Informer

Number 29, Winter 2009
ISSN 0950-4974

BCS INFORMATION RETRIEVAL

In This Issue

About Informer

Informer is the quarterly newsletter of the BCS Information Retrieval Specialist Group (IRSG). It is distributed free to all members. The IRSG is free to join via the BCS website (http://irsg.bcs.org/), which provides access to further IR articles, events and resources.

The British Computer Society (BCS) is the industry body for IT professionals. With members in over 100 countries around the world, the BCS is the leading professional and learned society in the field of computers and information systems.

Informer is best read in printed form. Please feel free to circulate this newsletter among your colleagues.

One of the great things about working in IR is the feeling that you're connected to a network of like-minded professionals around the world who share your passion for all things search-related. From the largest nations to the smallest regions, it seems that every corner of the globe has a research group or company working on IR.

So I'm particularly pleased to welcome as guest editor for this issue Monica Landoni, who brings us an overview of current IR research in Switzerland. Monica has drawn together a fascinating set of articles which truly demonstrate the vibrant nature of the IR community in this country.

Closer to home, I'd like to take the opportunity to announce a new feature for Informer: an "Industry News" section. As a complement our regular coverage of IR research, we'll be inviting representatives from start-ups and established businesses in web and enterprise search (e.g. Google, Yahoo, Autonomy, MS/Fast, Endeca and many more) to keep us up to date with all the latest product announcements, company news and case studies.

And of course, don't forget to we now have a BCS IRSG group on LinkedIn. This will be the place where book review offers and other announcements will initially be made, so if you want to be the first to hear, just sign up at http://www.linkedin.com/groups?gid=957047.

In the meantime, best wishes for great 2009,

Tony

Tony Rose, PhD MBCS CEng
Editor, Informer
Vice chair, IRSG
Email: irsg@bcs.org.uk

Introduction to the Special Issue - Information Retrieval in Switzerland

by Monica Landoni

When asked to organise this special issue on Information Retrieval research in Switzerland I went back in time to my experience with Swiss IR, starting from Eurospider, founded in 1995 by the Federal Institute of Technology (ETH) in Zurich, and its focus on cross language IR. Then, in 1999, when I was chair of what was then called BCS IR colloquium in Glasgow, I met, for my first time, researchers from University of Neuchatel presenting their work and finally our collaboration with the university of Geneva and their Content-Based Image Retrieval group, during a successful EU project, Webkit, 2002-2004. My experience of Swiss IR research goes a way back in time, but only since January 2007 I have been part of it as member of the newly established IR group in the faculty of informatics at USI, the University of Lugano.

This issue portrays a very healthy IR research scenario in Switzerland with a variety of different topics being covered and the caveat that there are many more active research groups, both from academia and industry, not to mention a Google research centre in Zurich that could not be included in this special issue because of limitations in time and space. This is indeed not a complete roadmap of IR in Switzerland but a way to stimulate readers into looking further and explore research in this country. For historical reasons I started contacting the very same groups I met in the beginning of my IR adventure. They all enthusiastically agreed to contribute and provided further pointers to active members of the growing IR community.

As for the content of this special issue, contributions span across different topics, from the continuous attention to multi/cross language by the group at Neuchatel university, a necessity in a country with four official languages, to the multimedia dimension that has been expanded to cover a number of creative applications as described in the contribution from the IDIAP institute in Martigny and to produce a range of strategies for multimedia content exploration and annotation by the Viper group at Geneva University. Our group at USI complete this picture by exploring the human, social and affective dimension of IR, via distributed IR, Blog search, Text Mining and evaluation of IIR. I would like to thank every group who contributed to the success of this issue; in order to keep up with stereotypes all Swiss authors submitted on time or earlier, and of course we Italians were to blame for any delay!

Dr Monica Landoni is a senior research fellow at the Faculty of Informatics since January 2007. Previously, she was a lecturer in the Department of Computer and Information Sciences of the University of Strathclyde since 1998 and spent 2005 on a sabbatical position at PARC (formerly Xerox Parc) in Palo Alto, USA. She has authored more than thirty papers in the area of Electronic Publishing and Hypermedia. In the past she has been involved in a number of projects, including the JRC Project "Visual Book", the SHEFC Project "MultiBrowser", and the EU Projects STAMP, PENG and REVEAL-THIS. Her special interests include digital books and libraries and interface design and evaluation. She was Project leader of EBONI (Electronic Books ON-screen Interface).

Information Retrieval Research at the University of Lugano

by Fabio Crestani

The Information Retrieval Group at the Faculty of Informatics of the University of Lugano (also known as Universita' della Svizzera Italiana) was established in 2007, after Fabio Crestani and Monica Landoni joined that university. It comprises of 10 members, including one professor, 3 postdoctoral fellows and 6 PhD students. The group has a strong connection with the University of Strathclyde in Glasgow, where Crestani and Landoni still hold a part-time position.

The group is engaged in research in different areas of Information Retrieval (IR), Digital Libraries and Text Mining. Research is currently funded by the Swiss National Science Foundation, the Haslerstiftung and the European Commission. Currently the four main areas of research are: distributed IR, blog search, text mining of Police data, and evaluation of interactive IR. In the following I will briefly overview the current work.

Distributed IR

Distributed Information Retrieval (DIR), also known as content-based federated search, is concerned with enabling a user to find unstructured or poorly structured documents by their semantic content using natural language queries expressing his/her information needs. The difference with standard IR is that documents are contained in a number of heterogeneous distributed resources, each with its own different retrieval engine.

When so many resources are available, the first information access task the user faces is *resource selection*. This is an ineffective manual task as users are often unaware of the contents of each resource in terms of quantity, quality, information type, provenance and likely relevance. People need accurate automatic resource selection tools to assist them in this task, but resource selection requires accurate *resource descriptions*. These descriptions can be built either manually or automatically, and are a real problem to derive in the case of non-cooperative resources, that is in the case of resources that enable access to their content by querying their search engines, but that do not provide any information about the content of their archives. This is the typical case for resources in the Deep or Hidden Web. A rough estimate put the size of these resources at 400 times that of the Visible Web. Once the resources have been selected and the query forwarded to them, the results returned by each one of them have to be merged by a process called *results fusion*, so that a single ranked list of results is produced and presented to the user, trying to maximise the overall retrieval quality.

"the effectiveness of IR systems can be greatly improved by adapting the system to the specific user tasks and needs"

A large body of research in the last 10 years has shown that the effectiveness of IR systems can be greatly improved by adapting the system to the specific user tasks and needs. This enables to satisfy the user information need taking into consideration the context in which the user need is placed, so personalising the interaction with the system. While a large body of work already exists for personalised and context dependent IR, the DIR research area has not yet considered issues of personalisation. One of the issue that make this kind of research hard is the difficulty in obtaining large personalised dataset, comprising documents, queries and relevance assessments related to one specific individual, rather that a group or an unspecified individual (e.g. TREC). The ADIR and ADIR+ projects, funded respectively by the Swiss National Science Foundation and the Haslerstiftung, aim at making personalisation research possible for DIR, studying how it is possible to build personalised test collections from data freely available, such as, for example, data available from del.icio.us. This work complements that carried out in the EPSRC project Nemo that is

concerned with designing, implementing and testing models of personalised content-based federated search that will be applied to retrieving information from the Deep Web. Together, these three projects are designing, implementing and testing advanced resource description, resource selection and results fusions methods that can be automatically personalised to the user task and user needs and that are specifically designed to access information held in non-cooperative and heterogeneous resources information resources, like those found in the Deep Web.

