MORNING SEMINARS

GIUSEPPE RICCARDI

"TALKING TO COMPUTERS: FROM SOUNDS TO WORDS, FROM ACTIONS TO INTERACTIONS"

UNIVERSITY OF TRENTO

ABSTRACT

Ever since computers were invented, scientists pursued the vision of "talking" to them and eventually adopting them. This vision has percolated beyond the scientific community and explored by science-fiction film makers such as Stanley Kubrick in "2001: A Space Odissey". Scientists from many disciplines have spent the last half-century investigating this research challenge. The issues scientists have been investigating are closely related to the fundamentals of human-human communication. What is the internal representation machines need to communicate with humans? What is the process of understanding and reacting to the human generated-signs, from sounds to gestures? What is the most efficient and effective protocol of communication? What is the motivation for a human to engage into a human-machine conversation?.

Scientists have decomposed the hard problem of "communication" into a "code-breaking" problem. Spoken words have been decomposed into acoustic or phonetic units that can be then aggregated into morphemes or words. Words can be then parsed searching for larger units (phrases) which are annotated with semantic features. Semantic interpretations are then learned to infer machine actions (e.g. computers will arrange for a flight reservation). As it occurs in human interaction, we expect machines to negotiate, cooperate and give feedback. In the early nineties, primitive interactive machines were built for small vocabulary tasks and limited capabilities. A decade later, computer can recognize millions of words continuously spoken in radio broadcasts, translate news text without requiring any language expertise and score very well at lip reading. What is the revolution that occurred in the last decade of research in human-machine communication? I will select the major milestones in the conversational system research as well as point out current limitations and future research challenges.

BIOGRAPHY

Prof. Riccardi received his Laurea degree in Electrical Engineering and Master in Information Technology, in 1991, from the University of Padua and CEFRIEL Research Center, respectively. From 1990-1993 he collaborated with Alcatel-Telettra Research Laboratories (Milan, Italy). In 1995 he received his Phd in Electrical Engineering from the Department of Electrical Engineering at the University of Padua, Italy. From 1993-2005, he worked first at AT&T Bell Laboratories and then AT&T Labs-Research where he worked in the Speech and Language Processing Lab. In 2005 joined the faculty of Engineering at University of Trento and is affiliated with the interdisciplinary Department of Information and Communication Technology at, Italy.

Prof. Riccardi's research on stochastic finite state machines for speech and language processing has been applied to a wide range of domains for task automation. He participated at the creation of the state-of-the-art AT&T spoken language system used in the 1994 DARPA ATIS evaluation. He and his colleagues have been pioneering the speech and language research in spontaneous speech for the well-known "How May I Help You?" research program. His research

on learning finite state automata and transducers has lead to the creation of the first large scale finite state chain decoding for machine translation (“Anuvaad”).

Prof. Riccardi has co-authored more than 70 papers and 25 patents in the field of speech processing, speech recognition, understanding and machine translation. His current research interests are stochastic language modeling, language understanding, spoken dialogue, machine learning and machine translation.

Prof. Riccardi has been on the scientific committee of EUROSPEECH, INTERSPEECH, ICASSP, NAACL and ACL. He has co-organized the IEEE ASRU Workshop in 1993, 1999 and 2001. He has been the Guest Editor of the IEEE Special Issue on Speech-to-Speech Machine Translation. He is on the Editorial Board of the ACM Transactions of Speech and Language. He is elected member of the IEEE SPS Speech Technical Committee (2005-2008). Prof. Riccardi is senior member of IEEE, ACL and New York Academy of Science.

Prof. Riccardi have received many national and international awards and more recently the Marie Curie Excellence award by the European Community.

GIUSEPPE DI FABBRIZIO

"STOCHASTIC NATURAL LANGUAGE GENERATION"

AT&T LABS – RESEARCH

ABSTRACT

Statistical methods are today predominant in spoken dialogue systems (SDS) and are quite mature for speech recognition, language understanding, and dialogue act detection tasks. However, spoken and natural language generation (SLG/NLG) still require a great deal of handcrafting, expert knowledge, and considerable cross-domain portability efforts to make the technology scalable for real systems. SDS are largely based either on template-based or rule-based generation techniques.

