Newsletter

Contents

News

AMI Knowledge and KnowHow Event

September 10, 2008

The AMI Consortium Technology Transfer work package

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

Applying AMIDA Content Hiding Slider to large scale online Synthetron discussions

MINI-PROJECT: SYNTHETRON (B) AND UNIVERSITY OF SHEFFIELD (UK)

This project is investigating applying technology developed by the AMIDA consortium for rapid browsing of meeting transcripts to assist the identification of key information from large scale Synthetron discussions.

Synthetron discussions are highly focused instant message style discussions centered on a very specific topic. As part of the process of these discussions participants can vote in agreement or disagreement with sentences which then are used to characterise and crystalise the discussion. Once the discussion is over Synthetron manually analyze the discussion and the key sentences in order to identify key themes and trends that have arisen throughout discussion. This analysis process is rapid for relatively small discussions, but as the number of participants and the amount of discussion increases the process of identifying key information becomes more complex.

The AMIDA consortium have developed dynamic compression techniques in order to aid users in rapidly recovering information from lengthy meeting transcripts. The algorithms that have been developed are able to make a very naive 'guess' as to the importance of parts of text within the transcript.

The user is then given an interface which allows them to view the transcript. The user is also given a slider which allows them to perform a transformation on the text according the amount of important information they wish to see. For example, in one implementation of dynamic compression, moving the slider to the left progressively removes any text which we think is unimportant leaving the user with a very high level overview of the transcript.

They can then identify parts of the transcript which are interesting to them, return the slider to the original position, and read the interesting parts in detail. In other implementations we alter the colour of unimportant parts of the transcript allowing the user to skip over these sections whilst browsing but still allowing them to read the full context should they wish.

All the alterations happen in real time, allowing the user to rapidly jump between high level overviews and the full detail. Thus the user can use the high level overview to identify key parts of the transcript and then use the full detail view to investigate these parts further.

This mini-project is investigating how dynamic compression technology can be applied to Synthetron discussions. The first challenge which has been addressed by the mini project is whether Synthetron discussions can be integrated into the dynamic compression framework developed for the AMIDA project. There are further challenges as a result of the difference between the two domains. As applied to meeting transcripts, dynamic compression can be used to identify key locations of information over the full transcripts. With Synthetron discussions the analysis process is more concerned with identifying key themes over the full transcripts. Therefore the mini project is investigating the use of different presentation approaches for Synthetron discussions.

Cover Story

Newsletter

Applying AMIDA Content Hiding Slider to large scale online Synthetron discussions (continued)

This analysis should also benefit AMIDA since these new presentation approaches may prove to be of use in the meetings domain - where identifying key trends across a series of meetings may prove useful.

Another focus of the mini-project is the subject of evaluating the dynamic compression technology. A problem with the lab studies of dynamic compression that have been undertaken so far is that they are often limited to a small and limited user base. Ideally user interface design should include both these types of studies

and large scale deployment of new technology. The mini project should allow a larger scale evaluation of the technology with real users who will hopefully benefit from using dynamic compression to perform their daily tasks.

For more information about this work, contact Simon Tucker at the University of Sheffield, s.tucker@dcs.shef.ac.uk. For more information about how to work with the AMI Consortium, contact Christine Perey at: *http://www.amiproject.org/business-portal/contacts*

Evaluating Oracle XML DB as a back end for the NITE XML Toolkit
MINI-PROJECT: ORACLE (F) AND UNIVERSITY OF EDINBURGH (UK)

One of the ways in which the AMI Consortium engages with industry is through «mini-projects» in which a Consortium member and a member of the Community of Interest jointly carry out work of mutual interest. In one of our current mini-projects, Oracle and the University of Edinburgh are evaluating the use of Oracle XML DB as a back end for the NITE XML Toolkit in order to see what advantages using it might bring.

