The (IM)2 Newsletter

Every month the (IM)2 Newsletter brings you the latest and hottest
scientific and administrative news about the (IM)2 NCCR and related topics

The first round of White Papers is a massive success, with 13 new contributions to (IM)2

A key feature of the (IM)2 NCCR, the concept of White Papers is meant as a dynamic management tool to allow (IM)2 to follow the rapid evolution of the state-of-art in the field of information technology. In addition to the 9 Individual Projects[1], (IM)2 members are given the opportunity once a year to submit new research proposals. White Papers are meant to either fill gaps in the original projects description, or to cover new areas.

The opportunity to submit White Papers also increases the number of projects leaders within (IM)2. Far from reducing the role of the IP Heads, it has allowed younger faculty members, group leaders or Tenure Track professors to receive funding for original and promising research projects.

Sixteen proposals were received and reviewed internally by the IP Heads and Management Board members. Based on scientific value, relevance to (IM)2, priority, and — ultimately — available funds, nine proposals have been selected for funding, and four more will be accepted in a revised form. In selecting the White Papers, the IP Heads have strongly emphasized the need to integrate the teams working on similar issues. A direct consequence is that every accepted White Paper is assigned to one (or more) IP, and it will be the responsibility of the IP Heads to enforce smooth integration of the new projects and to keep a global view of all research and development efforts in a given area.

We now give a brief description of the 13 accepted White Papers, sorted by domain.

Scene Analysis

Scene Analysis already has a dedicated IP, but is obviously an area of uttermost importance for (IM)2: no less than 5 original contributions will strengthen research in this area.

Prof. Pierre **Vandergheynst**, from EPFL-ITS, together with the team of Prof. Thierry Pun, UniGE/CVML, will develop a new joint framework of multiresolution image representation based on deterministic nonlinear image approximations and corresponding stochastic image models, and demonstrate its superior performance over the classical methods of multiresolution data analysis that are based on dyadic approximations and independent identically distributed (i.i.d.) stochastic image models.

Dr Jean-Philippe **Thiran**, also from EPFL-ITS, will lead two new projects. The first will investigate the theoretical aspects of multimodal signal processing, with direct contributions not only to (IM)2.SA but also to (IM)2.MI. The second aims at developing new scene analysis methods based on mathematical models of the behavior of the objects in the scene, especially in the case of human beings (traffic, crowds, etc.). It will be carried out in collaboration with Dr Michel Bierlaire, EPFL-IMA and MIT, a recognized expert in behavioral models.

Two seniors from the Computer Vision group at IDIAP, Dr Daniel **Gatica-Perez** and Dr Jean-Marc **Odobez** will develop methods and algorithms for the simultaneous tracking of people and recognition of their activity, using a multiple camera setup soon to be added to IDIAP's smart meeting room.

Finally, Prof. Luc **Van Gool** from the Computer Vision Group at ETHZ will develop self-learning trackers, that start off as generic trackers, but evolve towards trackers optimized for the tracking of a specific object that e.g. has been indicated in some initial video frame by the user or that has drawn the attention of the system in some other way. As the video proceeds, the trackers try to get better and better at specifically tracking 'their' object.

The three "tracking" proposals from EPFL, IDIAP and ETHZ build on different expertise and aim at different results, but will nevertheless work in close collaboration.

Multimodal Integration

Dr Samy **Bengio**, head of the Machine Learning group at IDIAP, will lead a small consortium in order to investigate the possibility of classifying spontaneous brain activity based on either reconstructed brain activity maps, or directly from EEG recordings. Indeed, the problem of accessing brain information is essential in any mid- to long-term human-machine interface development. The very fast communication channel of the brain is the paradigm for very innovative and unchallenged new interfaces between humans and non-humans. Moreover, many situations arise where classical interaction (with hands, speech, movements, etc) is difficult: when all other senses are active, for handicapped persons, or simply to be able to control one's DVD player while sitting on a sofa and just thinking!

Multimodal Dialogue Management

The (IM)2.MDM team lead by Prof. Susan **Armstrong**, UniGE-ISSCO, will add two components to their IP. The first deals with statistical detection of discourse and dialogue acts from unrestricted texts, in both human-to-human dialogue and human-to-machine dialogue, in both speech and text modalities and in a variety of languages — English and French initially — using robust techniques based on Machine Learning.

The second aims at automating the production and management of semantic knowledge and ontologies for dialogue understanding, providing formal ontologies to language processing software, and thus endowing a computer with some human understanding capacities.

Access and Content Protection

Dr Samy **Bengio**, IDIAP, and Dr Andrzej **Drygajlo**, EPFL-ITS, will work on improving the robustness of speaker verification with a statistical machine learning approach, with the objective of making it better suited for real-life applications.

Speech Processing

The speech processing part of (IM)2 will benefit from international collaborations with joint projects and access to available data and expertise from Prof. Christian **Wellekens**, Eurecom, and especially from the team of Prof. Nelson **Morgan**, ICSI-Berkeley, who has been working on similar projects including smart meeting minutes for several years.

Smart Meeting Manager

At the border between speech processing and smart meeting room, the proposal from Dr Iain **McCowan** and Darren **Moore**, from the Speech Processing group at IDIAP, will address the problems of acquiring clean speech, free of cross-talk, with minimal constraint upon the user(s), detecting the periods of voice activity for each user, and dynamically determining the location of each user. This will involve research and development work on the high-end audio recording equipment available at IDIAP, and will include contributions from the team of HEVs in Sion.