These approaches yield high quality output, but require costly maintenance and domain customizations. A trainable generation approach instead can learn generation rules from data and promises a more general cross-domain solution providing robust and adaptable models. Trainable sentence planning, for example, learns which combination of operations for aggregation and content ordering produces the highest.

BIOGRAPHY

Giuseppe Di Fabbrizio graduated in electrical engineering (MSEE) from the Politecnico di Torino, Torino, Italy in 1990. From 1990 to 1995, he was a senior researcher with Telecom Lab Italia (formerly CSELT), Torino. In January 1996, he joined AT&T Labs - Research in Florham Park, NJ, as senior researcher where he is currently member of the Speech Services Research Laboratory. During his carrier he mainly conducted research on spoken dialogue systems, natural language processing and speech services publishing more than 40 papers on these subjects and filed more than 20 US and EU patents. He was instrumental in the development and deployment of the AT&T VoiceTone® Dialog Automation product for the AT&T business enterprise customers. He received the AT&T Exceptional Achievement Award in 2003, the Most Innovative Speech Solution SpeechTek Conference Award in 2003, and the AT&T Research Excellence Award in 2002. Di Fabbrizio is a member of the International Speech Communication Association (ISCA), the Association for Computational Linguistics (ACL), and Senior Member of the Institute of Electrical and Electronics Engineers (IEEE).

GIUSEPPE CARENINI

"INTERACTIVE MULTIMEDIA SUMMARIES OF CUSTOMER REVIEWS"

UNIVERSITY OF TRENTO

ABSTRACT

Many organizations are faced with the challenge of summarizing large corpora of text data. One important application is evaluative text, i.e. any document expressing an evaluation of an entity as either positive or negative. For example, many websites collect large quantities of online customer reviews of consumer electronics. Summaries of this literature could be of great strategic value to product designers, planners, manufacturers and consumers.

In this seminar, I will first present and compare two approaches to the task of summarizing evaluative arguments. The first is a sentence extraction-based approach, while the second is a language generation-based approach. These approaches have been tested in a user study which indicates that an effective method for summarizing evaluative arguments must effectively combine them.

In the second part of the seminar, I will describe an interactive multimedia interface which presents the knowledge extracted form a corpus of evaluative arguments not only as a natural language summary but also in a hierarchical visualization mode. The interface is interactive in that it allows the user to explore the original dataset through intuitive visual and textual methods. Results of a formative evaluation of our interface show general satisfaction among users with our approach.

BIOGRAPHY

Laurea in Computer Science from the University of Milan, Italy (1988). Research Fellow IRST, Trento, Italy (1989-1991). MS and PhD from the University of Pittsburgh (USA) in Intelligent Systems in 1995 and 2000 respectively. Postdoc at the Computer Science Dept., University of British Columbia, Canada (2001-2002). Assistant Professor at the Computer Science Dept., University of British Columbia, Canada (2003-2005). Associate Professor, DIT, University of Trento (2006).

His interests lie at the intersection of Artificial Intelligence and Human-Computer Interaction. In particular, he is interested in improving human computer communication in the context of decision making. He has published approx. 50 papers in the areas of natural language generation, recommender systems, multimedia summarization and information visualization.

PAOLO BAGGIA

"TECH SPEECH TECHNOLOGIES AND STANDARS"

LOQUENDO

ABSTRACT

An overview of the evolution of speech technologies and speech applications in recent years and an in-depth evaluation of the impact of W3C and IETF standards
on speech application development. The VoiceXML and other W3C standards will be briefly introduced and the recent evolutions of the standards towards multimodal and video services will be described.

BIOGRAPHY

Paolo Baggia is currently Director of International Standards at Loquendo and actively involved in W3C and the VoiceXML Forum, where he is a member of the Board of Directors. Formerly, he lead the VoiceXML browser development for Loquendo's VoxNauta platform.