The NITE XML Toolkit (NXT) is open source software for marking up synchronized sets of recordings with many different kinds of annotations. Although there are other annotation platforms available, NXT is unusual in using stand-off XML to allow the work of annotation to be distributed among different sites at the same time, and in representing both temporal and structural relationships among the annotations. The AMI Consortium relies on NXT extensively to create and search annotations of

the multimodal recordings that underpin their scientific and technological advances.

NXT has its own query language, NQL, based on first order predicate calculus, with its own dedicated implementation. However, the queries could be expressed in the XQuery standard that has emerged since XT was designed and written. Using an XQuery-based implementation would have several advantages: it would make NXT easier to maintain, it would allow us to add more readily some query operators that users have requested, and it might improve search performance. NXT's current query implementation loads everything it is searching into memory, which limits the amount of material that can be searched at once.

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

3/4

Newsletter

Evaluating Oracle XML DB as a back end for the NITE XML Toolkit (continued)
MINI-PROJECT: ORACLE (F) AND UNIVERSITY OF EDINBURGH (UK)

Although many NXT users would find XQuery too cumbersome and difficult to use, it could translate NQL queries into XQuery, run them, and then translate the result sets back into the form that the current NQL implementation returns.

In this mini-project, we are evaluating the use of Oracle XML DB release 11g as the XQuery back end. In the first phase of the project, which is being carried out now, we are hand-translating a representative set of queries into XQuery, loading an NXT-format corpus into Oracle XML DB, and benchmarking the results against the current NQL representation. Oracle XML DB allows both structured XML or relational treatments of the same data, and has two possible storage models, giving a great deal of flexibility in how it can be used.

On the other hand, NXT data is different from most XML. Although it is standards-compliant, its use of stand-off means that many more individual files make up one coherent data set than is usually the case.

The first phase will tell us what gains we can expect from porting our implementation to Oracle XML DB and what sort of difficulties we might encounter. It will also test parts of the Oracle implementation that are less exercised by more mainstream users, and could potentially reveal new properties that affect their future target uses. The XQuery standard currently does not allow for queries that change the XML data; this facility is part of XQuery Update, which is currently a W3C Candidate Recommendation. In the second phase of the project, we will benchmarking using a pre-release copy of the next XML DB version to test whether its XQuery Update implementation could be used to implement NXT's manipulation of annotation sets as well. Although an NQL implementation that is limited to searching static data sets would still be useful, using XQuery Update could allow us to retire the existing NQL implementation, giving it the status of a reference implementation for the operators that it includes. As in the first phase of the project, this work will test the feasibility of such a replacement, but not include the development work that this would require.

For more information about this work, contact Jean Carletta at the University of Edinburgh, jeanc@inf.ed.ac.uk. For more information about how to work with the AMI Consortium, please contact Christine Perey at: *http://www.amiproject.org/business-portal/contacts*
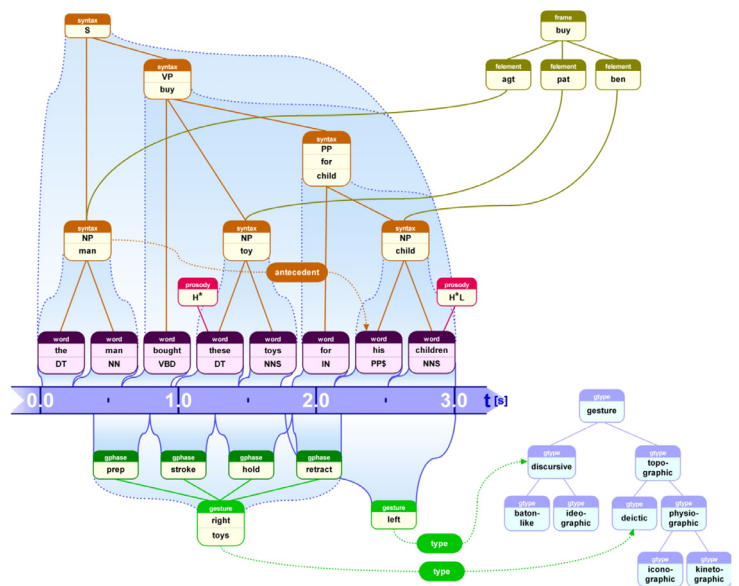
MLMI 2008 - 5th Workshop on Machine Learning and Multimodal Interaction
8-10.09.08, UTRECHT, THE NETHERLANDS

The 5th Joint Workshop on Machine Learning and Multimodal Interaction was held in Utrecht, Netherlands on September 8-10, 2008. It was a single track conference with approximately 100 attendees. This year's themes included: multimodal processing, speech and language, interactive systems, human interaction and activity as well as speech processing.

