Issue N° 14 / April 2008
Editor : François Foglia, francois.foglia@idiap.ch

Newsletter

www.amiproject.org
AMI c/o IDIAP Research Institute, Centre du Parc,
Rue Marconi 19, P.O. Box 592, CH-1920 Martigny
info@amiproject.org - www.amiproject.org

Contents

COVER STORY
• Demonstrations presented at the last AMIDA Review Meeting 1

FOCUS
• Demonstrations presented at the last AMIDA Review Meeting (continued) 2
• AMI COI Workshops provide content and context 3

INSIDE AMI
• News and Upcoming Events 4

News

MLMI 2008

5th Joint Workshop on Machine Learning and Multimodal Interaction

8-10 September 2008
Utrecht, The Netherlands

Invited talks

H. Bourlard, S. Renals
"Recognition and Understanding of Meetings - Overview of the European AMI and AMIDA Projects"
LangTech 2008
Rome, 28-29 February 2008

S. Renals, T. Hain, H. Bourlard,
"Recognition and Understanding of Meetings - The AMI and AMIDA Projects"
2007 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU2007)
Kyoto, Japan, December 9-13, 2007

S. Renals, T. Hain, H. Bourlard,
"Interpretations of Multiparty Meetings - The AMI and AMIDA Projects",
Hands-free Speech Communication and Microphone Arrays (HSCMA)
Trento, Italy, May 6-8, 2008

Cover Story

Demonstrations presented at the last AMIDA Review Meeting

The AMIDA first year review took place at the IDIAP Research Institute on February 7-8, 2008. On this occasion, the technology developed in the project was extensively demonstrated, both during a *demo session* on the first day, and through a *"surprise" demo* based on the first day's recordings and presented at the beginning of the second day.

Demo session

The *demo session* introduced the AMIDA reviewers, accompanied by many project members, to eight demos, some of them mainly based in one institution, but some others involving extensive collaboration between partners. These demos illustrated the following technologies and applications: Automatic Content Linking (outlined below), Mobile Meeting Assistant, Remote Monitoring, Meeting Summarization, Recognition of Visual Focus of Attention, Video Editing, Keyword Search in Speech, Remote Monitoring using a CE-HTML TV set.

Automatic Content Linking (ACL)

The Automatic Content Linking device is designed to perform searches at regular intervals, in a database of documents that can be prepared before a meeting, including for instance reports, emails that were exchanged, minutes and presentations from past meetings, and fragments of previous meetings. In this way, participants to a meeting receive in real-time suggestions about documents that are potentially relevant to the ongoing discussion. Participants are free to ignore them, or to start using them to enhance their discussion, e.g. with figures, precise facts, or decisions that were made in previous meetings. The same technology can also be used for offline meeting browsing augmented with related documents.

The search criterion is constructed based on the terms that were recognized automatically from the meeting discussion (some pre-specified terms can receive greater weight) and the results are presented as a list of document names ordered by relevance. A user interface offers the participants quick access to the content of the documents that are retrieved, if they believe that the contain valuable information to a given topic in the meeting.

Andrei Popescu-Belis, nominated head of WP6 by the AMIDA project board, described the ACL demo.

The Automatic Content Linking device relies on a number of modules, the main ones being the Document Bank Creator, the Document Indexer, the Query Aggregator, and the User Interface. These components exchange data through a subscription-based client/server architecture named the Hub. The components have been developed by members of the AMI Consortium at IDIAP, the University of Edinburgh, DFKI, and TNO.

Augmented Multiparty Interaction Distance Access (AMIDA) is an Integrated Project funded by the EC's 6th Framework Program, jointly managed by IDIAP (CH) and the University of Edinburgh (UK).

Issue N° 14 / April 2008
Editor : François Foglia, francois.foglia@idiap.ch
Newsletter

Demonstrations presented at the last AMIDA Review Meeting

OVERNIGHT DEMOS

The *"surprise" demo* involved overnight processing of audio, video, and slides that were recorded during the first day, thus demonstrating the adaptability of AMIDA technology to new, non scenario-based meetings. Technologies that were demonstrated were: Overnight ASR, Automatic Summarization of ASR Output, Storyboard Summary Visualizer.

Overnight ASR

Preparations for demonstrating the ASR system at the review started in earnest exactly two weeks prior at a 3 day ASR workshop at IDIAP (with researchers from USHEF, UEDIN, IDIAP, BUT). During this workshop, team members began porting the ASR system to a new unified modular framework. A trial run of the system was also carried out to assess the best audio acquisition setup and to make a dry run of the ASR. The audio acquisition that was adopted used an 8-channel circular microphone array and lapel microphone (for the main presenter). The modular ASR framework was also coupled to a new web-based ASR front-end developed and based at USHEF that allows audio to be uploaded via a web browser for processing and delivery of ASR output upon completion.