In 1989, he joined CSELT, Telecom Italia's R&D lab, and was involved in several research activities, including natural language, speech parsing, spoken dialog design, and language modelling. In 1999, he took part in a large-scale speech application: the automation of TreniItalia railway company's call centers. He holds a degree in computer science from the Università di Torino.

AFTERNOON SEMINARS

PAOLO FERRAGINA

"NEXT GENERATION SEARCH ENGINES"

UNIVERSITY OF PISA

ABSTRACT

In this talk we survey basic and advanced algorithmic solutions for the design of search engines that index textual and XML data. We then present two interesting applications in the context of Humanities Computing and Text Mining. Finally, we discuss future directions of research and challenging applications in which sophisticated data compression techniques and indexing data structures may turn to be key enabling technologies.

BIOGRAPHY

I am Associate Professor of Computer Science at the University of Pisa, and scientific coordinator of Signum, a research center on humanities computing of the Scuola Normale Superiore di Pisa. I got my Laurea degree (summa cum laude, 1992) and my PhD (1996) in Computer Science from the University of Pisa, and my Post-doc from the Max-Planck Institut fur Informatik (Saarbrucken, 1997-98). From 1998 to 2000, I've been Assistant Professor at the University of Pisa.

My research is mainly devoted to the design, analysis and experimentation of algorithms and data structures for storing, compressing, mining and retrieving information from large amounts of textual data like Web repositories, XML file collections, textual databases, genomic/DNA sequences. My research results received one US Patent (No. 6,434,566, 13 August 2002) and some international awards: "Best Land Transportation Paper Award" from IEEE Vehicular Technology Society (1995); "EATCS Doctoral Dissertation Thesis Award" (1997); "Philip Morris Award on Science and Technology" (1997); "Research Capital award" from the University of Pisa (2002).

I have served as PC member of many International Conferences on Theoretical Computer Science, specifically in the field of Algorithmics. I've been co-chair of International Conference on FUN with Algorithms (2004), DIMACS Workshop on the Burrows-Wheeler Transform (2004), Symposium on String Processing and Information Retrieval (2006). I've also been plenary speaker at CPM '04 and SPIRE '05, as well at various Workshops and Meetings on Algorithmics. I served as (co)editor of two special issues on the international journals: Theoretical Computer Science and Theory of Computing Systems (June 2006). I have also authored one chapter on "String search in external memory: Algorithms and data structures" in the Handbook of Computational Molecular Biology (CRC Press, Editor Srinivas Aluru), one Italian book on Cryptography (Bollati Boringhieri, 2001), and I'm one of the Area Editors of the Encyclopedia of Algorithms (Spinger, Editor Ming-Yang Kao).

Currently I'm leading two international projects: MIUR-FIRB Italy-Israel project on "Pattern matching and discovery algorithms on discrete structures, with applications to bioinformatics", and a three-years project funded by Yahoo! Research on "Data compression and indexing in hirerarchical memories".

PIERO COSI

"VIRTUAL HUMANS - EMBODIED CONVERSATIONAL AGENTS"

CNR- PADOVA

ABSTRACT

After a brief introduction on "Embodied Conversational Agents" (ECA) and "Virtual Humans" (VH), the state of the art of the most recent facial animation techniques will be described together with the last findings of Text-To-.Speech (TTS) synthesis research.
In particular, various applications of conversational agents will be introduced and possible subjects being involved in their development and implementation will be listed: movie companies and studios, animation software farms, intitutions, universities and research centers.

In conclusion, the recently attractive field of emotions will be analyzed and in particular the audio/visual rendering of emotions will be introduced.

Finally, future trends in this field will be suggested together with new possible applications.

BIOGRAPHY

Academic Background:
Doctoral Degree in Electronic Engineering (December 1981), (specialization in Computer Science) from Padova University (Italy). Thesis subject: LPC analysis synthesis and Pitch extraction.

Present status:
Researcher of Istituto di Fonetica e Dialettologia, C.N.R. Padova.(since March 1984).

Scientific Activity:
Basic and applied research on speech processing (Analysis, Synthesis and Recognition) and experimental phonetics.