Blog Search

The rise on the Internet of blogging, the creation of journal-like web page logs, has created a highly dynamic subset of the World Wide Web that evolves and responds to real-world events. Indeed, blogs (or weblogs) have recently emerged as a new grass roots publishing medium. The so-called blogosphere (the collection of blogs on the Internet) opens up several new interesting research areas. One of these areas is related to the fact that many blogs are created by their authors, not intended for any sizable audience, but purely as a mechanism for self-expression. Therefore, blogs often reflect the bloggers' opinions, ideas and emotions related to their daily life or to world events. In the OpiSoft and XMI projects, supported by the Swiss National Science Foundation we aim at exploring the information seeking behaviour in the blogosphere. In particular, we aim at designing, implementing and testing new models for opinion finding and opinion polarity detection (OpiSoft) and advanced models of evidence combination (XMI) that will be applied to different applications, including blog distillation.

Current approaches to opinion finding and opinion polarity detection rely heavily on the statistical analysis of the text of the blog posts to detect the semantic content, the presence of an opinion and estimate its polarity (positive, neutral, or negative). The two most effective approaches to the problem either build automatically a weighted dictionary of opinionated terms from a training sample of opinionated documents or use a pre-compiled list of subjective terms and indicators to generate an opinionated score for each document of the collection. In both cases user or expert input is extremely important, either in the form of relevance assessments used to build the training sample or in the selection of terms for the pre-compiled list. However, it is well known in IR that users, or even so-called experts, can provide only imprecise and vague feedback for relevance assessment, in the best of cases. Feedback on the presence of an opinion and its polarity is much more difficult than relevance feedback and is anyway subjective by its very nature. Statistical techniques are therefore not the best techniques to handle this level of imprecision and vagueness. Soft Computing provides a score of techniques that have been expressively designed to handle imprecision, vagueness, partial truth, and approximation. Therefore techniques falling under the umbrella of Soft Computing seem to be better suited to exploring the information seeking behaviour in the blogosphere and, in particular to the tasks of opinion finding and opinion polarity detection than purely statistical techniques.

“Soft Computing techniques seem to be better suited to exploring the information seeking behaviour in the blogosphere”

Nevertheless, none of the approaches to these tasks attempted in the TREC 2006 and 2007 blog tracks make effective use of Soft Computing techniques. In OpiSoft, working in close collaboration with the European Commission COST Action IC0702 on “Combining Soft Computing Techniques and Statistical Methods to Improve Data Analysis Solutions”, we are investigating the combination of Soft Computing and statistical techniques for feature selection for the tasks of opinion finding and opinion polarity detection, in order to create automatically a weighted dictionary of opinion-bearing terms that will be used to determine, in a query independent way, the topical-opinion of a blog post. This approach will be more robust to imprecision and vagueness present in training data and will be more easily interpretable that pure statistical approaches. This latter property is very important for media analysts as it enables them to understand the results and check them for plausibility.

Text Mining

Consider the following system that integrates into existing police databases. When a member of the public reports a crime such as vandalism, assault, burglary or other theft, a description is recorded. This will include what happened using both free text and ticking of boxes as well as the location, time and date. The system will then be able to then relate this unsolved crime to those already solved to prioritise suspects based on both the location of the crime and also the method employed. The system reveals not only known offenders who may have done it but also how likely each one is to be the culprit; this can be done in real time. Furthermore the police can enter subjective information to refine this process such as an individual is known to be particularly active in this area. If a clear suspect is identified, the system will indicate reasons for this by identifying in a visual way the features linking his past crimes to the new one. This information can help provide reasonable suspicion to obtain an arrest or search warrant from a court since past criminal activity is admissible here.

"The system reveals not only known offenders but also how likely each one is to be the culprit"

Often no obvious suspect emerges and this may be because the offender is hitherto unknown to the police. In this case the system is used to profile the offender and predict what likely attributes he or she may have such as gender, age, occupation (e.g. whether working or not), ethnic group and the likely distance travelled from the home base. Again the system identifies visually those features of the crime that lead to these inferences e.g. an untidy search suggests a burglar who has travelled. However such relationships may not be universal and only apply in a particular region or city; they will change over time. Thus each police force uses its own current data to recalibrate these models on a regular basis.

From time to time the police officers and their support staff take a more strategic view of the crimes being committed in their district. They can identify that specific behaviours are related by specific groups of individuals in the locality. The police can also observe trends over time revealing patterns and which were otherwise very difficult to detect manually.

All the examples given above can be achieved with text mining technologies developed by the iMOV project, funded by the EPSRC at the University of Strathclyde. Data recorded in police databases differs greatly between forces but certain common features exist. Most of the information about how a crime was committed is recorded as unstructured of semi-structured text. Using IR technologies and in particular continuous language models like those developed by iMOV, it is possible to link crimes, extract crime or estimate individual or group characteristics, that can be traced and modelled over time.

The results of the iMOV project are now starting to being applied and extended to analyse Swiss Police data, with particular regards to trans-border crimes to which Swiss public opinion is particularly sensitise, since Switzerland joined Schengen in December 2008. The challenges of this work will extend to deal with issues related to data integration (i.e. the databases of the Italian and Swiss Police) and multi-language text mining (i.e. data in Italian, French and German).

Evaluation of Interactive IR

Evaluation of interactive IR systems is a growing research area. The definition of novel methodologies to support design and running of user centred experiments is our crucial contribution to this research area. In particular, novel methodologies to measure usability of Interactive IR systems have been our research focus. We started from a previous project that produced a framework where classic IR measures like precision and recall, properly adapted to dynamic nature of web pages, were sitting next to more web specific factors such as: coverage, relevance ranking looking at how effective ranking is in guiding user to the most relevant pages and so avoiding information overload and accessibility, plus a directory assessment measure specific for electronic directories such as Yahoo!

This framework was used to run a large usability study looking at the way users, a large sample of under and post graduate students across disciplines, choose and use

search engines based on usability features. An evaluation experiment had been set up following the guidelines defined by previous study and findings showed a clear preference of users for Google in terms of usability and trust. Our current challenge, with new SNSF (Swiss National Science Foundation) funded project DEDUCE, is to use our experience to design novel methodologies to involve and support children when evaluating search tools designed for them. Children have specific needs and requirements that need to be taken into account when organising usability studies. For instance children could find it more difficult than adults to provide valid feedback since verbal communication, both in understanding and formulating sentences, so that evaluators have been exploring alternative methods, ranging from interpreting free drawings to using collections grids with "smileys" instead of grades. Also recurrent in literature is the recognition of the importance of facilitators with the caveats of how influential and delicate their role can be in terms of inevitable bias, as children have a natural tendency to please.

"An evaluation experiment showed a clear preference of users for Google in terms of usability and trust."

Besides, children have specific problems in expressing their feelings in terms of satisfaction or even just liking and disliking. A third element emerging from previous studies is the discrepancy between reported and observed usability when children are asked to provide subjective feedback versus direct observation. In general, HCI experts working with children have been very imaginative and innovative in coming up with alternative and creative techniques and tools to involve children and to overcome the linguistic barrier, something that could be beneficial to adults too, where instead evaluators have kept a more conservative approach.