There were 4 interesting keynote presentations:

1. «The Giraffe - A Mobile Telepresence Physical Avatar» by Stephen von Rump from HeadThere Inc. He gave an overview and some demos of their mobile video conferencing platform.

2. «Measuring interactivity on search engines is hard unless you're Yahoo, Microsoft, or Google» by Mark Sanderson from University of Sheffield. The talk was on how test sets and evaluation methodologies for search engines need to evolve so that they become more interactive with current users and web browsers.

3. «HCI2: Human-Centered Intelligent Human-Computer Interaction» by Maja Pantic from the Imperial College London and the University of Twente. The talk was on human centered user interfaces which are beyond the traditional keyboard and mouse. Maja talked in detail on how human communicative signals, especially facial expressions, could be used for such interfaces.

4. «Communicating with a distinctive embodied conversation agent» by Catherine Pelachaud from University of Paris. Catherine demonstrated the Embodied Conversational Agent (ECA) named Greta, which their research group has created. It is capable of communicating with users through both verbal and nonverbal cues.

The topics presented covered a broad spectrum of topics including multi-modal processing, design and evaluation of meeting browsers, analyzing human interaction and activity and speech processing. A good number of posters reflecting the themes of the conference were also presented over 2 sessions. MLMI hosted 3 satellite events as well - a speech recognition workshop for Dutch, an inter project meeting on the evaluation of space-time audio processing and an AMI knowledge transfer session.

Katayoun Farrahi, Samuel Thomas, Phil Garner: Idiap Reasearch Institute, Av. Marconi 19, CH - 1920 Martigny

2/4

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

Newsletter

News

AMI Knowledge and KnowHow Transfer Day at MLMI 2008 (AKKT)

8-10.09.08, Utrecht, The Netherlands

The AMI Consortium Technology Transfer work package

The AMI Consortium Technology Transfer work package hosted a one-day event on September 10, 2008 in Utrecht for AMI Scientists, Community of Interest members and a group referred to as the «friends of AMI.»

This was the first event in which the emphasis was more on placing the scientists in open positions with AMI CoI members or generating contracts for research for hire. There were 27 registered attendees on the day of the opening of the workshop.

The stated objectives of this day were:

- To provide AMI Community of Interest members and Friends of AMI an overview of the services for hire offered by AMI members in the domains of improving human-to-human communications during and between meetings, and the application of AMI technologies in business.

- To provide AMI Community ofInterest members a forum in which to explore existing or new AMI Mini-Projects, leading to applied research towards mutual objectives in a cooperative and supportive atmosphere.

- To foster the establishment or deepening of ties between AMI scientific research partners and the commercial entities in the Community of Interest such that Technology Transfer is stimulated and encouraged.

In the two days following the AKKT event, all the COI members and several of the Friends of AMI wrote by e-mail thanking for this opportunity to see the demonstrations and to meet with AMI candidates. Several used words such as «stimulating», «exciting», and «superb» when describing their feelings and impressions of AMI technologies and their thoughts about the event.

For additional information on the applications for technologies which are being developed or would like to learn their status for technology transfer, please contact *Christine Perey, cperey@perey.com*.

Selected publications

Machine Understanding of Human Behavior.
M.Pantic, A.Pentland, A.Nijholt, T.Huang
In Proceedings AI for Human Computing (AI4HC'07), workshop at IJCAI, pages 13-24, Hyderabad, India, 2007.