Finally, CSEM will also broaden the scope of the smart meeting manager by working on specialized multimodal extensions for medical consultations or meetings.

[1]See issue 0 of the (IM)2 Newsletter or the (IM)2 web site www.im2.ch for a description of the IPs.

The Research Group ISSCO at UniGeneva

ISSCO is a research group in the School of Translation and Interpretation (ETI) of the University of Geneva. ISSCO was founded by the Dalle Molle Foundation in 1972, based in Lugano, Switzerland. and is thus the oldest NLP lab in Switzerland. In 1976 it moved to Geneva and was attached to the University of Geneva. ISSCO is now permanently integrated into the Multilingual Information Processing unit (TIM) of the ETI where post-graduate studies involving translation technology can be undertaken.

ISSCO's initial mandate was research into cognition and semantics, with a special emphasis on man/machine communication. This rapidly led to specialization in natural language processing and, in particular, in multilingual language processing and translation issues. Over the years, the group has been active in projects and studies in knowledge representation, machine translation, natural language interfaces to database systems, lexicon and grammar development for NLP, artificial intelligence, dialogue modelling, belief system modelling, agent-based systems, lexical semantics, machine learning, corpus studies and evaluation.

ISSCO has directed and collaborated in a large number of European and Swiss projects including EUROTRA (European Machine translation project), MULTEXT (Multilingual lexical and textual resources), EAGLES (language engineering standards) and its follow-up project ISLE. ISSCO also works closely with the linguistic services of the federal administration of Switzerland in language technology matters.

People & Activities

The group is led by Prof. Margaret King who has more than 30 years of experience in the field and is a Fellow of the European Coordinating Committee for Artificial Intelligence. Prof. Susan Armstrong is head of the IM2 project on Multimodal Dialogue Management and has been active in the field for the past 20 years. Other members of the group include Dr. Pierrette Bouillon, Dr. Alexander Clark, Dr. Andrei Popescu-Belis and another 10 or so researchers and assistants. ISSCO regularly welcomes visitors for short and longer-term stays.

Initially, ISSCO was founded purely as a research group, but over the years has become active in teaching and consulting in translation technology. Already in 1974 the group was organizing tutorials, intensive study over a week of particular domains, such as computational semantics, formal linguistics and machine translation. The first contributions to regular University teaching in language technology date from 1978. More recently, post-graduate and continuing education courses have been established in Multilingual Information Processing. ISSCO also maintains the Swiss Language Observatory Web site and consults on language technology matters for international organizations and private industry.

The ISSCO group: Front: V. Tumelaire, S. Armstrong, V. Sauron, S. Seppälä. Middle: M. King, B. Cartoni, M. Starlander, G. Robert. Back: S. Halimi, P. Bouillon, A. Popescu-Belis, A. Clark.

Research Themes & Projects

Evaluation of machine translation software: Building upon the group's previous experience with standards for software evaluation (EAGLES), this activity focuses on the construction of a unified model for MT evaluation. The goal is to build a customizable quality model that will be applicable in a variety of contexts. Several hands-on workshops have been organized through the European ISLE project, to disseminate results and obtain feedback.

Spoken language translation: Building on previous work in MT at ISSCO, the work has been extended to spoken language translation in limited domains (based on initial collaboration with Telia Research and SRI). Current emphasis is on developing a robust environment for translation of fixed phrases in the medical domain. Recognition is constrained by a context free grammar (derived from a unification-based grammar) which serves both for source language analysis and target language generation.

Lexical semantics: In collaboration with the Computational Lexicon Working Group in the ISLE project, guidelines and standards for multilingual lexcions are under development as well as a lexical entry tool and database. Research work on the Generative Lexicon is a long-standing interest reflected in the bi-annual workshops (GL2001, GL2003) organized by ISSCO. In collaboration with French partners, work is underway on the automatic acquisition of lexical resources from corpora in the framework of inductive logic programming.

Machine Learning: Work in machine learning techniques for NLP has been ongoing in a European project on Learning Computational Grammars. Exploration of supervized as well as unsupervized methods have been applied in the domains of morphology, context-free grammars and language modelling. This work is also being applied to the problem of automatic bilingual terminological acquisition and text classification.

Machine-readable lexical resources: This activity aims at providing large scale access to high quality resources derived from printed material. For instance, electronic dictionaries from commercial publishers have been encoded and distributed through an Intranet server to the network of Swiss libraries (RERO). The XML format and the Unicode standard provide the technical foundations of this activity.

Corpus collection and annotation: For over more than a decade, the group has participated in various initiatives and projects dedicated to the collection and annotation of corpora for public distribution. Collections include the ACL/Data Collection Initiative, the European Corpus Initiative, the UN three language parallel corpus of conference meetings, the MLCC collection of parallel (9 languages) and comparable (6 languages) corpora, and the DIET linguistic test suite.

NLP tools and resources: In support of all of the aforementioned projects, ISSCO has a long-standing tradition in developing base tools and resources for multilingual information processing. These tools cover the range of core NLP processing tasks – segmentation, morphological analysis, part-of-speech tagging and parallel text alignment, including the linguistic resources for English, French, German and Italian.

sa