Project DEDUCE looks at the affective side of evaluation and possible ways to take that into account while running usability studies. Emotions and the ability to keep track of them explicitly play a critical role in evaluations. Traditionally there have not been any specific instruments for gathering this type of feedback, so that it could be used as a filter when interpreting feedback data. DEDUCE aims at filling this gap by designing and delivering tools to facilitate user participation in experiments raising level of engagement and interest. This should results in more accurate and representative feedback to be gathered by evaluators that will in turn allow for better understanding of users needs.

Prof. Fabio Crestani is a Full Professor and Director of PhD studies at the Faculty of Informatics of the University of Lugano since January 2007 and is the leader of the Information Retrieval Group. Previously he was a Professor at the University of Strathclyde in Glasgow (2000-2006), where he still holds a part-time Professor position. Crestani has over 20 years of experience in research in Information Retrieval, Digital Libraries and Text Mining. In these areas of research he has co-authored over one hundred refereed publications and edited nine books.

Contact information: Prof. Fabio Crestani, Faculty of Informatics, University of Lugano, CH-6900 Lugano, Switzerland. Email: fabio.crestani@unisi.ch

The IR Group at Univ Neuchatel

By Jacques Savoy

The main objective of the IR group at University of Neuchatel (Switzerland) is to propose, implement and evaluate various linguistic tools able to enhance IR performance when various languages other than English are involved. Also included among our aims are tools able to process bilingual IR (queries written in one language and documents in another) and multilingual IR (information items targeted are written in different languages).

These three aspects target a variety of scenarios. In multilingual countries such as Canada, Switzerland and India or in Europe in general, users search information on many different topics (law, economy, culture, education, leisure and travel), and the search context may be large international organizations or companies (e.g., WTO, European Parliament, Patent Office, etc.). While some users are perfectly bilingual, other can read documents written in another language but cannot formulate a query in that language or, at least they cannot provide reliable search terms in a form comparable to those found in the documents being searched. Moreover the target documents may be written in other languages or contain information in non-textual formats such as images, graphics and statistics. In all these cases, IR approaches must consider the multilingual nature of searchers, searches and targeted information.

Compared to the English language, other Indo-European languages make use of more complex morphologies. In the Spanish language for example the noun friend could be written as "amigo" (masculine) or "amiga" (feminine). While in some contexts distinctions such as these may be very important, when a query includes one form and relevant item being searched contains another, searchers expect the IR system to match the two related forms. Other distinctions involve diacritics (e.g., "leão" (lion in the Portuguese language), "Äpfel" (apples in German)), used to facilitate pronunciation and reflect exact meanings (e.g., "cure" (presbytery) and "curé" (parish priest) in the French language). Other variations applied in certain natural languages include suffixes to indicate number ("amigos" vs. "amigas" (friends)) and grammatical case (e.g, as in Latin with "rosa", "rosae" and "rosarum"). Thus, in order to increase retrieval effectiveness and make a better match between a user's query and pertinent documents, all these suffixes should be removed. Natural languages also make use of other alphabets, such as Cyrillic for the Russian language. In these cases IR methods need to interpret various encoding standards (e.g., ISO-8859-5, UTF-8 or *de facto* standards such as Windows-1251).

Languages not belonging to the Indo-European family are found in Europe, including Hungarian and Finnish. Both feature a greater number of grammatical cases (over 18 for Hungarian and 15 for Finnish). While numerous cases can make matching query request and documents more problematic, only the Finnish language causes any difficult problems, due to its frequently varying stems and various suffixes added. For example the noun "matto" (carpet, in the nominative case) becomes "maton" in the genitive (one "t" has been removed) or "mattojen" in the genitive plural (a "j" has been added).

"IR approaches must consider the multilingual nature of searchers, searches and targeted information"

Other European languages may also involve problems related to compound constructions (e.g., fireman, software). While in French the various components can be clearly identified (e.g., "porte-avions" (aircraft carrier) or in "chemin de fer" (railway)), variations in number and compound constructions can generate difficult problems when conflating the singular and plural forms to the same stem. For example in Italian the term "capoufficio" (chief clerk) becomes "capiufficio" in the plural form. In German the situation becomes more complex, especially from an IR point of view. Since compound construction is widely used and has many different forms, we may find

one form in a query ("Sicherheit am Computer" (computer security)) and yet another in documents ("Computersicherheit").

In other cultures people have adopted radically different writing systems, such as the ideograms found in the Japanese and Chinese languages. Moreover, word boundaries in these two languages provide IR systems with additional problems. For the Chinese sentence "我不是中国人" (I'm not Chinese), the IR system needs to known that the expression comprises the following words: 我 (I/me), 不 (not), 是 (to be) and 中国人 (Chinese person). Thus for IR systems effective automatic segmentation is a real concern for Chinese (and for Japanese to a lesser extent).

When considering bilingual or more generally multilingual IR systems, an automatic translation device is needed to cross the language barrier. Various solutions have been suggested, including bilingual machine-readable dictionaries (MRD), machine translation (MT) systems, or statistical translation models based on parallel corpora.

When employing MRDs to process an input sentence, we need to generate corresponding lemmas (e.g., "I saw men with a saw" → "I, to see, man, with, saw"). Moreover due to the dictionary's limited coverage we need to handle out-of-vocabulary (OOV) problems. Similarly, it could prove helpful to single out proper nouns and acronyms and translate them using a special dictionary (e.g., Gorbachev-Gorbachov, UNO-ONU). Moreover, certain input words could be ambiguous and MRDs may suggest more than one translation (e.g., the word "bank" could refer to a river or a financial institution).

As a second translation tool we could utilize machine translation (MT) systems. Although normally easy to use, they tend to perform poorly when translating entire documents and the resulting translated text is usually difficult to read. Moreover, when queries input are not complete sentences, or are simple expressions that make no grammatical sense, a translation could not be easily generated by the system.

As a third possibility, translation processes could be based on parallel or comparable corpora, presuming of course that they are available. For certain pairs (e.g., Finnish/Hindi) the required corpora could clearly be more difficult to find. The performance of these statistical translation approaches would however depend on very important factors such as source quality and size, not to mention various other cultural, thematic and time differences. For example the term "portable phone" must be translated as "portable" in France, "mobile" in Belgium, "cellulaire" in Canada or "natel" in Switzerland. When considering these various translation devices, uncertainty attached to the translation process, along with that attached to the match of queries and documents must be taken into account.

"Adding synonyms extracted from specific databases usually does not improve retrieval performance in biomedical domain"

The IR group at UniNE is also interested in analyzing retrieval effectiveness in dedicated domains (e.g., biomedicine, social science, blogs). When an IR system is deployed within a specific domain, we hope to improve its retrieval effectiveness by applying manually assigned descriptors, employing a specialized thesaurus or other dedicated terminological databases. Adding synonyms extracted from specific databases (LocusLink, AcroMed, Swiss-Prot, Gene Ontology, or MeSH) usually does not improve retrieval performance in biomedical domain, as is the case of related terms given by commercial search engines (e.g., Google). Manually defined descriptors do tend however to improve search quality (from 5% to 10%), and in the context of a bilingual search the translation of submitted queries can be improved by applying a domain-specific bilingual thesaurus. In the biomedical domain the lexical variants of search keywords is another important feature worth considering. A protein or a gene name may for example be known under different names ("NRAMP1," "SLC11A1") or spelling forms such as "Lsp1alpha," "lsp-1-a," "lsp-1a", etc., and a smart search system should therefore consider replacing "alpha" by an "a" or "2" by "ii", etc.