Review preparations also required the development of some new technologies and models:

- A speaker diarization system developed at ICSI was used for speech activity detection, segmentation and speaker clustering.
- The wordlist from our latest system lacked many key words. Notably, “AMIDA” and the some partners’ names were missing. Thus, the language model and dictionary were updated to better reflect the anticipated discussion topics.
- An automatic means for switching between the ASR outputs of the array or lapel was developed. It uses a comparison of speaker clusters from the two audio sources to select segments from the presenter (lapel) or audience (array).

The review meeting was captured using the audio setup established during the workshop two weeks prior. After cropping and resampling the array and lapel recordings at IDIAP, the data was uploaded to the web interface with results available a few hours later. Despite having our fair share of technical hiccups we were able to provide ASR output for subsequent processing.

Automatic Summarization of ASR Output

The aim of the summarization demonstration is the automatic derivation and display of information from recorded meetings to enable users to quickly find material from past meetings, and potentially even to catch up with a meeting that is ongoing.

The two types of information we derive for this demonstration are topic segments and extractive summaries. Both processes use ASR as input. When catching up with a meeting, topic segments can show a user what was discussed and provide links to appropriate parts of the meeting for a more detailed view. The purpose of extractive summary is to preserve the important parts of a meeting in a succinct manner, mimicking human meeting minutes.

With automatic extractive summarization, we aim to find the subset of utterances that best convey the important information in a meeting. Combining standard text-based features with prosodic, discourse and speaker cues, we find that significant improvements can be made in the quality of extractive summaries over a purely textual approach [1]. We have also found that effective term-weighting methods for summarization can be derived by considering how word usage varies between meeting participants [2]. In addition, we have demonstrated that an effective method for summarizing a meeting in progress is to adjust the utterance weights by referring to a very limited utterance context [3]. These approaches achieved significant results on the AMI and ICSI corpus data and during the overnight demo they were applied for the first time to a meeting of a completely different character.

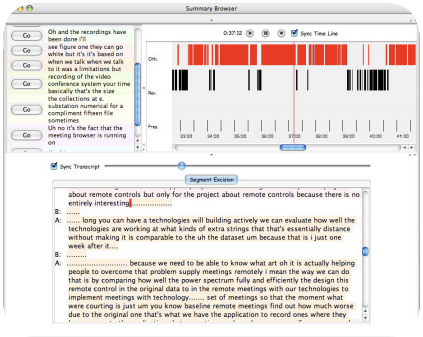

Automatic summarization of meetings

Topic segmentation divides up a meeting into segments at a high level, judging when participants change the subject of discussion. The problem can be formulated as detecting the speaker spurts that occur right before a new topic has been initiated. In the overnight demo, a preliminary version of our topic segmentater was used, leveraging the lexical cohesion statistics that are indicative of topic shifts. A more recent version of our topic segmenter integrates additional information about words, prosody, speaker movement and intention [3, 4].

The demonstrator can play the meeting audio in sync with the ASR transcript, which is in the bottom window. Above the transcript is a slider that allows a user to adjust the level of extraction: with the slider at the extreme right, the complete transcript is visible; as it’s moved left, the utterances judged less important are expunged until only those judged very important are visible. The topic window (top left) allows a top-down view of the meeting; clicking on a particular topic jumps directly to that part of the meeting.

To run these processes, the time-stamped words from the ASR group are first transformed into NXT format [5], which is the reference corpus format for the AMI project. This format is used by the extractive summary and topic-segmentation processes and also the input for the Storyboard demo. While we found some wrinkles in the transform process, there is nothing to suggest this kind of display cannot be brought online very quickly.

REFERENCES

[1] G. Murray, S. Renals, J. Carletta, and J. Moore, «Incorporating speaker and discourse features into speech summarization». Proc. ACL Human Language Technology 2006, New York, US, 2006.

[2] G. Murray, and S. Renals, «Towards online speech summarization». In Proceedings of Interspeech ‘07, Antwerp, Belgium, 2007.

[3] P. Hsueh and J. Moore, «Combining Multiple Knowledge Sources for Dialogue Segmentation in Multimedia Archives». Proc. ACL 2007, Prague, Czech Republic, 2007.

[4] P. Hsueh, J. Moore and S. Renals, «Automatic Segmentation of Multiparty Dialogue». Proc. EACL 2006, Trento, Italy, 2006.