UMBERTO BASSO

"DELIVERING A MULTIMODAL VIRTUAL ASSISTANT ACROSS DIFFERENT CHANNELS. THE EXPERIENCE OF A HUMAN DIGITAL ASSISTANT"

H-CARE

ABSTRACT

A multichannel virtual assistant should be accessible everywhere to everyone: no matter the device, no matter the user's skills. Caren, the Human Digital Assistant, interacts through 3D face animations and text-to-speech: she must be capable of a highly personalized interaction, where audio and video must be rendered in real-time, for the best user experience.

Caren will explain how she can be so dynamic in a carrier-grade scenario, where real-time 3D rendering, scalability to thousands of concurrent users, and integration with enterprise IT systems are the requirements for an effective implementation, and the device is not a side where she likes to depend on.

Tags: avatar, TTS, 3D, OpenGL, animation, real-time, rendering, encoding, streaming, HDA, Caren.

BIOGRAPHY

In 2005 he co-founded H-Farm, the incubator for technological innovation located between Treviso and Venezia.

He is the creator and managing Director of H-care, the company dedicated to the developing of the Human DIgital Assistant platform.

From 1998 to 2004 he contributed to E-TREE startup and become partner, Managing Director and Responsible for technology and Innovation, leading the company projects of developing internet portals.

Between 1995 and 1998 he took part to the realization of Mall Italy Lab, the first experimentation of virtual mall in Italy for Benetton Group.

He is now member of Equiliber scientific board for digital technology, an association dedicated to the planning of "an engine for equilibered culture" that studies the social, cultural and environment consequences of technology innovation.

Passionate of communication and network collaboration since the 80's and the first Bulletin Board System, he has always dreamed to do what he is now doing.

He attended Burning Man Project in 2003 and 2005.

Umberto Basso was born in 1969 in Treviso, where he lives and works.

WALTER BATTISTETTI

"SPEECH TECHNOLOGY: FROM THE LAB PROTOTYPE TO BUSINESS APPLICATIONS"

SPEECH VILLAGE

BIOGRAPHY

ABSTRACT

The talk will address the interdisciplinary factors that make speech interface design a crucial task. Spontaneous speech can induce a strong underestimation of the difficulty involved in speech interface design. We believe this is the reason speech application design has been for a long time an "hand-crafted task". Emerging standards is opening a new engineering era: a crucial factor for improving the quality of speech services and reducing the costs of development. We have focused our activity on engineering the process of developing sophisticated speech applications with highly engineered tools. In the talk we will discuss how a high-tech start-up can balance high-value technology products, financial resources and offering to the US, South American and European Markets.

BERNARDO MAGNINI

"RECENT ADVANCES IN OPEN DOMAIN QUESTION ANSWERING"

ITC-IRST

ABSTRACT

Open Domain Question Answering (QA) systems accept natural language questions (as opposed to keywords) and return to the user exact answers (as opposed to a list of scored documents).

Under the impulse of the QA track at the TREC evaluation campaign, QA has become both a hot research topic in Computational Linguistics and a challenging perspective for real use applications.

This seminar will report about current techniques and resources for Question Answering, including recent work on Cross-Language Questions Answering, where questions are expressed in a source language and the answer is searched in a document collection of a different language.

Finally, special attention will be put on evaluation methodologies, a necessary step in order to measure performance progresses of Question Answering systems.

BIOGRAPHY

Magnini is senior researcher at ITC-IRST (Istituto per la Ricerca Scientifica e Tecnologica) in Trento, where he coordinates the research line on Technologies for Text Processing (TexTec). He is contract professor at the University of Trento (since 2003) and at the University of Bolzano (since 2005).

His research interests are in the field of Natural Language Processing, with particular interests in question answering systems, word sense disambiguation and to the application of language technologies to the field of Semantic Web.

He is currently responsible of the multilingual question answering track at CLEF (Cross Language Evaluation Forum) and co-organizer of the Recognizing Textual Entailment challenge (PASCAL RTE-2).

He served as local organizer co-chair of EACL-06, the 11th Conference of the European Chapter of the Association for Computational Linguistics.