IR research should also extend beyond traditional limits. Unlike traditional document collections, the web and in particular the blogosphere is more subjective, and is also characterized by more diverse document structures and writing styles. Although the blogosphere may indeed contain objective information (facts), we are also interested in developing search tools to find answers based on opinions rather than relevant factual information. As such, relevant answers to the request "iPhone" may include factual and technological information (relevant but without opinions) but also more personalized (and subjective) features about the product (why it is useful, complaints, drawbacks of a specific function, personal experiences, etc.). Thus, instead of simply returning a ranked list of documents, we need to subdivide them into factual information items and information units expressing positive, mixed or negative opinions pertaining to the target entity.

Prof. Jacques Savoy is a Computer Science professor at University of Neuchatel (Switzerland). After having obtained his Ph.D., he jointed the Computer Science faculty at the University of Montreal (1987-93). His research interests include information retrieval for languages other than English, along with multilingual and cross-lingual IR. For many years he has participated in various evaluation campaigns (CLEF, NTCIR and TREC).
E-mail: Jacques.Savoy@unine.ch

Information Retrieval in Geneva: the *Viper* research group

By Stephane Marchand-Maillet

For more than 10 years, the *Viper* research group of the Computer Science Department of the University of Geneva has been working on advancing the field of Information Retrieval. The group is part of the Computer Vision and Multimedia Lab (CVMLab), originally specializing on Computer Vision problems. From there, the *Viper* group has spawned, with an initial motivation of transferring the lab's expertise in Computer Vision for reaching efficient Content-based Image Retrieval (CBIR). It now addresses a wider range of scientific challenges, from Content-based Multimedia Information Retrieval to Multimedia Content exploration and annotation via Knowledge Management and Discovering developments, all based on formalized and robust strategies, essentially derived from Machine Learning.

1. From classical information retrieval to multimedia information retrieval

Text information retrieval has a long history of development since the initial SMART system proposed in the late 60's. It has now reached a rather mature stage and exported itself in general contexts such as Digital Libraries and even expanded successfully to the Web context in the late 90's. Multimedia Information Retrieval has a younger history and is less attached to commercially successful systems. Multimedia Information here generalizes text to all other types of information, from visual content (*e.g.* images, drawing, 3D), audio content (*e.g.* speech, music) to the combination of all such as Web pages (*e.g.* considered as a compound of images and text) or video (*e.g.* a news broadcast as a compound of a spatiotemporal visual stream, speech, possibly music and text caption or teletext).

Multimedia Information Retrieval offers several challenges:

- First, most of the multimedia content is not text. The automated processing, indexing and retrieval of that content therefore calls for advanced techniques for the automated interpretation of audio, visual contents and alike. Hence, classical limitations of fields such as Computer Vision or Speech Understanding are directly or partly inherited by the field of Multimedia Information Retrieval. The lack of suitable interpretation capabilities is accepted as the *semantic gap*, the difference between the best interpretation reached from the physical features of the information at hand (color, pitch,...) and what a human operator would capture from that content ;

- Further, Multimedia Information is often complex and composed of several unit streams of information (e.g. audio and video streams), referred to as *modalities*. Hence, Multimedia Information Retrieval adds to the above interpretation challenge the problems of mixing correctly various simultaneous streams of information. This problem is that of Information Fusion where information gathered from multiple modalities should be fused into one unique characterization of the multimodal content. It is also generally recognized that further reductions of the semantic gap may only come from a good understanding of fusion mechanisms;

- The complexity of Multimedia Information retrieval also influences the efficiency of its handling. While text data may be stored and handled efficiently and document unit sizes are within acceptable ranges for storage (memory and disks), transfer, and analysis, this is less so with Multimedia Information. In the case of video for example, the size of one document only may well be equivalent to that of a complete text corpus. The scalability of processing and access strategies therefore become critical in this context;

- Multimedia Information Retrieval follows its textual counterpart in that it should be interactive. In most usage scenario, the acceptable delay between a query and a response is within the same range of that for textual systems. To scalable indexing and access, MIR therefore adds fast and accurate query processing. It is further recognized that proper user interaction is critical to a correct semantic interpretation of his/her queries via the concept of *relevance feedback*, where the user "trains" the system to his/her view of the information need;

- Finally, while text Information is interacted with easily through summaries for example (*e.g.* snippets, tagclouds), the case of MIR again makes the interaction more complex. Temporal streams (video, audio) should also be handled via summaries to enable proper relevance feedback acquisition. For example, if provided with a page of 20 video documents of 5 minutes each as initial result, the user may have to spend about 30 minutes before returning some list of positive and negative matches to better the results. Clearly, video summaries such as the joint use of keyframes and tagclouds are of help here;

An initial temptation is that of transferring directly the experience text IR to MIR. Models for text IR such as the classical TF-IDF indexing based on counting word occurrences in text documents are rather simple and have largely proved their efficiency. It is therefore natural to find MIR systems directly mapping these models onto other content. The Viper group has done so eight years ago when releasing the GNU Image Finding Tool (GIFT) as the first complete GNU opensource package[1] to enable Content-based Image Retrieval. Here, image parts characterized by color and texture at several scales become words and are indexed as text would be. The query mechanism is the Query-by-Example (QBE) paradigm where the user shows positive and negative examples of the sought content as a query. In this context, the Relevance Feedback mechanism is just a system to construct incrementally the proper query as the appropriate set of positive and negative examples.

