to as “naked media.” Consider the case of a large database of classroom video data contributed to TalkBank by Rich Lehrer from the geometry lessons of Carmen Curtis. This database consists of 200 hours of classroom video with no accompanying transcripts. It would take perhaps a full year to transcribe all of these sessions. On the other hand, the video can be prepared for streaming web access in about a month. Once the naked video is posted on the web, it can be target for collaborative commentary through ProjectPad. In cases of this type, collaborative commentary can operate effectively even without accompanying transcripts.

Three examples

To derive a more concrete idea about how collaborative commentary can impact specific research programs, we will survey its application to three specific topics, one from learning contexts and two from other topics that still have a conceptual relation to learning contexts.

Learning Contexts

Consider the following fictive case of how a researcher in this area will use ProjectPad to produce collaborative commentary⁴. Harriet Keck is a developmental psychologist specializing in children's concepts of number. She and her colleague Robin Clark are both interested in understanding how children solve problems such as $3 + 4 = ?$. Keck believes that children solve the problem in an internal mental model and then read out the solution to their fingers. Clark believes that children use their fingers to form external representations of the addends and then count their fingers visually. Keck's model predicts that children will count directly across the fingers, whereas Clark's model suggests that children will begin with placing one addend on each hand separately.

To explore this issue, Dr. Keck uses metadata search tools to find video cases in TalkBank format involving "four-year-old children AND counting". Exploring these videos using the DIVER tool (diver.stanford.edu), she finds that 70% of their gestures support her theory, whereas only 30% are in line with Clark's account.

Over the next several weeks, Keck and her colleagues use ProjectPad to link each case of finger counting to comments that also point to a brief report summarizing her conclusions. Not surprisingly, Clark disagrees with Keck's conclusions and responds by reinterpreting the same video cases that Keck has just analyzed. His analysis points to several counter examples that do not fit Keck's theory. He also argues for including trials that have no overt finger counting in the denominator. Keck, in turn, responds to Clark's criticism by asserting the gestures he has coded are inadvertent hand-movements and revises her paper to anticipate his objection. Keck and students submit their revised paper, including the video data, to the on-line edition of Cognitive Development. One of

the reviewers has a question on whether the authors have properly categorized a set of gestures from one of the videos. Keck responds with a close analysis of the gesture in question using a fine-grained analysis of the actual hand-movements. The reviewer is convinced by her response and the paper is published with links to the video data and analysis.

Although this scenario may seem a bit futuristic, it is not very different conceptually from forms of collaborative commentary we have already produced. One example is a special issue of the Journal of the Learning Sciences that focuses on learning about graphs and numerical distributions in a 7th grade classroom. The difference is that in this new framework, analyses will be directly linked to the data, rather than hidden within PDFs. Moreover, in this new framework, analyses will be directly accessible from browsers.

Gesture

Students of gestural communication are becoming increasingly interested in corpus-based research. With appropriate tools, researchers can create corpora coded for grammatical information, discourse structures, and facial expressions, as well as gestures. The resulting corpora can be used to test hypotheses concerning the relationship of the paralinguistic aspects of communication to speech and to meaning. To begin to address this need for a multi-modal corpus, the LDC TalkBank project developed FORM, a non-semantic, geometrically based annotation scheme that allows an annotator to capture the kinematic information in a gesture just from videos of speakers. FORM stores this gestural information using Annotation Graphs (AG), allowing for easy integration of gesture information with other types of communication information. The work so far has

produced 30 minutes of FORM-annotated videos of Brian MacWhinney teaching at CMU. This corpus has been published through the LDC and a second one is forthcoming. This corpus and others annotated with FORM can now be opened up to collaborative commentary that either promotes alternative coding systems, revises particular categories, or suggests competing interpretations of the meanings and functions of specific gestures.

Because the FORM corpus has been coded using the AG TableTrans application, it supports the method of storage of collaborative annotations developed by Ma, Lee, Bird, and Maeda {, 2002 #9653}. This system allows the user to store annotations from TableTrans in a central ODBC relational database accessed through SQL queries. The system eliminates problems involved in mapping from AG to SQL by precomputing a table of the transitive closure of annotations. This approach may provide useful methods for structuring and accessing the evidential database needed for a fuller system of collaborative commentary.

Oral Argumentation

Another major research area that is currently seeking to engage intensively in collaborative commentary is the field of legal argument analysis. The SCOTUS project at Northwestern is now providing transcripts linked to audio for the past 30 years of oral arguments at the Supreme Court of the United States. Collaborative commentary for this data set will focus initially on a collection of 15 cases in the area of patent law and 15 cases in the area of search-and-seizure. For each area, we will encourage the advocates who argued in the cases to provide commentary on their briefs and answers. We will then organize meetings of legal scholars, commentators, and discourse analysts to provide

evaluations of the shapes of the questioning by the justices and the responses of the advocates.

Conclusion

The TalkBank project has provided openly accessible databases for the study of spoken language interactions. We are now implementing support for collaborative commentary targeted to these databases. Construction of these new methods will open up many exciting new lines of investigation for each of the several disciplines studying human communication.

References

- Bird, S., & Liberman, M. (2001). A formal framework for linguistic annotation. *Speech Communication*, 33, 23-60.
- Bird, S., Maeda, K., Ma, X., Haejoong, L., Randall, B., & Zayat, S. (2002). TableTrans, MultiTrans, InterTrans and TreeTrans: Diverse Tools build on the Annotation Graph Toolkit. In *Proceedings of the Third International Conference on Language Resources and Evaluation* (pp. 1-7). Paris: European Language Resources Association.
- Ma, X., Lee, H., Bird, S., & Maeda, K. (2002). *Model and tools for collaborative annotation*. Paper presented at the Proceedings of the Third International Conference on Language Resources and Evaluation, Paris.
- MacWhinney, B. (2000). *The CHILDES Project: Tools for Analyzing Talk*. Mahwah, NJ: Lawrence Erlbaum Associates.
- MacWhinney, B. (in press). Multiple solutions to the logical problem of language acquisition. *Journal of Child Language*, 29.
- Nelson, K. (1998). *Language in cognitive development: The emergence of the mediated mind*. New York: Cambridge University Press.
- Sfard, A., & McClain, K. (2002). Special Issue: Analyzing tools: Perspective on the role of designed artifacts in mathematics learning. *Journal of the Learning Sciences*, 11, 153-388.

¹ The system described here has been designed and programmed by Franklin Chen.

² The discussion of various functions in collaborative commentary was developed in collaboration with Roy Pea and Ken Rosen at Stanford and Eric Hoffert of Versatility, Inc.

³ The TalkBank Viewer application was designed and programmed by Prabhu Raghunathan

⁴ The Dr. Keck scenario was constructed by Roy Pea.