[5] J. Carletta, S. Evert, U. Heid, J. Kilgour, J. Robertson and H. Voormann, «The NITE XML Toolkit: flexible annotation for multi-modal language data». Behavior Research Methods, Instruments, and Computers, special issue on Measuring Behavior, 35(3), 353-363, 2003.

AMI c/o IDIAP Research Institute, Centre du Parc, Rue Marconi 19, P.O. Box 592, CH-1920 Martigny,
info@amiproject.org - www.amiproject.org

Issue N° 14 / April 2008
Editor : François Foglia, francois.foglia@idiap.ch

Newsletter

Demonstrations presented at the last AMIDA Review Meeting

OVERNIGHT DEMOS *(CONTINUED)*

Storyboard Summary visualizer

SuVi - The Summary Visualizer is a constraint based multimodal layout system. It is an automatic layout tool for the visualization of multimodal summaries of the content of meetings as a storyboard. (See the comic-like image) This layout is well suited to represent the chronological sequence of events in a meeting in an appealing format. The storyboard contains original images taken from the video recordings of the meeting supplemented by text boxes and speech bubbles. The meeting is automatically segmented into the various topics covered and each new topic is marked by a text box. For each topic, the storyboard contains the important images, e.g., showing the main contributors and an excerpt of their contributions.

To connect SuVi with the AMI corpus we developed M2SuVi - Meeting to SuVi - which provides a generic interface to the AMI corpus allowing all these meetings to be summarized in a story-board style. In order to prove our generic approach, we created a storyboard summary based on the audio and video recordings of the AMIDA first review meeting overnight. SuVi and M2SuVi showed to be very robust with the new input data and the completely new domain. As the input was already provided in AMIDA's "NXT format" we had the first version of a story-board layout of the review meeting after around five minutes. After generating the first versions, however, we noticed a number of details, e.g., that the balloons were far too large and thus covered the whole image. We quickly added a parameter for the maximum length of text in a balloon and added «...» at the end to indicate that there is more text available. We also added other details such as page numbers to arrive at the final one-page and three-page versions within another two and a half hours. An example is shown to the left.

The content of meetings as storyboard

AMI COI Workshops provide content and context

The AMI Community of Interest is a valuable resource for technology transfer as well as for prioritizing the direction and pace of future research in AMI partner institutions. The challenge is to maximize the exchange between COI members and AMI partners in a very focused environment. Facing this challenge, the AMI Technology Transfer work package designed and conducted two invitation-only workshops.

Overview

The COI Workshop objectives were to:

- establish fruitful, collaborative working relationships between COI members and AMI Scientists
- chart the course of future tools, products and meeting processes using AMI Consortium technologies, and
- develop proposals for new joint projects involving COI and AMI Consortium

Planning for the workshops began in April of 2007 with the first one being conducted in Amsterdam immediately following the IBC conference (September 11-12, 2007) and the second in Martigny, Switzerland (February 4-5, 2008) immediately prior to the AMIDA first review. In parallel with the workshop planning, the project partners also developed the Mini-Projects program which was launched formally to the COI on September 12.

Format encouraged interaction

The two workshops were very similar in format and agenda. Formal presentations were kept to a bare minimum. Rather, the agenda emphasized direct contact between AMI scientists and the other types of participants (COI Members and Friends of AMI). After an introductory (plenary) session which covered technology transfer and the workshop guidelines, the scientists stayed at one table for individual meetings and demonstrations while the other participants rotated following the schedule which was prepared for them. Each schedule was developed based on areas of expressed interest (information gathered at registration using forms), the time which they had to participate and other factors.

After the demonstration or short presentation at the beginning of each breakout session, AMI Scientists were encouraged to ask engaging questions and to elicit feedback. The degree to which the dialog illustrated the scientists' mastery of the subject matter was judged by the COI member and friend of AMI. Content of discussions/remarks made during the breakouts contributed to a facilitated AMI Scientist brainstorming/reaction session on the afternoon of the second day.

During the second day of the workshop there were several small group and large group activities designed to increase communication and feedback. One session was dedicated to questions about mini-projects. A closing plenary involved all participants in discussing what they liked and disliked about the workshop format.

There appeared to be a wide range of understanding and interests from COI members ranging from those seeking an integrated prototype with all elements in an application scenario (or they thought perhaps that was what the Consortium would create/develop) to those who are interested in even smaller units (sub-units of the components shown) and to better define how AMI technologies fit into existing customer workflows and commercially available products.