[1] Still available at http://www.gnu.org/software/gift

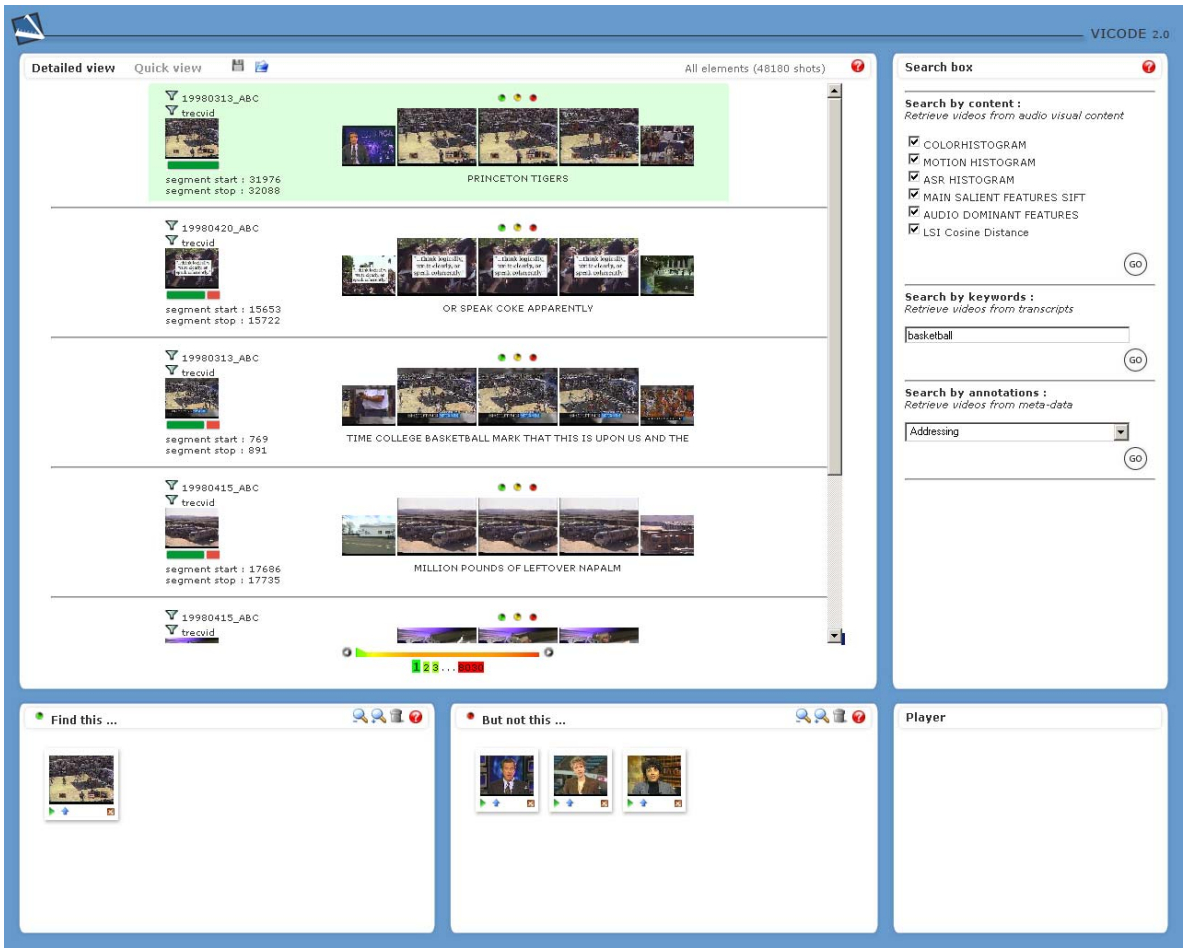

Figure 1: Viper's ViCoDE: a video Query-by-Example system

2. Information retrieval as a learning problem

Models designed for text IR adapt from relevance feedback, essentially using linear or Bayesian learning. As result, their learning capabilities are limited to well-behaved classes. That is groups of relevant documents that are rather similar in their content. This is largely not the case for MIR where relevance is a complex notion. For example, the various possible visual aspects of a given object make it impossible for any visual feature to group all views of that object into one compact cluster. This motivates the search for advanced learning techniques and the Viper group has developed an expertise in modelling retrieval problems using Machine Learning tools. We have developed a fully operational model considering positive and negative examples provided as query as samples of the positive (relevant) and negative (non-relevant) classes. The challenge is therefore to interactively train a flexible Machine Learning algorithm with few provided samples of a (potentially small) positive class against provided samples of a (potentially large and diverse) negative class. Our strategy is based on non-linear learning using generalization algorithms such as Support Vector Machines (SVM), Discriminant Analysis (xDA) and Boosting. We have adapted a number of techniques to this interactive unbalanced context.

Learning further enables various formal ways of performing efficient Information Fusion. We have employed either early or late fusion strategies to create image and video search engines. One recent achievement of the *Viper* group is the development of a large-scale multimodal retrieval engine called CMSE (Cross-modal Search Engine[2]) that acts over combination of audio, visual and text content. It has successfully been applied to the retrieval of Cultural Heritage (CH) material (essentially

[2] CMSE: Cross-modal Search Engine: http://viper.unige.ch/cmse

Figure 2: *Viper*'s browser for Cultural Heritage data in the MultiMATCH project. The browser adds visual similarity (vertically) to collection browsing (horizontally)

pictures of art pieces with descriptions from CH websites and museums) in the MultiMATCH European project[3]. The classical Open Computer Vision library is used to extract features such as colour, texture, face occurrence and our engine mixes accurately all needed modalities at query time. In such a context, the rather limited QBE paradigm is completed with keyword-based queries, the currently most natural way to formulate queries.

3. Extending the information access strategy

Content-based Retrieval is core for enabling scalable information access. However, not all scenarios are query-based. It may well be the case that the user simply wishes to explore the corpus at hand for *e.g.* opportunistic browsing (*i.e.* expressing his/her information need along the collection navigation). The Viper group is active in developing such a framework and has proposed the Collection Guiding principle as a way to guide the user within the virtual collection. That is, the collection of multimedia documents is pre-processed and a visit strategy for that (originally unorganized) content is automatically proposed. Our base strategy relies on viewing the collection as a network of inter-connected items and a visit over this network is modelled as a path reaching optimality. For example a Travelling Salesman Tour over the similarity networks will provide the user with a path based on "smooth" content transition from one item to the next. Adapted browsing interfaces may therefore be developed to make the best of that data organization (see e.g. Figure 2).

Organizing the content is not only good for its exploration, it also helps its description (group annotation, summarization, sampling). As active participant of the PeTaMedia[4] European Network of Excellence, the *Viper* group develops mining strategies for semi-automated annotation and content management over

[3] MultiMATCH : MultiLingual/MultiMedia Access to Cultural Heritage. http://www.multimatch.eu

[4] PeTaMedia : Peer-to-Peer Tagged Media : http://www.petamedia.eu

distributed networks. Our current developments focus on exploiting long-term interaction with multimedia content. We particularly exploit the context of social communities (in relation to the Web 2.0) to extract semantics from user interaction logs and project them onto shared items.

4. Summary

Providing efficient access to large-scale multimedia collections is a multi-faceted challenge, from content processing and indexing to annotation and exploration issues. We view all these issues as necessary and complementary themes to address to reach yet unavailable accepted performances. The *Viper* group is addressing these issues globally and has already provided software solution materializing its findings.

5. Further reading:

The Viper group's scientific publication list is available at
http://viper.unige.ch/publications

Selected items:

[1] Bruno, E., Moënne-Loccoz, N., & Marchand-Maillet, S. (2008). Design of multimodal dissimilarity spaces for retrieval of multimedia documents. IEEE Transactions on Pattern Analysis and Machine Intelligence, 30(9), 1520-1533.

[2] Kludas, J., Bruno, E., & Marchand-Maillet, S. (2008). *Exploiting Synergistic and Redundant Features for Multimedia Document Classification*. In 32nd Annual Conference of the German Classification Society - Advances in Data Analysis, Data Handling and Business Intelligence (GfKl 2008), Hamburg, Germany.

[3] Marchand-Maillet, S., & Bruno, É. (2005). *Collection Guiding: A new framework for handling large multimedia collections*. In Proceedings of the First Workshop on Audio-visual Content And Information Visualization In Digital Librairies, (AVIVDiLib05), Cortona, Italy.

[4] Moënne-Loccoz, N., Janvier, B., Marchand-Maillet, S., & Bruno, E. (2006). Handling Temporal Heterogeneous Data for Content-Based Management of Large Video Collections. *Multimedia Tools and Applications*, 31, 309-325

[5] Morrison, D., Bruno, E., & Marchand-Maillet, S. (2009). Capturing the semantics of user interaction: A review and case study. In *Web Emergent Intelligence*, Springer.