Maximum Likelihood and Maximum Mutual Information Training in Gender and Age Recognition System.
V.Hubeika, I.Szoke, L.Burget and J.Černocký
In Lecture Notes in Computer Science, volume 4629/2007, number 9, pages 496-501, ISSN 0302-9743, 2007.

Microphone Array Beamforming Approach to Blind Speech Separation.
I.Himawan, I.McCowan, M.Lincoln
Proceedings of the 4th Joint Workshop on Multimodal Interaction and Related Machine Learning Algorithms (MLMI 07), pages 295-305, vol. 4892, Springer-Verlag, 2007.

Minimum Mutual Information Beamforming for Simultaneous Active Speakers.
K.Kumatani, J.McDonough, U.Mayer, T.Gehrig, E.Stoimenov, M.Wolfel
In IEEE Workshop on Automatic Speech Recognition & Understanding, 9-13 December, Kyoto, 2007.

Mutually Coordinated Anticipatory Multimodal Interaction.
A.Nijholt, D.Reidsma, H. van Welbergen, H.J.A. op den Akker, Z.Ruttkay
Nonverbal Features of Human-Human and Human-Machine Interaction, 29-31 October 2007, pages 73-93, Springer Verlag, Berlin, Patras, Greece, 2008.

On the Contextual Analysis of Agreement Scores.
D.Reidsma, D.Heylen, H. J. A. op den Akker
In Proceedings of the LREC Workshop on Multimodal Corpora, pages 52-55, ELRA, ELRA, Marrakech, Morrocco, 2008.

Optimizing Bottle-Neck Features for LVCSR.
F.Grezl, P.Fousek
In IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 4729-4732 , 2008.

Physicality and Cooperative Design
V.Dhaval, D.Heylen
In MLMI08, 5th Joint Workshop on Machine Learning, Utrecht, The Netherlands, Springer LNCS proceedings, 2008.

Platform for Evaluation of Image Classifiers.
J.Silhava, V.Beran, P.Chmelar, Adam Herout, M.Hradis, R.Juranek, P.Zemcik
In Spring Conference on Computer Graphics, pages 103-109, Comenius University in Bratislava, Budmerice, SK, 2007.

Probabilistic and bottle-neck features for LVCSR of meetings.
F.Grezl, M.Karafiat, S.Kontar, J.Černocký
Proceedings IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 757-760, Hononulu, US, 2007.

Recognition and Understanding of Meetings Overview of the European AMI and AMIDA Projects
H.Bourlard, S.Renals
Techreport, LangTech 2008, February, 28-29, number 27, 2008.

Recognition and Understanding of Meetings: The AMI and AMIDA Projects
S.Renals, T.Hain, H.Bourlard
In IEEE Workshop on Automatic Speech Recognition & Understanding (ASRU), pages 238-247, 9-13 December, Kyoto, 2007.

Recognition of Dialogue Acts in Multiparty Meetings using a Switching DBN
A.Dielmann, S.Renals
In IEEE Transactions on Audio, Speech and Language Processing, pages 1303-1314, volume 16-5, September, 2008.

Robust Multi-Modal Group Action Recognition in Meetigs from Disturbed Videos with the Asynchronous Hidden Markov Model
M.Al-Hames, C.Lenz, S.Reiter, J.Wallhoff, J.Schenk, G.Rigoll
In IEEE International Conference on Image Processing, 2007 (ICIP 2007), pages II:213-216, San Antonio, TX, 2007.

Search in speech for public security and defense
J.Cernocky, I.Szoke, M.Fapso, M.Karafiat, L.Burget, J.Kopecky, F.Grezl, P.Schwarz, O.Glembek, I.Oparin, P.Smrz, P.Matejka
In IEEE Workshop on Signal Processing Applications for Public Security and Forensics, 2007 (SAFE '07), pages 1-7, IEEE Signal Processing Society, Washington DC, USA, 2007.

Sentiment classification with interpolated information diffusion kernels
S.Raaijmakers
Proceedings of the 1st international workshop on Data mining and audience intelligence for advertising, p.34-39, August 12, 2007, San Jose, California.