Following the first workshop, a field trip to TNO Soesterberg provided AMI Scientists as well as COI members an opportunity to learn more about and have a closer look at the research facility where TNO Human Factors has been studying the impact of AMI technologies on business meetings. The visit included a tour of several test environments developed for space and defense industry projects.

Issue N° 14 / April 2008
Editor : François Foglia, francois.foglia@idiap.ch
Newsletter

AMI COI Workshops provide content and context *(continued)*

Conclusion

In all respects the objectives were met and the feedback at the conclusion of the workshop indicates that the investment was valuable for all participants. The COI members who attended praised the AMI Consortium and members for an outstanding level of integration at the level of research efforts. They were impressed with the vision of the consortium and, at the same time, the focus which has been applied to advancing science in the area of automatic meeting annotation and analyses. Discussion with AMI COI members about the mini-projects program indicates that the topic was clearly communicated/well constructed.

COI members who attended an AMI Workshop

Cisco	Hewlett Packard	Spiderphone	Media Publisher	Nokia
Polycom	TANDBERG	Ceannard	CapGemini*	Synthetron*
IBM Haifa	Oracle	Telecats*	SMART Technologies	WebEx

* indicates a company who joined the COI as a result of coming to the workshop as a Friend of AMI

Friends of AMI who have attended an AMI COI Workshop

Sentient	The Value Web	Jean-Francois Raffestin	Stephen von Rump
Novartis	Orange (France Telecom R&D)	T-XChange	

Demonstrations shown at the COI Workshops

JFerret Meeting Browsers	Automatic Video Editing	Automatic Content Linking	Smart Access to Presentation Content	Extractive summarization
Temporal Compression of Meeting Audio	Visual Focus of Attention Recognizer	Virtual Presence Support	Storybook-style multi-modal summarization	Mobile Meeting Assistant
Meeting Metadata standardization	Meeting Evaluation System	Automatic Segmentation of Meetings	Keyword Spotting	

Amsterdam workshop stats at a glance

- The maximum number of participants was set at 50.
- Of the 45 participants who attended: 21 were AMI Scientists, 4 were technology transfer people from AMI partner institutions, 7 participants were from companies which are not currently COI members (“Friends of AMI”), 11 represented COI member institutions and two were coordinators of the workshop.

Martigny workshop stats at a glance

- The maximum number of participants was set at 40.
- There were 34 registrations. 18 were AMI Scientists, 2 were Technology Transfer representatives from AMI institutions, 4 participants were from companies which are not currently COI members (“Friends of AMI”), 8 represented COI member institutions and two were coordinators of the workshop.

News and Upcoming Events

AMI Career Day at MLMI 2008

Wednesday, September 10, 2008

The AMI Career Day is a satellite event of MLMI 2008 focused on matching up candidates with members of the AMI Community of Interest who are looking for people with AMI know how and skills to work in their companies.

The Career Day is intended for students or interns working in AMI labs and other facilities with expertise in related fields who are interested in the opportunities that are available in industry/business to develop and commercialize products or services using intelligent meeting technology. It is an opportunity for those seeking to turn their education and experience into an exciting and rewarding career. Conversely, AMI COI members such as Intel, Oracle, Hewlett Packard, Cisco, and SMART Technologies can increase their AMI (see http://www.amiproject.org/vendors) technology knowledge and expertise by hiring young professionals trained in AMI labs and other qualified institutions.

The event will feature face-to-face meetings between companies and candidates, presentations of multimedia meeting technology from the AMI Consortium and invited companies, assistance with job search and resume preparation, and a cocktail party. The exact schedule will be announced later on this page.

Registration

The event is open and free of charge to MLMI 2008 participants, who are however required to register before May 31 using the online form (www.mlmi.info).

Students from AMI Consortium labs, and in particular AMI Trainees, are encouraged to apply to the AMIDA Training Program for travel support to MLMI 2008 and/or Career Day. The AMI Career Day is supported by the AMI Technology Transfer and AMIDA Training programmes.

Contact : cperey@perey.com
p.green@dcs.shef.ac.uk

Klewel Award: European Seal of e-Excellence

Klewel received the 2008 European Seal of e-Excellence award at the CeBIT International Trade show on information technology and telecommunications (Hannover, Germany) from the European Multimedia Forum federation.

Mael Guillemot (Klewel), Philippe Wacker (European Multimedia Forum)

AMI c/o IDIAP Research Institute, Centre du Parc, Rue Marconi 19, P.O. Box 592, CH-1920 Martigny,
info@amiproject.org - www.amiproject.org