Dr. Stephane Marchand-Maillet received his PhD on theoretical image processing from Imperial College, London in 1997. He then joined the Institut Eurecom at Sophia-Antipolis localization and recognition. Since 1999, he is Assistant Professor in the Computer Vision and Multimedia Lab at the University of Geneva, where he is working on content-based multimedia retrieval as head of the Viper research group. He has authored several publications on image analysis and information retrieval, including a book on low-level image analysis. He and his group are currently involved in several joint international efforts for benchmarking content-based multimedia retrieval systems (including the Benchathlon and ImageCLEF). URL: http://viper.unige.ch Email: Stephane.Marchand-Maillet@unige.ch

Multimedia Meeting Processing and Retrieval at the Idiap Research Institute

By Andrei Popescu-Belis and Alessandro Vinciarelli

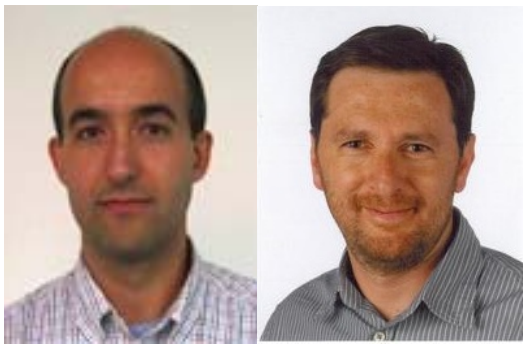

1. Multimedia Retrieval from Meetings

Indexing multimedia data such as meeting recordings can be made richer as progress in automatic processing of social interactions allows for annotation of new aspects of the data. The Idiap Research Institute has thus successfully applied its long-standing know-how in the field of perceptual artificial intelligence to the development and evaluation of meeting-related technology such as smart meeting rooms, media file servers, meeting browsers for accessing past meetings, and meeting assistants for supporting ongoing meetings. Idiap has also pioneered a wide range of approaches for the automatic understanding of human-human interactions, including the recognition of people roles, the detection of dominant individuals, the identification of collective actions and, most recently, the analysis of social signals (see Section 3 for more details).

Starting from the M4 European project[5], continuing with the (IM)2 Swiss National Centre for Competence in Research[6], and arriving at the AMI and AMIDA European Integrated Projects[7], automatic meeting processing and retrieval, for both offline and online applications, was gradually shaped and promoted as a benchmark and demonstration opportunity for mono- and multimodal processing technologies, as well as for commercial applications that are currently making their first, successful steps in the market. For instance, Klewel, an Idiap spin-off company[8], now provides a conference recording service allowing users to browse recordings of past events and search through them based on slide content.

To support this research context, Idiap has developed a complex infrastructure, including an instrumented meeting room which captures and synchronizes several audio and video streams – multiple time-aligned microphones and cameras, plus beamer, whiteboard and written notes – and a media file server which offers secure storage and access to large volumes of data and related manual or automatic annotations. A large collection of meeting recordings, with transcriptions and annotations, was made available to the scientific community in order to support further research and development in this area[9].

In the two following sections, we will focus on two recent promising research directions that build upon the expertise accumulated in the above described activities, while opening alternative and unexplored research avenues: Automatic Content Linking (Section 2) and Social Signal Processing (Section 3).

2. An Automatic Content Linking Device

An *Automatic Content Linking Device (ACLD)* is a system that provides tailored access to documents or fragments of recorded meetings, based on ongoing discussions. The device is intended to work online during a meeting, and to display documents potentially relevant to ongoing discussions in the same moments as these take place. Indeed, participants in a meeting often mention documents containing facts that are currently discussed, but only few documents are, in general, at hand. Search

[5] M4: Multimodal Meeting Manager, http://www.dcs.shef.ac.uk/spandh/projects/m4, 2002-2005.
[6] (IM)2: Interactive Multimodal Information Management, a Swiss National Centre of Competence in Research coordinated by Idiap, http://www.im2.ch, 2002-2009.
[7] AMI: Augmented Multiparty Interaction, 2004-2007, and AMIDA: Augmented Multiparty Interaction with Distance Access, 2007-2009, two EU IPs coordinated jointly by Idiap and the University of Edinburgh, http://www.amiproject.org.
[8] Klewel, http://www.klewel.com.
[9] The AMI Meeting Corpus: about 100 hours of recordings available at http://corpus.amiproject.org.

for the right piece of information could be performed using an information retrieval or document management system, but performing such operations during the meeting interrupts significantly the discussions. Moreover, even when all necessary documents are available, still recordings of past meetings are not at disposition, not to mention an effective device for searching through them. Alternatively, when browsing through recordings of previous meetings, users do not have the time to search also for additional information among the meeting documents. Therefore, an ACLD is expected to constantly suggest documents (including excerpts of previous meetings) that are potentially relevant to the ongoing discussion. Meeting participants are free to ignore the suggestions, but they can profit of them to enhance the discussion.

The above scenario is broadly related to the just-in-time retrieval approach, and to other speech-based information retrieval solutions. The originality of the ACLD lies in its robust approach to meeting retrieval, in the connection of the retrieval functionality to a data bank of meetings and documents which is enriched through meeting processing, and to the use of a context model to smooth the retrieval results in time.

"An ACLD is expected to constantly suggest documents that are potentially relevant to the ongoing discussion"

The ACLD developed at Idiap[10] currently demonstrates the basic concept of tailored access to a group's history using a set of four meetings from one of the groups recorded in the AMI Meeting Corpus. Although its primary use is for live meetings, the concept applies also to offline meeting browsing. In this case, the ACLD plays a meeting recording and processes it as a live meeting, treating segments from other, previously recorded, meetings and associated documents as a history to be linked. The final result of the ACLD is a number of project documents and fragments of previous meetings, relevant to the discussions, including reports, emails, and presentations given during other meetings, plus segments extracted from other meetings by segmenting them into 3-minute chunks.

The ACLD performs searches at regular intervals, using a search criterion based on the words that were recognized automatically from the meeting discussion, thanks to online automatic speech recognition (or offline in the case of browsing). If some pre-specified terms or keywords are recognized, then they receive greater weight in the subsequent query. The results are presented as a list of document names ordered by relevance (possibly empty), to which a user interface offers quick access. These functionalities are supported by modules that communicate through a subscription-based client/server architecture, called the Hub, which allows the connection of heterogeneous and remote software modules, and ensures that data exchange is extremely fast. "Producers" of annotations send triples to the Hub, which are received by the "consumers" that subscribed to the respective types; consumers can also query the Hub for past annotations and metadata about meetings.

3. Social Signal Processing

Social Signal Processing (SSP) is the new, emerging domain aimed at automatic analysis and understanding of nonverbal behaviour in social interactions[11]. SSP applies ***signal processing*** techniques, typically used to analyze and abstract the content of multimedia data, to sense and understand ***social signals***, i.e., complex aggregates of nonverbal behavioural cues (e.g., facial expressions, postures, gestures, vocal outbursts, etc.) through which we express our attitude towards others and social situations (e.g., agreement and disagreement, aggressiveness, empathy, politeness, etc.).

[10] Popescu-Belis A., Boertjes E., Kilgour J., Poller P., Castronovo S., Wilson T., Jaimes A. & Carletta J., "The AMIDA Automatic Content Linking Device: Just-in-Time Document Retrieval in Meetings", *Machine Learning for Multimodal Interaction V (Proceedings of MLMI 2008, Utrecht, 8-10 September 2008)*, LNCS 5237, Springer, p.272-283.

[11] See A.Pentland, "Social Signal Processing", IEEE Signal Processing Magazine, 24(4):108-111, 2007, and more recently A.Vinciarelli, M.Pantic, H.Bourlard, A.Pentland, "Social Signal Processing: State of the Art and Future Perspectives of an Emerging Domain", Proceedings of the ACM International Conference on Multimedia, pp. 1061-1070, 2008.

SSP is of interest to many domains (e.g., Human-Computer Interaction, Computer Mediated Communication, etc.), but its role is especially important in multimedia retrieval. The reason is that social interactions are an ubiquitous subject in multimedia material and their automatic understanding is a key step in bridging the semantic gap: radio programs are typically based on speakers interacting with one another (interviews, talk-shows, etc.), television programs rarely show something different from people involved in social interactions (movies, talk-shows, debates, etc.), and even application specific and less common types of data (meeting recordings, surveillance material, call centre conversations, etc.) revolve, in general, around social interactions between different individuals.

"Social Signal Processing is the emerging domain aimed at automatic analysis and understanding of nonverbal behaviour in social interactions"

It is therefore not surprising to observe in the literature a recent wave of works where social interactions are automatically analyzed to serve as a cue for content understanding. In particular, several approaches use the automatic recognition of roles (one of the most important and active areas in SSP) for structuring, summarization, and semantic segmentation of data like movies, news, and talk shows[12]. In such works, SSP is proposed as a new approach to tasks typically performed with other methods (most frequently text analysis techniques applied to speech transcriptions), but several contributions go further and actually propose the analysis of social interactions as a totally alternative way of looking at the multimedia indexing problem, at least for those data where interactions are the most important aspect of the content. These pioneering attempts include in particular several efforts made at Idiap to represent meetings in terms of social verticality dimensions[13], collective actions being carried on[14], and role distribution across meeting participants[15]. It is too early to say how far these new approaches will go, but for sure they will contribute to traditional content representation techniques (mostly based on text indexing) with a wide spectrum of social cues currently neglected in multimedia retrieval approaches.

Andrei Popescu-Belis is a Senior Researcher at the Idiap Research Institute in Martigny. His research interests are language processing by computers, language and multimodal resources, and the evaluation of linguistic and interactive systems. He currently works within the Swiss NCCR in Interactive Multimodal Information Management (IM2) and the European IP on Augmented Multiparty Interaction (AMIDA).

Alessandro Vinciarelli received his M.Sc. in Physics at University of Torino (Italy) in 1994, and his PhD in Applied Mathematics at University of Bern (Switzerland) in 2003. Between 1994 and 1999 he has worked in several companies and research laboratories in Italy and the United States (Accenture, Italian Institute for Bioelectronics, Elsag Electronics). He has joined IDIAP Research Institute in 1999 as a PhD student and, since 2005, is Senior Researcher in the same institution.

[12] See in particular S.Favre, H.Salamin, J.Dines and A.Vinciarelli, "Role recognition in multiparty recordings using social affiliation networks and discrete distributions", in Proceedings of International Conference on Multimodal Interfaces, pp. 29-36, 2008, and references therein.

[13] See D.Jayagopi, S.Ba, J.-M. Odobez and D. Gatica-Perez, "Predicting two facets of social verticality in meetings from five minutes time slices and nonverbal cues", in Proceedings of International Conference on Multimodal Interfaces, pp. 45-52, 2008, and references therein.

[14] See I.McCowan, D.Gatica-Perez, S.Bengio, G.Lathoud, M.Barnard, and D.Zhang, "Automatic analysis of multimodal group actions in meetings", IEEE Transactions on Pattern Analysis and Machine Intelligence, 27(3) :305-317, 2005.

[15] See N.Garg, S.Favre, H.Salamin, D.Hakkani-Tur, and A.Vinciarelli, "Role recognition for meeting participants: an approach based on lexical information and social network analysis", proceedings of ACM International Conference on Multimedia, pp. 693-696, 2008.

Forthcoming Events

Edited By Andy MacFarlane

Search - the Science Making Tomorrow's World
An evening talk on the past present and future of search – part of science week 2009. This is a joint event with the BCS North London Branch.
BCS Covent Garden, London, U.K., 9th March 2009.
http://www.nlondon.bcs.org/

Unlocking Audio 2: Connecting With Listeners
Of interest to members interested in issues such as knowledge organisation for sound archives or audio retrieval.
British Library Conference Centre, London, U.K., 16th to 17th March 2009.
http://www.bl.uk/unlockingaudio

First International Workshop on Aspects in Evaluating Holistic Quality of Ontology-based Information Retrieval (ENQOIR 2009) – part of APWeb-WAIM 2009 conferences.
Of particular interest to members working in the area of knowledge organisation and IR.
Suzhou, China, 1st to 4th April 2009.
http://events.idi.ntnu.no/enqoir09/

Second International Workshop on Location and the Web (LocWeb 2009), in conjunction with CHI 2009
Of interest to members working in the areas of web based Geographical and/or mobile search.
Boston, USA, 5th April.
http://ifgi.uni-muenster.de/locweb2009/

31st European Conference on Information Retrieval (ECIR 2009)
The IRSG's annual conference focused on all aspects of IR.
Toulouse, France, 6th to 9th April 2009.
http://ecir09.irit.fr

18th International World Wide Web Conference (WWW 2009)
The top web conference of interest to members working in the area of web search, including a workshop on semantic search.
Madrid, Spain, 20th to 24th April 2009.
http://www2009.org/

The 17th European Symposium on Artificial Neural Networks Advances in Computational Intelligence and Learning (ESANN 2009) - special session on Preference Learning
Of interest to members working areas such as learning to rank etc.
Bruges, Belgium, 22nd to 24th April 2009.
http://www.dice.ucl.ac.be/esann/

6TH International Conference, Information Technology: New Generations (ITNG 2009).
Of interest to members working on various issues in web search including logs etc.
Las Vegas, Nevada, 27th to 29th April 2009.
http://www.itng.info/Taska/WEB_TECHNOLOGIES_TRACK.htm

The Second International Conference on MOBILe Wireless MiddleWARE,Operating Systems, and Applications (MOBILWARE 2009)
General mobile computing conference of interest to members working in the area of mobile search.
Berlin, Germany, 28th to 30th April 2009.
http://www.mobilware.org/

4th International Symposium on Location and Context Awareness (LoCA 2009)
Of interest to members working in the areas of Geographical and/or mobile search.
Tokyo, Japan, 7th to 8th May 2009.
http://loca2009.context-aware.org

The Seventh International Conference on Pervasive Computing (Pervasive 2009)
A conference of interest to members working in the area of mobile search.
Nara, Japan, 11th to 14th May 2009.
http://pervasive2009.org/

The Alberto Mendelzon Workshop on Foundations of Data Management
Of interes to members more interested in the technical aspects of search.
Arequipa, Peru, 12th to 15th May 2009.
http://db.cs.ualberta.ca/amw09/

3rd Int'l AAAI Conference on Weblogs and Social Media (ICWSM 2009)
Of interest to members working in the area of social search.
San Jose, California, 17th to 20th May 2009
http://www.icwsm.org/2009/

The 1st International Conference on Mobile Lightweight Wireless Systems (MOBILIGHT 2009)
General mobile computing conference of interest to members working in the area of mobile search.
Athens, Greece, 18th to 22nd May 2009.
http://www.mobilight.org

The 15th International Conference on Auditory Display (ICAD 2009) and 6th International Symposium on Computer Music Modeling and Retrieval (CMMR2009)
Of interest to members working in the area of music search.
Copenhagen, Denmark, 18th to 20th May 2009.
http://www.icad09.dk/

The Fourth International Conference on Internet and Web Applications and Services (ICIW 2009)
A general web conference of interest to members working in the area of web search.
Venice, Italy, 24th to 28th May 2009.
http://www.iaria.org/conferences2009/ICIW09.html

Libraries in the Digital Age (LIDA 2009)
Of interest to members working in the area of Libraries and search.
Zadar, Croatia 25th to 30th May 2009.
http://www.naaclhlt2009.org

The 2009 IEEE International Symposium on Mining and Web (MAW-09)
A general Data Mining conference of interest to members working in the area of text mining.
Bradford, UK, 26th to 29th May 2009
http://www.enews. ece.uvic. ca/conf/MAW09

Joint Annual Meeting of the North American Association for Computational Linguistics (NAACL) with the Human Language Technology Conference (HLT) – (NAACL HLT 2009)
Of interest to members working in the area of NLP and search.
Boulder, Colorado, U.S.A, 31st May to 5th June 2009.
http://www.naaclhlt2009.org

The First International Conference on Security and Privacy in Mobile Information and Communication Systems (MobiSec 2009)
General mobile security and privacy conference of interest to members working in the area of mobile search particularly in the area of personalisation.
Turin, Italy, 3rd to 5th June 2009,
http://www.mobisec.org/

12th International Conference on Artificial Intelligence & Law (ICAIL 2009)
Of interest to members applying AI and IR techniques to search problems in Law.
Autonomous University of Barcelona, Spain, 8th to 12th June 2009.
http://idt.uab.cat/icail2009/

Joint Conference on Digital Libraries (JCDL 2009)
Of interest to members working in the area of Digital Libraries.
Austin, TX, USA, 15th to 19th June 2009.
http://www.jcdl2009.org/

The 7th conference on Next Generation Information Technologies and Systems (NGITS'2009)
A general information systems conference with a theme on information retrieval and filtering.
Haifa, Israel, 15th to 18th June 2009.
http://mis.haifa.ac.il/~ngits2009

Joint Second International Conference on Mathematics and Computation in Music (MCM 2009) and the John Clough Memorial Conference will take place
Of interest to members working in the area of music retrieval.
New Haven, Connecticut, USA, 19th to 22nd June 2009.
www.mcm2009.info

Biennial Conference of the British Chapter of the International Society for Knowledge Organization (ISKO 2009)
Of interest to members working in the area of classification and knowledge organisation.
London, U.K. 22nd to 23rd June 2009
http://www.iskouk.org/conf2009/index.htm

The 17th International Conference in User Modeling, Adaptation, and Personalization (UMAP 2009)
A conference of interest to members working in area such as personalisation and user models for IR.
Trento, Italy, 22nd to 26th June 2009.
http://umap09.fbk.eu

The Second Information: Interactions and Impact (i3) International Conference
Of interest to members working in the area of context and IR.
Robert Gordon University, Aberdeen, on 22nd to 25th June 2009
http://www.i3conference2009.org.uk

The 7th international Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt'09)
General mobile computing conference of interest to members working in the area of mobile search.
Seoul, Korea, 23rd to 27th June 2009.
http://www.wiopt.org

The Fifteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'09)
A general data mining conference of interest to members working in the area of text mining.
Paris, France, 28th June to 1st July 2009.
http://www.kdd.org/kdd2009/

32nd Annual International ACM SIGIR Conference
The big IR conference, with all themes on the subject of interest to members.
Boston, MA, U.S.A, 19th – 23rd July 2009.
http://sigir2009.org/

Industry News: Endeca unveils Digital Asset Navigator solution

Endeca Technologies, Inc., a search and information access software company, recently unveiled its new Digital Asset Navigator solution, a new search and information access solution designed for media companies, content publishers, advertising agencies and marketing departments, and the latest addition to Endeca's media and publishing-specific information access solutions. Designed to complement and extract latent value from Digital Asset Management (DAM) initiatives, it provides information access capabilities to explore, find and evaluate images, graphics, audio/video files, and other high-value digital assets. As a result, creative, publishing and marketing professionals can reduce production costs, speed time-to-market, avoid compliance issues, and reduce related royalty payments through greater reuse of assets.

The solution is a direct result of new investments designed to deliver new search and information access capabilities, solutions and tools, for media and publishing businesses, announced in late September. In a separate announcement, Endeca reported that Discovery Communications has selected and deployed the Digital Asset Navigator solution.

Legacy vendor-based and homegrown DAM solutions, designed to manage, categorize and store digital assets, and address related workflow and processes, rely on lower-end, commodity search engines and rigid navigational hierarchies, that hinder discovery of digital assets and cloud visibility into the overall collection. In addition, other information relevant to digital asset reuse, including syndication rights, licensing rights, prior usage, related marketing campaigns, advertisers, target demographic etc., are often unavailable or stored in other, siloed source systems.

Specific features and functionality include:

- Market leading search - helps people quickly find items of interest by taking advantage of the latest innovations in information retrieval, including "look ahead" suggestions, spell correction/suggestions, synonyms and compound dimensions
- Guided Navigation capabilities – encourage exploration and discovery by displaying attributes, characteristics and other metadata such as production type, scene type, location, orientation, etc. as dynamic navigation options, letting users browse large collections of digital assets and refine long lists of search results
- Content Spotlighting - promotes relevant related content, low-cost alternatives, tips and more based on a user's profile, search terms and browse path
- Simple integration with popular DAM, ECM, databases and enterprise applications - streamlines application development and ensures a 360-degree view of relevant information by including content adapters for OpenText Artesia, Interwoven, MediaBeacon, ODBC, JDBC
- Native security - ensures permission-based access and record-level security to ensure that users only see what they have the ability to use, reducing legal and compliance costs

The user experience and functional design of the new Digital Asset Navigator solution was created through close collaboration with several of our largest media and publishing clients. Each expressed the same frustration with their digital asset management solutions: lack of visibility into the collective information assets stored within," said Steve Papa, CEO of Endeca. "As a result, we've been able to rapidly create a solution that unlocks the latent value in existing digital asset management investments, while providing visibility across all of the information needed to make reuse and production decisions."

Headquartered in Cambridge, MA, Endeca has operations in North America, Europe and Asia. For more information: endeca.com or info@endeca.com.

Featured Job: Enterprise Search Consultant

Conchango is one of the UK's leading business and technology solution providers and the UK leader in Agile development using Scrum methodology. We are looking for an Enterprise Search Consultant with a passion for technology to help deliver innovative projects and contribute to the wealth of knowledge in our teams and communities. As an Enterprise Search Consultant your role will involve:

- Implementing innovative enterprise search projects using some of the industry's leading enterprise search solutions, such as Microsoft Fast, Endeca, Sinequa and Mercado
- Consulting on Enterprise Search strategy from both a business and technical perspective
- Advising clients on Enterprise Search vendor selection processes
- Developing an understanding of the critical success and failure drivers with enterprise search implementations
- Advising on best practices when designing Enterprise Search solutions
- Identifying possible project or application design decisions that might impact a client's ability to fully realise their enterprise search strategy
- Having a thorough understanding of testing and deploying search in production environment
- Tracking and resolving problems efficiently and effectively
- Providing a consultancy role on search technologies
- Taking an active role in communities (internal interest groups) to spread technical and project knowledge, to help improve best practices and to help identify areas of IP reuse

For further details see: http://www.conchango.com/join-us/?article=search-consultant

Contacts

Web: http://irsg.bcs.org/
Email: irsg@bcs.org.uk
Subscriptions: http://irsg.bcs.org/membership.php
ISSN: 0950-4974

To subscribe, unsubscribe, change email address or contact details please visit http://irsg.bcs.org/ or email irsgmembership@bcs.org.uk.

The IRSG is a specialist group of the British Computer Society.
To automatically receive your own copy of Informer, simply join the IRSG via the IRSG website.